

# Representational Analyticity

JACK C. LYONS

---

**Abstract:** The traditional understanding of analyticity in terms of concept containment is revisited, but with a concept explicitly understood as a certain kind of mental representation and containment being read correspondingly literally. The resulting conception of analyticity avoids much of the vagueness associated with attempts to explicate analyticity in terms of synonymy by moving the locus of discussion from the philosophy of language to the philosophy of mind. The account provided here illustrates some interesting features of representations and explains, at least in part, the special epistemic status of analytic judgments.

## 1. Introduction

Although Kant may have been the first philosopher to explicitly define analyticity in terms of concept containment (Kant, 1929, B10), he was summarizing a long tradition in doing so. It is a varied tradition, however, for different authors understood the term ‘concept’ differently, and ‘containment’ must be understood differently as a result. The sense in which one Platonic Form might contain another, for example, is likely to be quite different from the sense in which one Lockean idea might contain another.

Similarly, the apsychologism dominant throughout most of twentieth century Anglo-American philosophy mandated that concepts are taken to be word meanings, Fregean senses, or some other abstract objects, and the notion of containment had to be correspondingly and radically reenvisioned. According to a standard post-Fregean view, a proposition (more often: statement) is analytic just in case it is a logical truth or could be transformed into such by substitution of synonymous expressions, where two expressions are synonymous if and only if they have the same definition, alternatively, if and only if they express the same (Fregean) sense. Talk about concept containment has often been retained in this apsychologistic tradition but it cannot be read literally. The concept *unmarried* would be contained in the concept *bachelor*, in Frege’s words, ‘as plants are contained in their seeds, not as beams are contained in a house’ (Frege, 1953, p. 101). Obviously containment is no longer doing theoretical work here; it is synonymy that bears the explanatory burden.

---

Thanks to Chris Hill and two anonymous referees at this journal for helpful comments on drafts of this paper. Earlier versions were presented to the philosophy departments at Florida State University and the University of Arkansas; thanks also to those audiences.

**Address for correspondence:** Department of Philosophy, University of Arkansas, Fayetteville, AR 72701.

**Email:** jclyons@uark.edu

Famously, however, the arguments of Quine and others (e.g., Quine, 1951, 1960; White, 1950) have brought this standard account into serious disrepute. There seems to be no non-circular way of making sense of the notions of definition, sense, and synonymy, in such a way that they can be used to explicate the notion of analyticity. Since the present concerns are neither historical nor with the standard post-Fregean view, I will not discuss these arguments but will simply suppose that any attempt to understand analyticity in terms of sense, definition, or synonymy is doomed to failure. These arguments, however, are aimed at the standard post-Fregean view, and they leave the older concept-containment understanding of analyticity intact.

In 'Two Dogmas of Empiricism' Quine presupposes *psychologism*, dismissing the Kantian concept-containment account of analyticity as too vague and metaphorical to even bother refuting. Concept containment, however, need not be mysterious or metaphorical. If we just take concepts to be psychological entities, in particular, those mental representations the tokening of which is partially constitutive of the tokening of an occurrent belief, the notion of containment becomes tractably clear, and a structuralist theory of analyticity becomes possible. Concept containment becomes a species of representation containment, and there is nothing particularly metaphorical or mysterious about this. Frege's influential complaint with concept containment was mainly that it allowed for too few analytic truths. Viewing concept containment as a species of representation containment makes it possible to reply to this objection as well, at least in part, though it will take the whole of the present paper to show this.

My goal here is to resurrect the old understanding of analyticity as concept-containment. My own personal concerns with analyticity involve epistemology and the philosophy of mind primarily; this makes it appropriate to 'psychologize', concentrating on analytic judgments rather than analytic statements. The consequences of this choice of target are far-reaching, and the resulting notion of analyticity will look very different from the post-Fregean notion. I cannot emphasize enough that the account provided here is *supposed* to look very different; it is not intended to have any immediate consequences for linguistics or the philosophy of language, absent certain substantive claims (which I will not pursue here) about the relation between thought and language.<sup>1</sup>

The present approach departs from that of the post-Fregean tradition in three crucial respects. First, I will take as a starting point Kant's first formulation of

---

<sup>1</sup> Not everyone seeks an analytic/synthetic distinction for the same sorts of reasons. Jerrold Katz, for example, is quite explicit about taking analyticity to be a linguistic rather than an epistemological phenomenon: 'the epistemological significance of the [analytic/synthetic] distinction is nothing like what it is commonly taken to be' (Katz, 1992, p. 11). Katz ends up with a theory of analyticity according to which some analytic statements are only contingently true (Katz, 1998), since analyticity has only to do with sense properties and sense doesn't determine reference. For Katz, although it is analytic that all cats are animals, it is possible that all cats are (not animals, but) Martian robots. Katz's general project is not the one pursued here.

analyticity (1929, B10) in terms of concept containment, rather than his second formulation (1929, B12) in terms of non-contradiction. Second, I will take the target to be analytic judgments, or beliefs, rather than analytic statements. Finally, I will take concepts to be mental entities rather than abstract objects; concept containment will be understood accordingly. The resulting conception of analyticity will be more modest, but also I hope more tenable, than the standard post-Fregean conception. The present account of analyticity illuminates some interesting features of representations—features that have generally gone unnoticed in discussions of representation—even apart from the epistemological consequences of the analytic/synthetic distinction. Whether, in the end, a structuralist theory of analyticity can accommodate all of the truths we want to count as analytic will have to remain an open question. But I do hope I will have shown that the notion of concept containment can be made fairly clear and that it is far more powerful than one might have thought.

I present a two-part account of representational analyticity. In section 2, I present a rough but intuitive account of concept containment, to illustrate the basic ideas and to show the range of applicability in different kinds of representational systems. The first approximation presented there suffices to suggest the epistemological consequences of the general view and begin the long argument for the claim that concept containment supports more analyticities than Frege thought. There are, however, counterexamples to the first approximation, the elimination of which demands a more rigorous specification of much of what the intuitive account takes for granted. In section 3, I give that more rigorous account. The kind of containment at issue is, in the end, neither so simple as that of beams in a house nor so mysterious as that of trees in seeds. Once this structuralist conception of analyticity has been sufficiently articulated, I will return, in section 4, to some of the epistemological implications of analyticity thus construed.

### 1.1 Terminological and Methodological Matters

Although I want to eschew such notions as sense, definition, and synonymy, I will be helping myself to the notions of representation, extension, reference, necessity, and for the sake of convenience, propositions. By ‘representation’ I simply mean a physical state or object with content, i.e., an extension or truth conditions. I will be using the term ‘extension’ a bit more broadly than is sometimes done; I mean to include all the relevant non-actual, as well as actual, objects in the extension of a representation.<sup>2</sup> This is partly the result of a semantic intuition (I think that non-actual chairs are chairs) but also because I will need a relatively fine-grained notion.

---

<sup>2</sup> Thus ‘extension’ will be used here in roughly the way that Lewis (1986) uses ‘property’. The main reason for not simply proceeding in terms of Lewisian properties is that extensions are extensions *of* some representation, and properties are not connected to representations in the right sort of way.

I will assume, as do many others, that there is nothing problematic about appealing to semantic concepts without already assuming an analytic/synthetic distinction.<sup>3</sup> Many, if not most, of the authors working on theories of content in the philosophy of mind explicitly deny the analytic/synthetic distinction (e.g. Harman, 1982; Fodor and LePore, 1992). I will not be defending or assuming any particular theory of content, though certain theories are more naturally suited to certain kinds of representation than others (and vice versa). There seems to be nothing circular about taking the content of a particular representation as read, and this is standard operating procedure in philosophy and psychology.

Similarly, it seems quite possible to understand a distinction between necessary and contingent truths without presupposing an analytic/synthetic distinction. It is important that propositions are not understood here as the *senses* of declarative sentences, although they may be understood along the lines of Lewis (1986) as structured constructions out of possibilities. (I will want to distinguish various necessary truths from each other, so it is not enough to view propositions as mere unordered *sets* of possible worlds.) As my present concerns are not with metaphysics, I will assume—or at least hope—that the metaphysical assumptions just invoked serve only the purpose of expository convenience and that the major claims of this paper could be made without them.

Finally, I should reiterate that I will be using the term ‘concept’ in the older philosophical usage, rather than the contemporary cognitive scientific usage, to denote those mental representations the tokening of which is constitutively involved in the tokening of an occurrent belief or judgment. In the contemporary concept literature in cognitive science, the term is usually used to pick out underlying knowledge structures that are causally—rather than constitutively—implied in the formation of judgments.<sup>4</sup> Consequently, the major theories offered in the concept literature (e.g. prototype theory, exemplar theory, theory-theory) may not exhaust the legitimate possibilities concerning what I am here calling concepts.

