Representational Specialization:
The Synthetic A Priori Revisited
In *The World in the Head,* Oxford Univ. Press, 2010
Robert Cummins
UIUC

0. Introduction

Anyone familiar with the history of western philosophy is familiar with the idea that our knowledge of the world is limited, by our representational resources, especially our perceptual and conceptual resources. The empiricist version of this theme has it that all of our ideas are derived from perceptual or reflective *experience*. Locke, Berkeley and Hume all held that all of our ideas are derived from sense and reflection, and subsequent empiricist philosophy has not challenged this in any serious way. Rationalists have allowed for innate ideas of various sorts not derived from experience, but have not challenged the claim that what can be known is constrained by what concepts and percepts we can have. Indeed, this seems a truism: how could one possibly know about what one cannot conceive or perceive?

I think it is possible to challenge this truism, but I do not intend to do that here. Rather, I want to focus on its critical applications, i.e., attempts to rule out this or that as unknowable on the grounds that there is no plausible source for the underlying representations. Those uses of the "truism" have been, I think, largely unsuccessful. Perhaps the most famous example of what I have in mind is the final paragraph of Hume's Enquiry:

> When we run over libraries, persuaded of these principles, what havoc must we make? If we take in our hand any volume; of divinity or school metaphysics, for instance; let us ask, Does it contain any abstract reasoning concerning quantity or number? No. Does it contain any experimental reasoning concerning matter of fact and existence? No. Commit it then to the flames: for it can contain nothing but sophistry and illusion. (Hume, 1748)

A. J. Ayer famously dismissed ethics and theology on the same grounds. (Ayer, 1936)

Critiques of this nature cannot hope to get off the ground unless there is a route to claims about what can be represented that is prior to claims about what can be known: For example, there must be a route to the claim that all of our ideas are derived from sense and reflection (a claim badly in need of unpacking before its credentials can be examined) that is prior to—i.e., does not presuppose—the claim that we cannot know about things that we cannot experience. Evidently, any story about representation that takes it to be, or to depend on, a kind of knowledge will undermine a critique that moves from the alleged origins representations to limits on representational resources and thence to limits on what can be known.

1. A Kantian thought experiment

To prepare the ground for the discussion to follow, I begin with a thought experiment that owes its inspiration to Kant's *Aesthetic*. Imagine a perceptual system—call it *Cubic*—that represents objects by coloring the surface of cubes in a 3-dimensional packed array of such cubes, like a Rubic's cube that is 10,000 X 10,000 X 10,000. This gives colored blobs of various sizes and shapes in a three dimensional space. It is easy to take this the wrong way and imagine the system *looking at* such a cube, noticing that you cannot see the cells inside, and concluding that there is no perceptual representation of the insides of objects on this scheme. On the contrary, the scheme I am imagining here contains information about the colors of the insides of things that are larger than one cell. An uncolored cell inside a large object representation represents an unresolved ambiguity between, say, an apple (red outside, white inside) and a hollow plastic apple (red outside, hollow inside). It is also, for some of us, anyway, to think of this symbolically, as a three dimensional matrix with cells bound to RGB values. That would be a mathematical description of Cubic, but Cubic itself, as I want you to imagine it, involves 10K cubes packed together, and a mechanism for coloring them. It is, if you like, an *implementation* of the matrix scheme.

If we assume that representations of this sort are the only representations that Cubic can generate, then a number of propositions will be synthetic and a priori for Cubic:

V1. Every object is colored.

V2. Every object has a determinate size and shape.

V3. No two objects can occupy the same place at the same time.

V4. Every object has a determinate location relative to every other object.

V5. Every object is a determinate distance from every other object.

Of course, Cubic does not represent any of these propositions. Indeed, Cubic cannot represent any propositions at all.[1] Its only representational targets are colored shapes in three-dimensional space. Nevertheless, *we* are in a position to see that these propositions are somehow inevitable for Cubic, even though they are, in fact, all empirical, and all false. These propositions will be *a priori* for Cubic in the sense that Cubic cannot represent a counter-example to any of V1-V5. A representational target for Cubic that violates any of V1-V5 cannot be accurately represented in Cubic. If Cubic attempts to represent such a target—e.g., something like an electron that has a size but no shape, no color, and no determinate location relative to its neighboring particles, it will get it massively wrong. The result will be a case of forced error: an attempt to represent a target that Cubic does not have the resources to represent accurately (Cummins, 1996, ch. 2). V1-V5 are, of course, synthetic propositions, as the electron case illustrates. So we have here a toy case of the synthetic a priori. It is a *relativized* a priori—relativized to Cubic, or the representational scheme Cubic uses—but so also was Kant's notion relativized to *human* cognitive architecture.

It is, of course, odd to talk of propositions a priori for Cubic, given that it cannot represent propositions.[2] It is tempting to say, instead, that V1-V5 must hold of the world *as it is represented by Cubic.* This is harmless enough, provided we mean by it simply that Cubic cannot represent a state of affairs that would be a counter-example to any of V1-V5. It is not harmless if we understand this to mean that there is something called Cubic's phenomenal world, the *world-as-represented-by-Cubic,* and that V1-V5 are true—indeed necessarily true—of *that.* There is no such world, and so this kind of talk is not harmless: it wrongly suggests that there is a "world", other than the actual world, that Cubic represents, and represents just as it is—a phenomenal world about which Cubic cannot be mistaken, or anyway, about which it cannot be mistaken in the ways picked out by V1-V5. We need to insist that Cubic represents the actual world, and represents it in a way that ensures that no representation of a counter-example to V1-V5 can possibly occur. From our God-like positions, we can see that the representations Cubic

[1] It might *encode* propositions (see below), but I am assuming here that Cubic has no means to exploit such a possibility.

[2] Do not be confused by the fact that *we* can read off (infer) propositions from one of Cubic's representations. Cubic has nothing like a *language;* it has no resources for representing propositions.

constructs are necessarily inaccurate in certain respects, but, to borrow a phrase from Descartes, this is nothing to Cubic. V1-V5 are part of the synthetic a priori content inherent in the form of P's intuitions and their representational functions—their proprietary targets.[3] It is because the targets of P's representations are objects in space that V1-V5 are about objects in space. It is because of the particular form of Cubic's representations that it must represent those objects in a way that satisfies V1-V5.

You don't have to be the Village Empiricist to suspect that Cubic cannot escape the limitations imposed by V1-V5 without modification to its resources for perceptual representation. But it can. Let us now add a non-perceptual scheme for propositional representation, a language of thought: a scheme that will enable the expanded system to represent, among others, the propositions expressed by V1-V5 and their negations. Call this expanded system P-Cubic (Propsitonal-Cubic). It is a tempting Kantian idea to think of P-Cubic's *experience,* its epistemic access to the world, as consisting of its perceptual representations, and to point out that P-Cubic could never have an experience that directly contradicted any of V1-V5. This would not be exclusively a contingent fact about the external world, but also and correlatively a consequence of the form of P-Cubic's representational resources. V1-V5 would be synthetic for P-Cubic—not conceptual truths—but they would be a priori in the sense that no disconfirming experience is possible for P-Cubic. For P-Cubic, colored shapes in a Cartesian 3-space would be the a priori form of outer intuition, i.e., of perceptual object representation. The form of P-Cubic's perceptual representations constrain not only what it can accurately perceive, but also, indirectly, what it is rational for it to think is true. P-Cubic could contemplate the negation of, say, V1, but could never have an experience disconfirming V1. Or so goes one version of the Kantian story and its empiricist predecessors. It is beguiling, but it is fundamentally mistaken.

Begin with a relatively trivial point. We have to distinguish between an experience of an uncolored object, and an experience of an object as uncolored, i.e., between experiencing a colored object, and experiencing an object *as colored*. By construction, it is impossible for P-Cubic to experience an uncolored object *as uncolored*. But P-Cubic might well experience an uncolored object, because P-Cubic might experience uncolored objects, yet experience them as colored. These will all be cases of forced error, of course, and so it might seem that P-Cubic could not *know* that it was experiencing an uncolored object, and hence that P-Cubic could not know there were such things.[4] But this is mistaken: P-Cubic might well *discover* the existence of uncolored objects, their locations and their trajectories, without ever experiencing such objects as such, or even without experiencing them at all. Even a rudimentary naïve mechanics, gleaned from what *is* available in perception, could provide P-Cubic with persuasive evidence of what is not, e.g., an imperceptible wall on a pool table. What makes this possible is precisely the fact that P-Cubic can represent and evaluate propositions that are true of possibilities that it cannot experience as instances of those possibilities, but for which P-Cubic can have persuasive indirect evidence. Because P-Cubic's propositional representations do not have the same contents as P-Cubic's perceptual representations--because its concepts are not *copied* from its percepts, but do apply to the objects represented in those percepts--P-Cubic is in a position infer how things should look on the pool table if there is an invisible quarter circle barrier around one of the pockets. This, more or less, is how we generally find out about the unperceived.