## 2. Concept Containment as Representation Containment

Kant had two notoriously distinct characterizations of the analytic/synthetic distinction: one in terms of concept containment, and one in terms of non-contradiction. The basic idea behind the concept-containment type of analyticity, of course, is that ‘ $(\forall x)(Fx \supset Gx)$ ’ is analytic if and only if the concept of *F* contains

<sup>3</sup> Sometimes it is claimed that the rejection of analyticity carries with it the rejection of meaning altogether. This, however, seems a bit strong, especially if we take a theory of meaning to be merely a theory of reference. The term ‘meaning’ is ambiguous, but when it is used for something less than the combination of sense and reference, it is usually used to denote sense as opposed to reference. I have no objection to the claim that the failure of the standard accounts of analyticity render the very notion of a Fregean sense suspect.

<sup>4</sup> A possible exception is Fodor (1981, 1998), who may be using the term in the way I do here. This would explain why the resulting theory is so unsatisfying from the standpoint of someone who seeks an explanation for how it is that we manage to categorize objects.

the concept of G. ‘All bachelors are unmarried’ and ‘ $(p \ \& \ q) \supset p$ ’ are supposed to be analytic on both accounts. ‘ $p \supset (p \vee q)$ ’ is analytic on the second account but not on the first, unless ‘concept containment’ is read in a highly metaphorical sense, like the sense in which the seed contains the tree (cf. Katz, 1992). Again, how we understand containment will depend on how we understand concepts. I propose that we take the terms ‘concept’ and ‘containment’ at face value. Thinking of concepts in the ordinary, psychological, sense will make it far easier to see what containment would amount to.

This, I think, is actually the understanding of analyticity at work in Locke. For Locke, a concept is simply a certain kind of idea, and an idea is simply a certain kind of representation. Concept containment, then, is merely a certain kind of representation containment, and there is nothing very mysterious about representation containment. This will require quite a bit of elaboration, but the basic idea is clear enough. A rough sketch of Locke’s theory—or a suitably anachronistic version of it—is worth reviewing for the sake of illustration.

## 2.1 Locke on Concepts

According to Locke (1975), experience provides us with a stock of primitive, or simple, ideas; they are simple in the sense that they are not and cannot be constructed out of other ideas. Complex ideas, on the other hand, are constructed, presumably recursively, out of simple(r) ideas. The method of composition that Locke has in mind seems to be nothing more than mere concatenation of the simple ideas. As such, complex ideas contain their constituent ideas in a very real and literal sense, in the way that a set contains its elements. To take one of Locke’s own examples, the idea typically associated with the English word ‘gold’ is a complex representation built out of simpler constituents like the idea of yellow, the idea of fusibility, etc. I’ll write this as follows:

$$[\text{YELLOW}] \wedge [\text{FUSIBLE}] \wedge [\text{METAL}] \wedge [\text{SOLUBLE IN AQUA REGIA}] \wedge [\text{FIXED}] \wedge [\text{MALLEABLE}].^5$$

For Kripkean and other reasons, the idea under consideration has a different extension from that of the English word ‘gold’, but the theory in question is one about thought, not about language. To avoid confusion, I’ll use ‘gold\*’ to refer to all and only yellow, fusible, malleable, fixed metals that are soluble in *aqua regia*.

Locke’s theory is thought by many (e.g. Rosch, 1975; Smith and Medin, 1981; though see also Armstrong *et al.* 1983), to be an empirically implausible account of

---

<sup>5</sup> Of the list of component ideas, [YELLOW] is probably the only one that isn’t simply shorthand for another such concatenation of simple ideas, and whether Locke’s concept empiricism leaves him the resources to decompose each of the others into simple ideas is an issue that need not concern us here. Concept empiricism is a theory about the *origins* of simple ideas, while the view under consideration is a theory about the *structure* of complex ideas.

the structure of our concepts. Again, I am not using ‘concept’ in that sense, and I argue below that Lockean-type representations are alive and well in cognitive science. For the time being, however, let us imagine, for the sake of illustration, that our concepts *are* just simple ideas or concatenations of them. Then it seems clear enough that my judgment that gold\* is yellow is an analytic judgment. The concept:

[YELLOW] ^ [FUSIBLE] ^ [METAL] ^ [SOLUBLE IN *AQUA REGIA*] ^  
[FIXED] ^ [MALLEABLE]

contains the concept:

[YELLOW]

in a very straightforward sense, and thus, given the way the Lockean representational scheme works, it is necessarily true that anything in the extension of

[YELLOW] ^ [FUSIBLE] ^ [METAL] ^ [SOLUBLE IN *AQUA REGIA*] ^  
[FIXED] ^ [MALLEABLE]

is in the extension of:

[YELLOW].

The judgment that gold\* is yellow is trivial, in the sense of not being very informative about the world, but also in the sense of being extremely easy to get right. Extracting constituents from a complex representation is an entirely syntactic, and therefore mechanical, task. These are the properties that we wanted out of analytic judgments, and there is clearly some kind of analyticity involved in this sort of judgment. Since it may not be the recently conventional, post-Fregean, sort of analyticity, I’ll call it ‘representational analyticity’: that kind of analyticity that holds in virtue of representation containment. To better understand representational analyticity, we merely need to get clear on the notion of representation containment.

The analytic/synthetic distinction was traditionally supposed to have epistemological implications; analytic inferences were supposed to be so epistemically secure precisely because of their triviality. The present notion of representational analyticity can account for this more easily than the standard post-Fregean account. Whether ‘rational animal’ is synonymous with—or even coextensive with—‘man’, ‘human being’, or ‘person’ is very hard to know. It is very easy, however, to know that anything in the extension of ‘rational animal’ is in the extension of ‘rational’, even if it is unclear what the extension of ‘rational’ is. Similarly, one need not

know anything about language to know that anything in the extension of:

$$[\text{YELLOW}] \wedge [\text{FUSIBLE}] \wedge [\text{METAL}] \wedge [\text{SOLUBLE IN AQUA REGIA}] \wedge [\text{FIXED}] \wedge [\text{MALLEABLE}]$$

is in the extension of:

$$[\text{YELLOW}].$$

This we can know merely by knowing how the Lockean scheme works, i.e. by knowing what the function is that assigns extensions to complex representations on the basis of their constituents.

I want to reiterate how different a theory of analytic judgments is from a theory of analytic sentences. It is easy to view Locke's theory as making a claim about the sense of the word 'gold'. Although I doubt that this is what Locke had in mind, it doesn't really matter if he did. The account would do just as well without any mention of sense or of the *word* 'gold'. We would have concepts and analytic relations among them even if we didn't have a language (extreme Whorfians notwithstanding); we just wouldn't be able to talk about them. It is possible that ordinary English (i.e. without words like 'gold\*' or the special orthography introduced above) is incapable of accurately translating any of the judgments formed in the Lockean scheme; this however, would be a fact about (verbal) belief *attribution*, not a fact about the metaphysical or epistemic status of the *judgments* in question.

## 2.2 How to Think About Representational Analyticity

Claims about representational analyticity must be relativized to a particular representational scheme.<sup>6</sup> Representation containment is largely a syntactic issue, and different representational schemes will have different syntaxes. What is represented by a complex representation in one scheme may be represented by a primitive representation in another scheme, and consequently, propositions that are representationally analytic with respect to the first scheme may not be representationally analytic with respect to the second.

The analyticity of the judgment that all gold\* is yellow, for example, depends on certain contingent features of the Lockean scheme. In particular, it depends on the fact that the only way to represent gold\* in this scheme is via a complex representation of yellow, fusibility, and the like. If the Lockean scheme were to include a syntactically primitive representation of gold\* as well, there would be two ways to form the judgment that gold\* is yellow: one by tokening the primitive gold\* representation, and one by tokening the complex gold\* representation.<sup>7</sup> Only the

<sup>6</sup> A representational scheme is, roughly, a system of representations along with the rules governing them. A more precise definition is given below in section 3.1.

<sup>7</sup> I am assuming for the moment that judgments are individuated by their propositional content and that these two judgment tokens would have the same propositional content. This makes my task harder but results in a more general account of representational analyticity. I consider the possibility that judgments are more finely individuated in section 3.4 below.

former, however, would involve representation containment and consequently the mechanical extraction of constituents.

We can suggest, as a very first approximation to the notion of representational analyticity, the following:

- (A1): The proposition  $\lceil (\forall x)(Fx \supset Gx) \rceil$  is representationally analytic with respect to representational scheme *S* iff every representation in *S* of something's being *F* includes as a part (though not necessarily a *proper* part) a representation of its being *G*.

((A1) is in serious need of clarification and even modification, but it should serve for the time being to fix ideas for the ensuing discussion.) Representational analyticity is defined here in terms of parts, rather than proper parts, in order to include the class of what Locke calls 'identical propositions', where something is predicated of itself (e.g., 'gold is gold'; 'all horses are horses').

It is important for some purposes to have a non-relative understanding of representational analyticity. This is easily accomplished:

- (NR): *p* is representationally analytic *simpliciter* iff there is some representational scheme with respect to which *p* is representationally analytic.

Representationally synthetic propositions (*simpliciter* or with respect to some scheme) are just those that are not representationally analytic.

I am saying that something is representationally analytic *with respect to* a given representational scheme, rather than *in* that scheme. There is a reason for this: I want to explore the general notion of representation containment, and to that end, it is important to consider even representational schemes that are not capable of expressing propositions. The Lockean scheme, for example, might not contain the representation:

[IS].

This would not affect the relation between:

[YELLOW] ^ [FUSIBLE] ^ [METAL] ^ [SOLUBLE IN AQUA REGIA] ^  
[FIXED] ^ [MALLEABLE]

and:

[YELLOW],

although it would affect the scheme's ability to express that relation. Thus, I will say that the Lockean scheme *supports* certain analyticities, even if it is not capable of *expressing* them.

It is worth taking a moment to show how a theory of mind might make room for beliefs and judgments, even when the representational scheme in question is incapable of expressing propositions. According to Locke, knowledge is ‘nothing but the perception of the connexion and agreement, or disagreement and repugnancy of any of our Ideas’ (1975, IV.1.2). This ‘perception’, however, need not be a third idea. Instead, it could be some kind of relation to the relevant ideas.

Consider a standard sentential functionalism (e.g. Fodor, 1975, 1985), according to which occurrent belief is a certain functionally defined two-place relation between a cognizer and a propositional representation—a sentence in *Mentalese*. Ramsey-Lewis functionalism (Lewis, 1970), on the other hand, takes occurrent belief to be a one-place property of the cognizer. These two theories offer two options: build all the belief content into the representation, or build it all into the functionally defined relation. A third alternative is to build some of the belief content into the representation and some of it into the relation. What we might call ‘Lockean functionalism’ would view occurrent belief as a three-or-more-place functionally defined relation among a cognizer and two or more mental representations.<sup>8</sup>

Although coherent, Ramsey-Lewis functionalism suffers some problems that Lockean functionalism can avoid. Being a non-representationalist theory, Ramsey-Lewis functionalism is at odds with contemporary psychological theorizing, which depends heavily on positing representations. Nor can it explain any of the observed systematicities of thought or inference. Certain cognitive capacities tend to co-occur; anyone who is capable of entertaining the thought that John loves Mary is capable of entertaining the thought that Mary loves John (Fodor and Pylyshyn, 1988). Certain representationalist theories explain this by claiming that we possess mental representations that can be put into various constructions with each other; on a non-representationalist view, this phenomenon is an utter mystery. Depending on which representations are posited, and which systematicities are in fact discovered, Lockean functionalism can avoid these difficulties.<sup>9</sup>

Although Lockean functionalism is an idea very much worth pursuing, all that matters for the present purposes is that it is left as a possible option. Though I will be committing myself to some form of representationalism in

---

<sup>8</sup> Strictly speaking, all that’s necessary is a relation between a cognizer and one or more representations. The belief that gold is yellow might involve a relation to the gold representation and the yellow representation, while the belief that gold is gold would involve a very similar relation to just one representation: the gold representation. The formulation in the text, however, better illustrates the difference between this sort of view and sentential functionalism. Obviously, I am not attributing functionalism to Locke himself.


<sup>9</sup> Problems of systematicity can, of course, also arise for representationalist theories—witness the debate over connectionism. However, a generic sort of Lockean functionalism would have the resources to handle these problems, provided that it posits the right sorts of representations.

what follows, I will not assume that all the representations involved in belief are propositional representations. The mere possibility of Lockean functionalism shows that it is at least possible to have propositional attitudes without having propositional representations.<sup>10</sup> By taking seriously the possibility of Lockean functionalism, we greatly increase the number of representational schemes that might support representational analyticities, for many important schemes, including standard connectionist schemes, do not contain propositional representations.


**2.3 Examples in Other Schemes**

Since I want to explore the general notion of representational analyticity, it will be helpful to look at a few different kinds of representational schemes and the representational analyticities they support. Any scheme that contains complex representations will support some representational analyticities or other.

For example, in a Venn diagram scheme, the representation of the proposition ‘nothing is A’, as shown in Figure 1, contains a representation of ‘no B is A’ (Figure 2). Therefore, it is representationally analytic with respect to this scheme that if nothing is A, then no B is A.



**Figure 1** Venn diagram: ‘nothing is A’



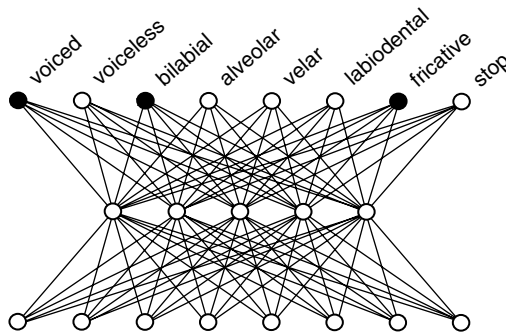
**Figure 2** Venn diagram: ‘no B is A’

<sup>10</sup> For a different way of getting the same result, see Cummins (1992, 1996).

It is representationally analytic with respect to English that all yellow dogs are yellow, and it is representationally analytic with respect to the propositional calculus that ' $p \ \& \ q$ ' implies  $q$ .

A more important and interesting kind of representational scheme is that employed by some connectionist networks. A *microfeatural scheme*, as I will call it, is a kind of distributed representational scheme, where the activation of individual nodes represents certain low-level features (microfeatures), and activation patterns of whole layers of nodes represent sets of those microfeatures. NETtalk (Sejnowski and Rosenberg, 1987), for example, is a network that takes strings of text as input and maps them to the appropriate phonemes. The phonemes are represented at the output layer as a distributed representation across phonetic features. A simplified portion of the network, tokening a representation of a voiced bilabial fricative, is illustrated in Figure 3. It is representationally analytic with respect to the NETtalk scheme that all voiceless bilabial stops are bilabials.<sup>11</sup>

I claimed above that Lockean schemes are alive and well in today's cognitive science. This is because these connectionist schemes just *are* Lockean schemes of a certain kind. The primitives are perhaps a bit 'smaller'—more 'micro'—than the ones Locke himself envisioned, but this is the only difference. Both make complex representations out of mere concatenations of primitives. Neither is capable of



**Figure 3** Portion of network representing voiced bilabial fricative

<sup>11</sup> As with the gold example above, I have not claimed that it is representationally analytic with respect to this scheme that all [p]s are bilabials. Whether it is even true—let alone analytic—that all [p]s are bilabials depends on what a bilabial is; if a bilabial is anything that has some particular kind of *sound*, then it is true; if, however, a bilabial must be something that results from actually pressing two lips together, then it is not true, for [p]s can be electronically generated, and these would not be bilabials. It is tempting to describe a certain output pattern as a [p] representation, but this is only an accurate description if a [p] is the same thing as a voiceless bilabial stop. One must consult a phonologist to determine whether all [p]s are bilabials, but not to determine whether all voiceless bilabial stops are bilabials.

representing propositions, though Lockean functionalism may render this deficit innocuous. In any case, microfeatural schemes are taken quite seriously in contemporary cognitive science, so Lockean schemes are doing well, at least outside of the concept literature.<sup>12</sup>

There are also at least some representationally analytic truths about mathematics. In his famous argument against the analyticity of mathematics, Kant argues that the concepts of seven and five are not contained in the concept of twelve. Whether this is true depends on the structure of the concept of twelve. Consider a slightly modified hash-mark scheme. The one I have in mind is just like the normal one, except that here we don't make every fifth slash a diagonal one. Every number is represented by that many vertical slashes. In such a scheme, the representation of seven:



contains a representation of four and a representation of three; thus it is representationally analytic with respect to such a scheme that three plus four equals seven (and that seven minus three equals four, etc.). All the truths of addition and subtraction of the natural numbers (as well as the laws of associativity and commutativity) will be representationally analytic with respect to the hash-mark scheme. In Roman numerals, matters are more complicated. It is representationally analytic with respect to Roman numerals that one plus two equals three, but not that two plus three equals five ('III' contains 'I' and 'II', though 'V' does not contain 'II' or 'III').

Several different schemes will be required to capture the truths of multiplication and division for the natural numbers. The way the decimal system handles multiples of ten gives us almost what we are looking for. Consider a 'pseudo-decimal' scheme, in which:

- '0' represents the number 10 (rather than 0)
- '1' represents the number 1
- '2' represents the number 2
- etc. . . .
- '9' represents the number 9
- when any numeral  $n$  is placed directly to the left of '0', the resulting numeral represents the quantity  $10 \times n$ .

' $9 \times 10 = 90$ ' and ' $90 \div 9 = 10$ ' are representationally analytic with respect to that scheme. It is clear that another scheme will effect the same results for multiples of eight, and so forth for the rest of the natural numbers.

---

<sup>12</sup> One reason for the differential treatment is that Locke's theory seems to be attempting to do what the microfeatural schemes never purported to do, like provide word meanings, definitions, and the like.