Thus, though P-Cubic's knowledge is limited by the form and representational function of its perceptual scheme, it is limited in a rather subtle and complex way, not in the simple way

---

[3] See Cummins, 1996, chs. 1 and 8, for more on targets.

contemplated by Hume. Hume gets his critical leverage by assuming that ideas are copied from impressions, the implication being that the only representational targets available to us are those available to the perceptual systems. This is more than empiricism about evidence; it is a doctrine about mental representation for which Hume has no defense whatever.

2. What is implicit in representational form and function.

There are two morals I want to draw from this Kantian thought experiment. The first has already been drawn, viz., that we *can* know what we cannot experience. This, obviously, is not news, but I think the discussion of P-Cubic illustrates the point in a way that exhibits nicely how radical empiricism can seem unavoidable if you assume a kind of representational monism. Unlike Cubic, P-Cubic deploys two schemes that have different targets and different forms. Kant's representational pluralism—percepts and concepts—gave him the wherewithal to avoid Hume's claim that it is impossible to know about what you cannot represent perceptually, but it is arguable that he did not fully exploit it.

The second moral, of more interest to me here, is this: *every* representational scheme generates representations with a certain proprietary form or format that defines that scheme and constrains both what can be represented and how the representations can be processed.[5] These constraints can and should be conceived as enforcing presuppositions about the scheme's target domain (the domain it is supposed to represent), assumptions that are implicit in the form of the scheme's representations in the way that V1-V5 are implicit in the form of the representations generated by Cubic. Notice that there are two correlative parts to this idea. One is that representational schemes are aimed at proprietary targets, and their normal uses presuppose that the representations the scheme generates represent those proprietary targets and nothing else. The other is that the formal properties that define the scheme constrain how its targets are represented in a way that can be thought of as enforcing presuppositions about the target domain. I am going to call this pair of ideas the *representational specialization thesis,* or just the *specialization thesis,* for short.[6] Here are three simple examples of what I have in mind.

Maps

Your standard paper road map assumes that its targets are street and intersection structures. Because the map itself is flat, it gets linear distances more or less wrong in more or less hilly terrain. Ordinary globes of the Earth involve a comparable assumption. Notice how form and function interact here. A floor plan whose target is a flat two-dimensional structure has the same formal properties as the street map, but is limited to targets that do not violate the assumptions about distance built into the format. Stairs are not in the target domain, except where they begin and end, hence do not engender forced distance errors.

Notations for Sound

Standard musical notation is useless for representing most sounds: speech, the sound of a construction site, the sound of a receding train whistle. Attempts to represent any of these with standard musical notation will result in forced error.

---

[5] The words 'form' and 'format' are used here as a generalization of what is generally called syntax. Syntax is what we call the form or format of a structured propositional representation—e.g., a sentence.

[6] A precursor of this thesis is to be found in Haugeland, 1991.

This is because there are assumptions about the target domain that are implicit in the scheme—that pitch is not continuous, for example: only pitches in the chromatic scale are representable. Since the violin is not fretted, it can play pitch sequences that cannot be represented discretely. Moreover, the notation abstracts away from the difference between a violin and an oboe playing a middle C. This is no problem provided the music meant for the violins is given to the violin section and the music meant for the oboes is given to the oboe section, and the two sections are tuned together.

Phonetic writing of the sort to which there is a ready transcription to what you are now reading represents speech as a sequence of word length phonetic sequences. It is hopelessly inaccurate if used for representing any sound other than speech. The system simply assumes that everything is a phonetic element of speech—typically, though not necessarily, in a particular language. Those are its proper representational targets. The result is an abstract normalization of the actual wave forms—exactly what is needed for the sound to formant mapping required for communication.

Palette Systems

These are familiar from the graphics programs available on every computer. Presentation software, for example, allows a certain set of standard shapes: circles, lines, rectangles, callouts, etc., which can be sized and (sometimes) rotated. The frustrating assumption is that everything you want to draw is some relatively simple combination of these shapes at various sizes and rotations. These schemes also assume that their targets are two or three-dimensional. There are also specialized pallet systems for such things as flow charts and circuit diagrams, their intended targets and attendant assumptions being obvious from their functional designations. Their formal properties specializes them in ways that make them mostly useless for targets not in their intended domains.

The specialization thesis should be more or less obvious from the fact that we have developed a large variety of representational systems for different representational targets. No one would attempt to represent a Beethoven sonata with the resources of map-making or even graphed waveforms, nor would one attempt the architectural plans for a new building in musical notation. Yet the idea persists that target neutral representation is somehow possible; that language (or symbol strings) in particular, is somehow target neutral and hence fundamental. Thus, while my demystified version of the synthetic a priori may make sense to empiricists, it may seem to be of little importance: if there is target neutral representation, then the specialization thesis is false in its full generality, and the phenomena it emphasized can be dismissed as a curiosity that has no deep philosophical implications.