A pictorial scheme is one where the representation represents in virtue of an isomorphism between it and its representanda; photographs, dioramas, sketches, paintings, scale models, and the like are pictorial representations in this sense.<sup>13</sup> A quasi-pictorial scheme is one that includes pictorial as well as non-pictorial elements, e.g. wiring diagrams, maps, political cartoons. Pictorial and quasi-pictorial schemes support various representational analyticities in virtue of their part-whole relations: a picture of a circle and a triangle contains a picture of a triangle. But the more interesting kind of analyticity supported by such schemes involves the 'identical propositions'. In a standard map scheme, a representation of Borneo being (roughly) to the east of Sumatra contains as a part (though not a proper part) a representation of Sumatra's being (roughly) to the west of Borneo. Thus it is representationally analytic with respect to map schemes that if  $x$  is to the east of  $y$ , then  $y$  is to the west of  $x$ , and that if  $x$  is to the north of  $y$ , then  $y$  is to the south of  $x$ . Similarly, it is representationally analytic with respect to pictorial schemes that if  $x$  is to the right of  $y$ , then  $y$  is to the left of  $x$ . None of these is representationally analytic with respect to English.

The preceding examples illustrate something of the power and the ubiquity of representational analyticity. There are a great many representational schemes, each supporting its own analyticities. The Fregean worry that containment won't generate enough analyticity is at best in need of argument, at worst, misplaced. Representational analyticities are supported by a wide range of representational schemes, though different schemes will support different analyticities. This is why the non-relative definition of representational analyticity from section 2.2 is important: to capture all the representational analyticities there are, we will have to appeal to more than one scheme.

### 3. A More Refined Account

Though the foregoing offers a hint of what representational analyticity can do, the notion itself is in need of clarification and reformulation. I mentioned in section 2.2 that there were some problems with the first approximation:

---

<sup>13</sup> It is possible, of course, to *use* pictures in such a way that they represent something to which they are not at all isomorphic. A spy might put out a photograph of his grandmother in order to indicate to one of the dinner guests that the goods have been delivered. I will ignore such cases in what follows, and will always have in mind pictorial representations *qua* pictorial representations, i.e. those that represent in virtue of an isomorphism. It is fine with me that a particular representation token can fall into more than one scheme at a given time.

John Kulvicki pointed out to me in a recent conversation that many representational schemes embody an isomorphism between the representations and the representanda. Microfeatural and traditional Lockean schemes serve as good examples. However, the difference between these and pictorial representations is that the latter represent *in virtue of* this isomorphism; the former do not. Because the primitives of a microfeatural scheme have their extensions independently of the other primitives, the isomorphism makes no independent contribution to the semantics of the complex representation, though it is certainly important for other reasons.

- (A1):  $\lceil (x)(Fx \supset Gx) \rceil$  is representationally analytic with respect to representational scheme  $S$  if every representation in  $S$  of something's being  $F$  includes as a part (though not necessarily a *proper* part) a representation of its being  $G$ .

Although (A1) seems to give us the right answers concerning the examples from section 2.3, it is inadequate; this is because it proceeds in terms of parts, rather than in terms of constituents. As a result, it allows too many propositions to count as representationally analytic.

The English word 'dog' (or at least that string of letters) is a part of the word 'dogmatic'. As such, it follows from (A1) that it is representationally analytic with respect to English that only dogs are dogmatic. This, of course, is unacceptable. A different but related problem is that 'non-dog' contains the word 'dog'. Thus, (A1) also has the consequence that it is representationally analytic that all non-dogs are dogs.

The problems here are due to the fact that 'dog' is a part of, but not a constituent of, 'dogmatic', and although it is a constituent of 'non-dog', it is not the right kind of constituent. An adequate formulation of representational analyticity will have to be sensitive to constituents of various kinds. It is first necessary to clarify the notion of a representational scheme.

### 3.1 Representational Schemes and Constituency

A *representational scheme* is a set of representations, together with a syntax and semantics governing them. A *scheme specification* for a given representational scheme will specify:

- (a) all primitives of the scheme; this may include both genuinely representational primitives (symbols that have extensions) and logico-syntactic operators,
- (b) a set (possibly null) of syntactic combination rules, describing which combinations of primitives count as well-formed representations,
- (c) an extension for each representational primitive (for most schemes this would have to be effected indirectly, by means of a general theory of content for the scheme), and
- (d) a set (possibly null, corresponding to the one in (b)) of semantic composition rules, which will provide an extension for each complex representation as a function of the operators (if any) and the extensions of the representational primitives.

A scheme specification need not accomplish all these separately, so long as it accomplishes all of them. Together, (a) and (b) constitute the syntax of the scheme, (c) and (d) the semantics.<sup>14</sup>

<sup>14</sup> For an illustration of the applicability of the notions of syntax and semantics to non-language-like representational schemes, see Shin, 1994.

To illustrate, consider the somewhat unusual case of pictorial schemes and what would be needed to specify one. Different kinds of pictorial schemes will contain different primitives. Typically the primitives are merely pixels (or something like them in the case of analog pictorial representations), though we could also make pictorial schemes where all the pictures need to be made out of a certain stock of basic shapes. For example, the primitives of the Tinker-toy kits used in chemistry classes to make models of molecules are wooden balls of various colors, with sticks and springs for constructing more complicated shapes out of them. The rules for well-formed representations are more salient in such schemes. A single ball is a well-formed representation in the scheme, but a single stick isn't. Two balls joined by a stick is a well-formed representation in the scheme, but two balls and a stick, all lying next to each other on the table isn't.<sup>15</sup>

A standard story concerning the semantics of pictorial schemes is that a representation  $r$  represents some object  $o$  if and only if  $r$  instantiates some relational structure  $\langle A, R_1, \dots R_n \rangle$ ,  $o$  instantiates some other relational structure  $\langle B, S_1, \dots S_n \rangle$ , and these two relational structures are isomorphic under some function  $f$  (something like this is found in Cummins, 1992, 1995).<sup>16</sup> However, this is better viewed as a constraint on schemes qua pictorial than as a full scheme specification; a pure isomorphism view is not likely to work for all pictorial schemes. The Tinker-toy example above shows that well-formedness is a matter of more than just instantiating some relational structure or other. Also, representations notoriously tend to instantiate a number of different relational structures; consequently, too

<sup>15</sup> Pictorial schemes resemble traditional Lockean and microfeatural schemes in at least the following respects. None of these schemes contains operators. All of them, however, are capable of representing conjunctions (e.g. a square and a triangle) even without an 'AND' element. 'AND' is not a primitive in the scheme; rather it's built into the composition rules that the concatenation of an  $A$ -representation and a  $B$ -representation is an  $A$ -and- $B$ -representation. Finally, due to the absence of operators, none of these schemes has any syntactic complexity to speak of.

<sup>16</sup> A relational structure is an ordered  $n + 1$ -tuple consisting of a set of objects and  $n$  relations, each of which has that set as its field. For a physical object to instantiate a relational structure in the present sense is simply for that structure to accurately (though perhaps not exhaustively) describe that object's physical structure. Isomorphism is roughly just a matter of there being a mapping from every element of one relational structure to a unique corresponding element of the other. More precisely, a relational structure  $\langle A, R_1, \dots R_n \rangle$  is isomorphic under a function  $f$  to another relational structure  $\langle B, S_1, \dots S_n \rangle$  whenever  $f$  maps  $A$  one-one onto  $B$  and for every  $n$ -tuple of elements in  $A$ ,  $\langle x_1, \dots x_n \rangle \in R$  iff  $\langle f(x_1), \dots f(x_n) \rangle \in S$ .

Swoyer (1991) takes many different sorts of representational schemes to represent in virtue of isomorphism, and claims (p. 493) that there is therefore nothing special about pictorial representations. There is an important difference, however, between representation in virtue of an isomorphism between the representational *scheme* and the representandum domain on the one hand, and representation in virtue of an isomorphism between the individual *representation* and its representandum on the other. The term 'analog' is unfortunately sometimes used to describe representational schemes of the first sort, sometimes of the second sort. Pictorial schemes are unique in that they exhibit both kinds of isomorphism.

many things will be isomorphic to too many other things. This last problem can be solved by including more detail in the scheme specification. The representational scheme will substantially pare down the number of isomorphisms by specifying which features of the representation are representational features.

Intuitively, not all of the features of a given representation are representational features. The fact that the Tinker-toy model parts are wooden does not mean that the models represent water molecules as having wooden components. In a red monochromatic scheme, we might have a red picture of a cube (because everything in the picture is some shade of red), but this will not represent the cube as being red, because color is not a representational feature of monochromatic schemes (Haugeland, 1990).<sup>17</sup> If, for example, the colors in the picture don't have representational significance, then the scheme will require that the *relevant* relational structures instantiated by the picture are only those where the color of the picture is not specified by the first element of the relational structure. Similarly, in false-color satellite images, color does represent, but it doesn't represent *color*; instead it represents some kind of geological or geographical feature.<sup>18</sup> Thus the scheme ought to provide additional constraints on the relevant functions under which two structures are isomorphic. If only certain instantiated structures count, and only certain isomorphisms among these count, it is far more likely that we will get univocal content out of pictures.

This bears on an objection that will have occurred to some readers: suppose I have a scheme consisting of a number of pictures, and all the cat pictures just happen to have the cats sitting. Is it therefore analytic with respect to that scheme that all cats are sitting? It depends on whether the position of the cats in the pictures is a representational feature of the pictures. The mere fact that the cats are sitting in the pictures does not indicate that the pictures represent the cats *as* sitting. If the pictures don't (all) represent the cats as sitting, then there is no counter-example to (A1). If the pictures do (all) represent the cats as sitting, then what is analytic in that scheme is not that all cats are sitting, but that all *sitting* cats are sitting, which presumably ought to count as analytic.

One might object that I am taking too much for granted here, that one cannot determine the content of a representation without a specification of the downstream processes that put that representation to use.<sup>19</sup> Certainly a proponent of use theories will attribute the additional (extra-isomorphism) constraints on pictorial schemes to downstream processes, but I make no such assumption here. The present purpose is not to give a theory of content or to say what makes certain features representational features. The purpose is to explore the consequences of a representation's having a certain extension and a certain structure. Taking such

---

<sup>17</sup> This is, of course, a general phenomenon, and not one restricted to pictorial schemes. In English and propositional logic, font doesn't have representational significance, though we could easily create variants in which it does.

<sup>18</sup> Such images would therefore fall under quasi-pictorial schemes, with color playing the same sort of role as words on a map.

<sup>19</sup> This objection was raised by a referee for this journal.

notions for granted, the content of a representation is as much a matter for stipulation as its structure.

In (purely) pictorial schemes, the semantics is typically given all at once; the specification of the relevant structures and isomorphisms applies generally, to the primitives as well as to the complex representations. In more language-like schemes, where the mapping from primitives to their representanda is arbitrary, these will have to be handled separately. Thus, some theories of content (e.g. Fodor, 1987) explicitly aim only at (c), leaving (d) for some other occasion.

Clause (d), though a fairly standard characterization of compositionality (see, e.g. Fodor and Pylyshyn, 1988), is false; it can't distinguish idiomatic from genuinely compositional expressions. Consider natural languages. For any idiomatic expression which consists of finitely many words and which has a determinate extension, there will always be *some* mapping or other from the extension of those words to the extension of the idiom. Appeal to functions here is vacuous; we'll need a more constrained understanding of constituency.

Roughly and intuitively, the difference between idiomatic expressions (where the parts are merely parts) and compositional expressions (where the parts are actual constituents) is that in idiomatic expressions, the meaning of the whole has nothing to do with the meanings of the parts. I'll say that  $R^*$  is a *constituent* of  $R$  in scheme  $S$  iff:

- (i)  $R^*$  is a part of  $R$ , and
- (ii) holding the composition rules constant, if the meaning in  $S$  of  $R^*$  were to change, the meaning in  $S$  of  $R$  would also change.<sup>20</sup>

If the word 'dog' were to change in meaning, 'non-dog' would change correspondingly. However, 'dogmatic' would (or at least might) continue to mean what it does. So 'dog' is a constituent of 'non-dog' but not of 'dogmatic'. Or so it is with English, anyhow. We could easily create a language in which 'dog' is a constituent of 'dogmatic'; for that reason, claims about constituency have to be relativized to representational schemes.

This notion of constituency applies to nonlinguistic schemes as well. Recall the scheme that Paul Revere and his friend reportedly used to signal when the British were coming. His friend would hang one lantern from the church tower if the

---

<sup>20</sup> It follows from this definition of constituency that every representation is a constituent of itself, since if its meaning were to change, its meaning would change. This is just what is needed to handle identical propositions. This definition also allows constituents to be recursively embedded in larger constituents.

Since only some parts of a given representation count as constituents, the extraction of constituents from a given representation, though still mechanical, may not always be so trivial as suggested in section 2.1 above, though I have little to say here about the complications. Extraction will require a certain kind of downstream mechanism that is tuned to constituency, not mere parthood. This will be easiest for those schemes where the easily separable parts always are constituents, as in microfeatural representations.

British were coming by land and two if they were coming by sea. Here the coming-by-sea representation contains the coming-by-land representation, but only as a part, not as a constituent. They could easily agree to let one lantern mean that the British aren't coming at all, and the pair of lanterns might still represent their coming by sea.

Due to clause (i), the present definition of constituency is fairly strict.<sup>21</sup> In fact, fewer things satisfy it than satisfy Fodor and McLaughlin's necessary condition for something's being a classical constituent: 'for a pair of expression types E1, E2, the first is a *Classical* constituent of the second only if the first is tokened whenever the second is tokened' (1990, p. 186).<sup>22</sup> This is true even if we add something like my clause (ii) to their characterization. If Fodor and McLaughlin mean to be restricting classical constituency to Language of Thought schemes, they will need further necessary conditions; otherwise pictures and microfeatural representations will have classical constituents.

The above definition of constituency makes it possible to distinguish genuinely complex representations from primitives:  $R$  is a *complex representation* in  $S$  iff  $R$  can be divided without remainder into two or more constituents in  $S$ . All other elements (representational elements as well as operators) in the scheme are *primitives*. The composition rules adverted to in (d) are only defined for complex representations, so the definition of complex representations solves the vacuity problem with (d), and consequently paves the way for a solution to the two problems that started this section.

### 3.2 Kinds of Constituents, and a More Precise Account of Representational Analyticity

The composition rules for a given scheme will frequently treat different combinations of representations differently. This allows us to distinguish different kinds of constituents, depending on how they contribute to the extensions of the complex representations in which they occur. For example, we can say that  $R_1$  is a *conjunctive constituent* of  $R$  in  $S$  iff  $R$  is a complex representation consisting of  $R_1 \dots R_n$  as constituents (in  $S$ ), and the extension of  $R$  in  $S$  is specified by the composition rules as being the intersection of the extensions of  $R_1 \dots R_n$ . To say that a representation *consists of*  $R_1 \dots R_n$  as constituents in  $S$  is to say that the representation in question can be divided without remainder into  $R_1 \dots R_n$ , where  $R_1 \dots R_n$  are constituents of that representation in  $S$ . All the constituents of:

<sup>21</sup> See, for example, van Gelder (1990) for a more lenient notion, viz., that of functional constituency. Van Gelder claims that connectionist representations exhibit only functional constituency. This is a mistake, as the example of microfeatural representations shows.

<sup>22</sup> Fodor and McLaughlin explicitly allow that being a part is one way of being a classical constituent, although not the only way. Absent additional necessary conditions, there will be at least as many classical constituents in their sense as there are constituents in the present sense.

[YELLOW]  $\wedge$  [FUSIBLE]  $\wedge$  [METAL]  $\wedge$  [SOLUBLE IN AQUA REGIA]  $\wedge$   
[FIXED]  $\wedge$  [MALLEABLE],

for example, are conjunctive constituents. ‘cat’ is a conjunctive constituent of ‘black cat’ (as is ‘black’). All the constituents of any representation in the traditional Lockean scheme are conjunctive constituents; the same holds for microfeatural representational schemes. This is also the case for pictorial representations, though this is typically not noticed, perhaps because the compositionality of a pictorial representational scheme is a consequence of the isomorphism semantics and does not need to be explicitly built into the scheme by way of composition rules.

Not all constituents, however, are conjunctive constituents. I will say that  $R_1$  is a *negated constituent* of  $R_2$  in  $S$  iff  $R_2$  is a complex representation containing  $R_1$  as a constituent (in  $S$ ), and the extension of  $R_2$  in  $S$  is specified by the composition rules as being a subset of the complement of  $R_1$ . So, for example, ‘dog’ is a negated constituent of ‘non-dog’; ‘gun’ is a negated constituent of ‘fake gun’. This last example illustrates the reason for appealing to subsets in the definition of a negated constituent. There may be fewer fake guns than non-guns (the former may need to resemble guns to some extent), even though all fake guns are non-guns.

There are many other types of constituents, and I won’t bother naming them all. We can distinguish different kinds of representational analyticity, based on the differences in the kinds of constituents involved. Consider a few examples:

1.  $\lceil \text{All } F\text{s are } G\text{s} \rceil$  is representationally analytic with respect to  $S$  iff every representation in  $S$  of something’s being  $F$  contains, as a conjunctive constituent, a representation of its being  $G$ .
2.  $\lceil \text{No } F\text{s are } G\text{s} \rceil$  is representationally analytic with respect to  $S$  iff every representation in  $S$  of something’s being  $F$  contains, as a negated constituent, a representation of its being  $G$ .
3.  $\lceil x + y = z \rceil$  is representationally analytic with respect to  $S$  iff every representation in  $S$  of  $z$  consists of a representation of  $x$  and a representation of  $y$ , where the composition rules of  $S$  specify the extension of a complex numerical representation as being the sum of its constituents.
4.  $\lceil m \times n = z \rceil$  is representationally analytic with respect to  $S$  iff every representation in  $S$  of  $z$  consists of a representation of  $m$  and a representation of  $n$ , where the composition rules of  $S$  specify the extension of a complex numerical representation as being the product of its constituents.