## 3. Target Neutrality.

There are a variety of routes to the idea of target neutrality. The most influential, I think, is the idea that maps and musical notation are just special purpose “languages”, or that all non-linguistic systems must be, or can be, linguistically interpreted. Linguistic imperialism does not, by itself, yield target neutrality, but it makes it more plausible by dismissing all of the obvious cases of representational specialization as mere differences in language. Specialization, as the example of Cubic makes clear, is invisible from inside a single scheme system. A system using an exclusively propositional scheme could allow for the proposition that its scheme is specialized. But it couldn't, of course, represent a non-propositional target. Such targets would be as inaccessible as colorless objects are to Cubic.

A more direct route to target neutrality is the idea that language is, potentially anyway, universal, limited in its representational power only by its currently limited but extendable lexical resources.[7] One can, after all, refer to a Beethoven sonata, or specify each of the notes in an extended linguistic description, which, though cumbersome, would, it seems, specify the same piece of music. And one can pixelize a picture and render it as a set of equations, or a set of verbal expressions (the RGB value of the third cell in the second row is r,g,b). To understand what is wrong with this kind of thinking we need to understand the distinctions between representing on the one hand, and reference and encoding on the other.

## 4. Reference, Representation and Encoding[8]

I have been urging the thesis that representational schemes are inevitably suited to some targets to the exclusion of others; that they are *representationally specialized.* This rather obvious fact is obscured by two phenomena that are often confused with representation:

1. Reference: One can *refer* to just about anything.
2. Encoding: One can use representations of one kind to *encode* representations of another kind, and hence to carry information about other targets indirectly.

*3.1 Reference.* Representations, as I understand them here, share structure with the things they accurately represent. Thus, an accurate street map of a city shares geometrical structure with the network of streets it maps. Sentences, on some views, share structure with the propositions they express, and, hence, can be said to represent those propositions. Terms, however, do not share structure with the things to which they refer. The phrase 'the street and intersection structure of Chicago' does not share geometrical structure with its referent.[9]

P-Cubic, as we have imagined it, cannot represent colorless objects, but it can refer to them. It can, therefore, theorize about such objects—invisible objects—and gather indirect evidence for their existence. Imagine, for example, that, every now and again, the balls on a pool table behaved as if they had collided with an invisible ball. A little Newtonian mechanics, and the judicious use of barriers, could create an overwhelming case for such invisible balls. (Think of Brownian motion!) Lacking a sense of touch, P-Cubic could not experience colorless objects as colorless, but it could certainly formulate and confirm propositions about such objects. This fact, while important, has no power to undermine the specialization thesis unless you think that reference is representation.  I do not want to argue over the word, but I do want to call attention

---

[7] Syntax, one might have thought, at least limits the representable *propositions*. And so it does. But look what happens in the interesting cases, e.g., modality. Friends of modality will say that non-modal languages are specialized to non-modal propositions. Enemies of modality will insist that there are no modal propositions to represent. No non-controversial instance of specialization emerges. And yet one might say, with more than a little irony, that it is a priori for those cleaving to a Quine-approved syntax that there are no modal propositions.

[8] For more on encoding, see "Systematicity and the Cognition of Structured Domains," Ch. #, this volume.

[9] The tendency to confuse reference with what I am calling representation, and hence referential content with representational content, makes it seem that all representation must be linguistic in character.

to a distinction between the kind of relation a street map of Chicago has to the street and intersection structure of the city, and the kind of relation that the phrase ‘the street and intersection structure of Chicago’ has to the street and intersection structure of Chicago. I’ve written extensively on this elsewhere (Cummins, 1996), so I won’t belabor the point here. All that is required in the current context is the fact that a successful reference to the street and intersection structure of Chicago need not carry any information about what that structure is. It is possible, of course, to describe, even specify, that structure linguistically, but that is not a case of representing that structure, but of *encoding* it.