Formulations analogous to the last two can be given for subtraction and division as well. Even so, the list is incomplete and not very well unified. It would be nice to have a single formulation that is general enough to handle all the instances of representational analyticity.

The post-Fregean theory of analyticity viewed analyticity as truth by definition. Representational analyticity, however, is a certain kind of truth by scheme specification. The immediately preceding formulations provide a way of understanding

this. Supposing that a proposition asserts a relation among properties (i.e. the extensions of certain representations), a rough and intuitive notion of representational analyticity emerges. In the above examples, the proposition in question asserts a relation among two or more properties, and the representational scheme is such that that same relation is what determines the extensions of the relevant representations.

More precisely,

(RA): That  $X$  bears  $R$  to  $Y_1 \dots Y_n$  is *representationally analytic* with respect to  $S$  iff for every representation,  $a$ , in  $S$ , which has  $X$  as its extension,

- (i)  $a$  consists of some constituents  $z_1 \dots z_n$ ,
- (ii)  $Y_1 \dots Y_n$  are the extensions of  $z_1 \dots z_n$ , respectively,
- (iii) the scheme specification for  $S$  entails that the extension of  $a$  is a function of the extensions of  $z_1 \dots z_n$ , in particular, that the extension of  $a$  bears  $R$  to the extensions of  $z_1 \dots z_n$ .

The  $z_i$ s can themselves be complex representations. 1–4 above are merely instances of this more general formulation.

(RA) ensures that only necessary truths will come out representationally analytic. Since the extension of a representation, as I'm using the term, includes all the possible objects to which the representation refers,  $a$  couldn't have  $X$  as its extension unless  $X$  bears  $R$  to  $Y_1 \dots Y_n$  in all possible worlds. It is less clear whether all necessary truths are representationally analytic; in fact, I will leave this question open. I have only scratched the surface of mathematics, for example. There are, however, far more representationally analytic propositions than one might initially think, certainly more than Frege thought.

Anything that the post-Fregean account of analyticity could capture as true in virtue of concept containment can also be shown to be representationally analytic. Suppose that there are defining, hence metaphysically necessary and sufficient, conditions,  $C_1 \dots C_n$ , for something's having property  $P$  (e.g., suppose that 'bachelor' is synonymous with 'unmarried man'). Then there will be a Locke-like scheme that does not contain a primitive representation of  $P$ , but can only represent something's having  $P$  by representing its having  $C_1 \dots C_n$ . With respect to such a scheme, any proposition of the form 'All  $P$ s are  $C$ s' is representationally analytic. I am not at all sure that there are any interesting such analyticities and don't know how to find out—this is one of the major shortcomings of the Fregean understanding of analyticity—but if there are, the present account can handle them as well as the post-Fregean account can.

I mentioned at the beginning of section 2 that ' $p \supset (p \vee q)$ ' is generally thought not to be analytic in virtue of concept containment. It is, however, if we read concept containment as representation containment in some scheme. All we need is a scheme consisting only of (a) a primitive representation,  $x$ , which means that  $p$  or  $q$  (i.e. is true in just those worlds in which ' $p \vee q$ ' is true), (b) another

representation,  $\gamma$ , which means that not- $q$ , and (c) an ‘&’ operator. In this scheme, the only way to represent  $p$  is via one of the following representations:

$$x \ \& \ \gamma \qquad \gamma \ \& \ x$$

and each representation does contain  $x$ , which is a representation of  $p \vee q$ , as a conjunctive constituent.


Finally, although the standard view had it that concept containment analyticities were special cases of truths of logic, it may turn out to be the other way around. At least some truths of logic, including some of the more complicated ones, can be captured as instances of representation containment. The following, for example, is not *representationally* analytic with respect to the first order predicate calculus:

$$[(\forall x)(Ax \supset Bx) \ \& \ \sim \exists y(Cy \ \& \ By)] \supset \sim \exists z(Az \ \& \ Cz).$$


It is, however, representationally analytic with respect to a Venn diagram scheme. The Venn diagram representation of  $(\forall x)(Ax \supset Bx)$ , shown in Figure 4, conjoined with the representation of  $\sim \exists y(Cy \ \& \ By)$ , shown in Figure 5, is shown in Figure 6. This representation contains the representation of  $\sim \exists z(Az \ \& \ Cz)$ , shown in Figure 7. Again, I will leave open the question of whether all necessary truths are representationally analytic with respect to some scheme or other. I am willing to allow for the possibility of representationally synthetic necessary truths, although the apparent existence of such may simply be due to a lack of creativity in devising the necessary schemes.

### 3.3 Naturalism and Representational Analyticity

Quine thought—rightly—that semantic notions are mysterious if left unanalyzed, unnaturalized. This likely had a lot to do with his distrust of the analytic/synthetic distinction, for the standard post-Fregean appeal to synonymy and definition not only invokes semantic notions, but the worst kind of semantic notions: those that do not seem to lend themselves to naturalization. However, now that we have the beginnings of some naturalistic theories of content, there is no reason not to use them in offering a naturalistic explication of the analytic/synthetic distinction. If the analytic/synthetic distinction can be naturalized, then much of the murkiness surrounding it will have been removed.




**Figure 4**  $(\forall x)(Ax \supset Bx)$




**Figure 5**  $\sim \exists y(Cy \ \& \ By)$

The present account does take the notion of content, of extension, as given. This, however, is one of the semantic notions that seems most amenable to naturalization by a good theory of content. If there is anything that a naturalistic theory of content should deliver, it will be an account of reference, and the extension is just the set of actual and possible things to which a representation refers. Fodor's asymmetric dependence theory, for example (Fodor, 1987, 1990) fairly explicitly offers a theory that maps representations onto properties (1987,



**Figure 6**  $(\forall x)(Ax \supset Bx) \ \& \ \sim \exists y(Cy \ \& \ By)$



**Figure 7**  $\sim \exists z(Az \ \& \ Cz)$

pp. 98–101), which are more or less what I have been calling extensions. This particular theory does not attempt to specify the composition rules (p. 98), but there is no reason to think that this is beyond the pale of a naturalistic semantics.

Of course, asymmetric dependence theory is used here merely as an example—a teleosemantics or any other naturalistic semantics would do just as well. Even a use theory, like conceptual role semantics, might fit the bill, provided it is naturalistic, and although naturalism is usually taken as a constraint on theories of content anyway (Cummins, 1989; Fodor, 1987), it is one that use theories are going to be tempted to violate. Fodor and LePore (1991, 1992) have argued that the only way for a conceptual role semantics to avoid an implausible holism would be to include an analytic/synthetic distinction as part of an idealization condition, to distinguish content-fixing inferences from non-content-fixing inferences. They object to this on the grounds that the analytic/synthetic distinction is untenable. Another problem, however, is that the analytic/synthetic distinction is itself a semantic distinction, so the resulting theory would therefore no longer be a naturalistic one. If there is a *naturalistic* way of distinguishing inferences that determine content from inferences that don't, the resulting conceptual role semantics could play the role required here.

### 3.4 Fine-grained Individuation of Belief

I have been supposing so far that beliefs/judgments are individuated by their propositional content. Because of this, the account has proceeded in terms of propositions, where the latter are quite coarsely individuated. This approach makes it possible to provide a theory of the (representational) analyticity of propositions, but it also has its limitations. Suppose that it is metaphysically necessary that all and only humans are rational animals. Then the Mentalese formula 'RATIONAL ANIMALS ARE RATIONAL' would express the same proposition as the formula 'HUMANS ARE RATIONAL', at least given a certain way of individuating propositions. We might, however, not want to count these as the same judgment, for precisely the sorts of reasons that motivate an interest in representational analyticity in the first place. The inference from 'RATIONAL ANIMAL' to 'RATIONAL' ought to be credited with a much higher degree of certitude than the inference from 'HUMAN' to 'RATIONAL', and this would be true even if 'RATIONAL ANIMAL' is synonymous with 'HUMAN', since synonymy is so much harder to detect than constituency. Yet the *proposition* 'rational animals are rational', i.e. the proposition 'humans are rational', is not representationally analytic with respect to this scheme, for 'rational' is only a conjunctive constituent of some, but not all representations of humanity.

The intuitive idea that needs to be captured is this: the *representation*  $\lceil$ All *F*s are *G*s $\rceil$  is representationally analytic with respect to *S* if and only if *G* is a conjunctive constituent of *F* in *S*. The more general account of section 3.2 can be adapted. The changes are quite small if we assume that the representations in question are propositional:

(RA<sub>R</sub>): The representation  $\lceil aRz_1 \dots z_n \rceil$  is representationally analytic with respect to  $S$  iff

- (i)  $z_1 \dots z_n$  are constituents of  $a$ , and
- (ii) the scheme specification for  $S$  entails that the extension of  $a$  is a function of the extensions of  $z_1 \dots z_n$ , in particular, that the extension of  $a$  bears  $R$  to the extensions of  $z_1 \dots z_n$ .

A slightly different version is needed if we do not assume that the representations are propositional, but the modifications are fairly straightforward, and I will not bother with the details except to note that the formulation ought to proceed in terms of the representational analyticity of *inferences*, where an inference is understood as a move from one representation to another (not from one content to another), e.g. from:

$$[\text{RATIONAL}] \wedge [\text{ANIMAL}]$$

to:

$$[\text{ANIMAL}].$$

Thus, we have two ways of characterizing representational analyticity: one that is defined over propositions and another that is defined over representations. So we can have a theory of analytic judgments whether we take judgments to be individuated by propositional content or by the representations involved. For ease of exposition, I will ignore these newer complications in what follows and act as if the bearers of representational analyticity are propositions, or judgments individuated by their contents.

#### 4. Analytic Knowledge

Although I think that representational analyticity is interesting enough in its own right to merit discussion, I mentioned earlier that the account developed here would have epistemological consequences. The fact that there are infinitely many schemes indicates that for any given scheme, many propositions that are not representationally analytic with respect to *it* will nonetheless be representationally analytic, in virtue of being representationally analytic with respect to some other scheme. An obvious question arises: suppose that somewhere in Plato's heaven there is a representational scheme with respect to which  $p$  is representationally analytic but that none of us has ever even imagined such a scheme; why should its mere existence carry any epistemological weight?

The answer is that it shouldn't. There is no reason to think that all analytic truths are analytically knowable for us; rather, what is analytically knowable for us should

depend on what kinds of schemes we are capable of using.<sup>23</sup> The general idea here is that:

(AK): *p* is *analytically knowable* for some cognizer *C* iff there is some scheme *S* which is available to *C* and with respect to which *p* is representationally analytic.

I think that (AK) is intuitively appealing, and I have intentionally stated it imprecisely in order to maximize this intuitive appeal. In particular, the notion of availability could use some sharpening.

There are two main senses of 'availability', each with something in its favor. I will say that scheme *S* is *weakly available* to *C* iff *C* is a competent user of *S*. I don't have much to say about how competent user status should be understood, but the idea is familiar enough from appeals to competent speakers of a language, and so on (though I am not sure that anyone has ever made much sense of the notion). I am more interested here in contrasting it with a stricter understanding of availability. I will say that *S* is *strongly available* to *C* iff *S* is one of the mental representational schemes *in which C* forms beliefs, i.e. if it is items in *S* that *C* tokens in tokening occurrent beliefs.

The analysis is stated thus to leave room for the important possibility that we have more than one scheme of mental representations and that more than one of these is actually involved in thought, belief, and inference. This possibility should be attractive to any proponent of 'mad dog modularity' (Cosmides and Tooby, 1992; Segal, 1996; Sperber, 1994) as the different modules are likely to use different representational schemes (Bever, 1992), though it is also compatible with other views. Researchers in different fields of psychology frequently posit different representational schemes: hash marks for reasoning about numerosity (or perhaps only for small numerosities), pictorial representations for imagery tasks, and so on. There is no reason to think these theorists are disagreeing about which representational schemes the mind uses; rather the assumption seems to be that the mind may use multiple schemes, depending on the task in question.

There are two very different ways in which we may *use* representational schemes without thereby *thinking in* these schemes. To think in a scheme is for items of that scheme to be the representational relata involved in belief. There are some mental representations that never enter into beliefs in this way, as well as some 'external' representations, which we can use as cognitive aids but which are not literally parts of beliefs.

Some of the mental representational schemes we use may not be schemes in which we think. Psychologists typically posit a large number of unconscious mental representations, which, though causally implicated in belief formation, are

---

<sup>23</sup> I will use 'analytically knowable' for 'representationally analytically knowable,' as the longer phrase would be cumbersome.

not the representations that are tokened in tokening a belief. Recall that according to representationalist functionalism, belief is a certain functionally defined relation to one or more mental representations. Some of my mental representations, however, are such that I never stand in *that* relation to them. The relation in which I stand to a 2½-D sketch (Marr, 1982) while processing visual information is very different from the one in which I stand to my 'car' representation while tokening the belief that my car is low on gas.

Readers familiar with the cognitive scientific literature on concepts will have noticed that I have said little about prototypes, exemplar sets, whole theories and the like. One reason is that they do not support any interesting representational analyticities, as far as I can tell. More importantly, however, these representations have much the same status as the 2½-D sketch, or the frames of classical AI; they are knowledge structures that are causally implicated in belief-formation, but not necessarily the actual relata of the belief relation. According to the exemplar theory, for example, a concept is a set of (possibly hundreds or more) representations of individual members of some category. This set is presumably not what is tokened in tokening a conscious, occurrent belief, since such a set would never fit into working memory.

So 2½-D sketches, frames, and likely even what many psychologists tend to call concepts are the wrong sorts of mental representation for literally entering into the belief relation, though they are certainly capable of causally contributing to belief. The other kind of representation we can use without thereby thinking in is a weakly available 'external' representation. Suppose that we reason in English (improbably) but that we learn to use Venn diagrams in reasoning. Such learning need not require that the reasoning systems are actually tokening Venn diagram representations. It may be that our use of the diagrams involves the tokening of English descriptions of Venn diagrams, rather than the diagrams themselves. What is representationally analytic with respect to the descriptions will be quite different from what is representationally analytic with respect to the diagrams. The English sentence 'the circle labeled "A" is completely filled in' does not literally contain the sentence 'the area of overlap between the circle labeled "A" and the circle labeled "B" is filled in'. In this way, a scheme could be weakly available without being strongly available. We would be thinking *about* Venn diagrams, but not thinking *in* them.

In pursuing the epistemological implications of representational analyticity, we need not choose between strong and weak availability, but we must recognize an epistemological difference between them. The analyticities supported by strongly available schemes are more knowable, more certain, than the ones supported by the merely weakly available schemes. One reason is simply that there is more room for error, since the use of an 'external' scheme (one that is merely weakly available) will always involve two representational schemes. I am happy to allow that both strong and weak availability suffice for analytic knowability, but we should distinguish strongly analytically knowable propositions from weakly analytically knowable propositions, corresponding to the different kinds of availability.

It is an empirical question—and one for which we have few if any uncontroversial answers—which representational schemes are strongly available to us. Again, the set of strongly available schemes will be a subset of the set of mental representational schemes, but here are a few plausible candidates for strongly available representational schemes:

- Fodorian Language of Thought (Fodor, 1975, 1981)
- Various kinds of connectionist representations, including microfeatural representations (e.g. Rumelhart and McClelland, 1986)
- certain pictorial schemes (e.g. Kosslyn and Pomerantz, 1977; Kosslyn and Koenig, 1992) or quasi-pictorial schemes (e.g. Paivio, 1971)
- traditional Lockean schemes (e.g. Armstrong, Gleitman, and Gleitman, 1983)
- hash-mark-like schemes (e.g. Wynn, 1992)

These are actually only classes of schemes, rather than particular schemes, for these broad characterizations leave a number of free parameters concerning primitive size and so forth. Still, they offer a sense of the empirical possibilities. Some of these schemes will support few, if any representational analyticities, though some support some very interesting representational analyticities.<sup>24</sup> And despite occasional claims to the contrary, the most likely possibility is that the mind utilizes a wide variety of different representational schemes (possibly differing for different tasks or for different systems or modules), each supporting its own analyticities. If few or none of the representational schemes used by the mind support any interesting representational analyticities, this is itself an important finding.

I have been focusing here on strong analytic knowability, because it is an interesting but neglected issue; this, however, should not overshadow the importance of weak analytic knowability. The term ‘weak’ here is used only to distinguish it from the other kind. Weak analytic knowledge, in fact, is quite similar to the kind of analytic knowledge provided by the standard post-Fregean view, according to which analytic knowledge is knowledge in virtue of linguistic competency. Such knowledge would be analytic (and hence a priori and the rest), but still not quite as epistemically secure as strong analytic knowledge. The fact that there is an even more secure kind of analytic knowledge, however, should not be taken to denigrate the epistemic status of weak analytic knowledge or to place it on a par with (mere) empirical knowledge. Learning how to use Venn diagrams really does improve one’s epistemic position *vis-à-vis* logic, just not as much as would

---

<sup>24</sup> Smolensky-style tensor product representations (1991), for example, are connectionist representations that are basically like Gödel numberings: none of the representations are complex in the present sense of the term. Consequently, the scheme supports no representational analyticities at all aside from identical propositions. Karen Wynn’s hash-mark-like scheme for representing numerosity, on the other hand, supports all of the analyticities involving addition and subtraction of the natural numbers mentioned in section 2.3.

having one's brain rewired to have Venn diagrammatic representations built into the functional architecture of the reasoning processes. Assuming the failure of the post-Fregean understanding of analyticity, representational analyticity may very well provide the only way to understand either kind of knowledge.

An interesting consequence of the epistemological theory sketched here is that it becomes an empirical question which propositions are (strongly) analytically knowable for us. This, of course, follows straightforwardly from the fact that it is an empirical question which schemes are strongly available to us. This may strike some readers as problematic.

It is clear enough that the present view commits me to denying the famous KK principle:

(KK): If *S* knows that *p*, then *S* knows that *S* knows that *p*.

This by itself is not too worrisome. Any good externalist or epistemological naturalist will deny the KK principle (e.g. Nozick, 1981), and non-naturalist internalists are free to do so as well. However, I am also committed to denying the following, somewhat weaker principle:

(KA): If *p* is analytically knowable for *S*, then *S* knows that *p* is analytically knowable for *S*.

There are independent grounds, however, for denying this weaker claim. (KA) implies that it is impossible for someone be a competent user of a scheme without knowing that she is a competent user of that scheme, and this seems wrong.<sup>25</sup> Surely it is possible to doubt (and even to have good reason to doubt) one's competence regarding a scheme of which one is in fact a competent user.

This, of course, is only the beginning of an epistemology that takes representational analyticity seriously. But even this sketch ought to be enough to show that representational analyticity is a phenomenon worthy of our attention.

## 5. Conclusion

Unlike the post-Fregean account, the traditional account of the analytic/synthetic distinction, developed by Locke and others, was supposed to have not only philosophical, but also psychological consequences. The notion of representational analyticity bolsters a distinction between two importantly different kinds of psychological processes: those that involve the extraction of constituents, and

---

<sup>25</sup> Even on the standard post-Fregean account of analyticity, analytic knowability is supposed to be merely a matter of competent use of the language, where competent use involves knowing the relevant definitions. Absent *something* like this, there is no significant epistemological distinction between analytic knowability and regular knowability, in which case, (KA) is as strong and as implausible as (KK).

those that involve the tokening of different representations in response to certain given representations. Having a relatively precise and detailed account of representational analyticity gives us a much clearer understanding of this important psychological distinction.

As mentioned from the outset, a theory of representational analyticity is less ambitious than the sort of theory of analyticity that dominated most of post-Fregean philosophy. Representational analyticity will be of little use in resurrecting logical positivism or providing an idealization condition for a conceptual role semantics, for instance. But these may very well be things that cannot be done by any theory of analyticity; the post-Fregean account never lived up to its own ambitions, after all.

And representational analyticity does have philosophical consequences. The account of representational analyticity illuminates important features of representations, including mental representations. Furthermore, I have been arguing that the epistemological work of the traditional analytic/synthetic distinction can be done by the notion of representational analyticity, and better than by the post-Fregean conception. Representational analyticity explains how certain things are as easy to know as they are and how they can be known independently of experience. A surprising and important result of the present view, but one I have argued we should be prepared to accept, is that it is an empirical question which propositions are representationally analytic for us.

Most contemporary philosophers seem to think that Quinean arguments show the analytic/synthetic distinction to be untenable. This indicates the prevalence of the attitude that if a conception of analyticity is to be viable, it must be a roughly post-Fregean one. I have tried to offer something different here, proceeding without reliance on such dubious notions as sense or synonymy. If nothing is analytic, then it may very well turn out that nothing is a priori, and this would be a victory for the radical empiricist. Representational analyticity, however, would guarantee the possibility of *some* a priori knowledge, thus rendering an acceptance of the Quinean arguments compatible with the denial of radical empiricism. The decline of the post-Fregean conception of analyticity should not lead us to abandon the more traditional analytic/synthetic distinction or what it aimed to do.

*Department of Philosophy  
University of Arkansas*

## References

- Armstrong, S.L., Gleitman, L.R. and Gleitman, H. 1983: What some concepts might not be. *Cognition*, 13, 263–308.
- Bever, T.G. 1992: The logical and extrinsic sources of modularity. In M. Gunnar and M. Maratsos (eds), *Modularity and Constraints in Language and Cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates.

- Cosmides, L. and Tooby, J. 1992: Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, and J. Tooby (eds), *The Adapted Mind*. New York: Oxford UP.
- Cummins, R.C. 1989: *Meaning and Mental Representation*. Cambridge, MA: MIT Press.
- Cummins, R.C. 1992: Conceptual role semantics and the explanatory role of content. *Philosophical Studies*, 65, 103–27.
- Cummins, R.C. 1996: *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.
- Fodor, J.A. 1975: *The Language of Thought*. New York: Crowell.
- Fodor, J.A. 1981: The present status of the innateness controversy. In J.A. Fodor, *Representations*, Cambridge, MA: MIT Press.
- Fodor, J.A. 1985: Fodor's guide to mental representation. *Mind*, 94, 55–97.
- Fodor, J.A. 1987: *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, J.A. 1998: *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J.A. and LePore, E. 1991: Why meaning (probably) isn't conceptual role. *Mind & Language*, 6, 328–43.
- Fodor, J.A. and LePore, E. 1992: *Holism: a Shopper's Guide*. Cambridge, MA: Blackwell.
- Fodor, J.A. and McLaughlin, B. 1990: Connectionism and the problem of systematicity: why Smolensky's solution doesn't work. *Cognition*, 35, 183–204.
- Fodor, J.A. and Pylyshyn, Z.W. 1988: Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28, 3–71.
- Frege, G. 1953: *The Foundations of Arithmetic*. J.L. Austin, trans. Oxford: Blackwell.
- Harman, G. 1982: Conceptual role semantics. *Notre Dame Journal of Formal Logic*, 23, 242–56.
- Haugeland, J. 1990: Representational genera. In W. Ramsey, S. Stich, and J. Rumelhart (eds), *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kant, I. 1781/1929: *Critique of Pure Reason*. N.K. Smith (trans.). New York: St. Martin's Press.
- Katz, J.J. 1992: Analyticity. In J. Dancy and E. Sosa (eds), *A Companion to Epistemology*. Cambridge, MA: Blackwell.
- Katz, J.J. 1998: *Realistic Rationalism*. Cambridge, MA: MIT Press.
- Kosslyn, S.M. and Koenig, O. 1992: *Wet Mind*. New York: Free Press.
- Kosslyn, S.M. and Pomerantz. 1977: Imagery, propositions, and the form of internal representations. *Cognitive Psychology*, 9, 52–76.
- Lewis, D.K. 1970: How to define theoretical terms. *Journal of Philosophy*, 67, 427–46.
- Lewis, D.K. 1986: *On the Plurality of Worlds*. Cambridge, MA: Blackwell.
- Locke, J. 1975: *An Essay Concerning Human Understanding*. P.H. Nidditch (ed.), Oxford: Oxford UP.
- Marr, D. 1982: *Vision*. New York: Freeman.
- Nozick, R. 1981: *Philosophical Explanations*. Cambridge, MA: Harvard UP.
- Paivio, A. 1971: *Imagery and Verbal Processes*. New York: Holt.

- Quine, W. V. O. 1951: Two dogmas of empiricism. *Philosophical Review*, 60, 20–43
- Rosch, E. 1975: Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192–233.
- Rumelhart, D.E., McClelland, J.L. and the PDP Research Group. 1986: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1: Foundations*. Cambridge, MA: MIT Press.
- Segal, G. 1996: The modularity of theory of mind. In P. Carruthers and P. Smith (eds), *Theories of Theories of Mind*. Cambridge: Cambridge UP.
- Sejnowski, T.J. and Rosenberg, C.R. 1987: Parallel networks that learn to pronounce English text. *Complex Systems*, 1, 145–68.
- Shin, S. 1994: *The Logical Status of Diagrams*. Cambridge: Cambridge UP.
- Smith, E.E. and Medin, D.L. 1981: *Categories and Concepts*. Cambridge, MA: Harvard UP.
- Smolensky, P. 1991: Connectionism, constituency, and the language of thought. In B. Loewer and G. Rey (eds), *Meaning in Mind: Fodor and his Critics*. Oxford: Oxford UP.
- Sperber, D. 1994: The modularity of thought and the epidemiology of representations. In L. Hirschfeld and S. Gelman (eds), *Mapping the Mind: Domain Specificity in Thought and Culture*. Cambridge: Cambridge UP.
- Swoyer, C. 1991: Structural representation and surrogate reasoning. *Synthese*, 87, 449–508.
- van Gelder, T. 1990: Compositionality: a connectionist variation on a classical theme. *Cognitive Science*, 14, 355–84.
- White, M.G. 1950: The analytic and the synthetic: An untenable dualism. In S. Hook (ed.), *John Dewey: Philosopher of Science and Freedom*. New York: Dial Press.
- Wynn, M.K. 1992: Evidence against empiricist accounts of the origins of numerical knowledge. *Mind & Language*, 7, 315–32.