3.2. *Encoding.* A transparent example of encoding is the use of pixel matrices to encode pictures and hence the targets they represent. A *symbolic specification* of a binding of RGB values to cells in a matrix is itself a scheme the proper targets of which are cell coordinates and RGB values. But the representations it *encodes* are not (typically) representations of cell coordinates and RGB values, but of objects in space. These objects might themselves be representations. One might digitally encode a picture of a musical score, for example. When the result is displayed, i.e., decoded into a picture, one looks right through the picture: a good picture of a score *is* a score, and a good picture of a map *is* a map. A good picture of a statue, on the other hand, isn’t a statue, and does not represent what the statue represents, though it provides some good clues. Digitized text, when displayed, yields token sentences that represent propositions. It is tempting, but mistaken, to think that a symbolic representation of pixel values is itself a representation of objects in space (digitized picture) or of propositions about something other than pixel values (digitized text). But these cases involve what John Haugeland once called a dimension shift (Haugeland, 1978), i.e., a shift in the targeted domain of representation. A picture of your aunt Tilly’s face represents her visual appearance in two dimensions—i.e., a particular two-dimensional projection of her facial topography; the pixel specifications are about pixels and their values. The two systems are semantically disjoint, except for the bizarre case in which one digitizes a text (or picture of a text) that specifies pixel values.

Another familiar example of encoding is the use of Gödel numbers to encode symbol strings.[10] Imagine a system that takes as input a sentence represented as an ordered list of words and returns its phrase structure, represented as a labeled bracketing of the underlying formants.

$$\langle\text{Wilbur loves the toy}\rangle \rightarrow \langle[[\text{Wilbur}]^{NP}[[\text{loves}]^{V}\ [[\text{the}]^{D}\ [\text{toy}]^{N}]^{NP}]^{VP}]^{S}$$

Corresponding to this is a system that takes as input the Gödel number of a sentence and returns the Gödel number of its labeled bracketing. On the assumption that the function that takes sentences onto their phrase structures is computable, it follows that there is a purely arithmetic function—call it g-parse—from the Gödel number of a sentence to the Gödel number of the corresponding labeled bracketing.[11] If we did such a calculation in standard decimal notation, for example, we would be generating references to integers, not to sentences or words. Moreover,

---

[10] Gödel numbering: To encode a string of words, e.g., ‘Mary loves John and Mary’, encode each word as a natural number and its position in the string as an prime exponent, the $n^{th}$ prime corresponding to the $n^{th}$ position. Let ‘Mary’ = 4, ‘loves’ = 1, ‘John’ = 3 , ‘and’ = 2. Since ‘Mary’ occurs in the first and fifth positions, we get $1^4$ and $7^4$. ‘loves in position 2 yields $2^1$. John in position 3 yields $3^3$. ‘and’ in position 4 yields $5^2$. The string is encoded as the product of these numbers: $1^4$ X $7^4$ X $2^1$ X $3^3$ X $5^2$ = 2441. Since every number has a unique prime factorization, factoring a number gives us the words and their positions (assuming the exponents have been assigned to words).

[11] g-parse is, for example, expressible as a composition of plus, times and minimization.

the calculation we did would depend on algorithms defined over decimal numerals, and, at a more abstract level, over integers, a kind of processing that has no application to words and sentences at all. The scheme would assume that every object in the domain has a unique prime factorization, and (at the level of notation) that shifting left implements multiplication by ten. All of this *specialization*—proprietary targets and proprietary processing—becomes invisible if we think of the system using Gödel numbers as "doing the same thing" as the system that uses words and labeled brackets. But it is not doing the same thing, as is evident from the fact that g-parse is insensitive to whether or not the input number encodes a grammatical sentence, or even whether it encodes a sentence or words at all. (Remember: no decoding takes place, just arithmetic.) It is thus not built into the form of this scheme that everything in the domain is a word or sentence or syntactic category.

Because encodings do not, by themselves, presuppose what is built into the form of the schemes they encode, there is a sense in which they allow one to float free of representational specialization, at least to float free of the presuppositions of the scheme *encoded*. Encoding allows for contents that cannot be represented or referred to in the encoded scheme. We could allow for a variant of P-Cubic that encodes its spatial representations by binding locations (cells) to objects, and colors to cells. It is natural for us to think of such a scheme as symbolic and propositional—as a set of equations. But it need not work that way. We might have visual cells, object cells and color cells all oscillating (activation or intensity or…) at distinctive frequencies. An object is encoded by the set of visual cells oscillating at the same frequency as a given object cell. A cell is colored red if it oscillates at the same frequency as a given color cell—the "red" cell. This would allow for uncolored objects (oscillation corresponding to no color cell). It would not allow for objects whose representations are not uniformly colored, though it would allow for masses of variegated color that are not objects. As this example makes clear, encodings impose their own constraints on the encoded system's target domain, and on its own domain: the representations of the encoded system. The presuppositions of the encoded system may be violated, but not those of the encoding system.[12]

Encodings, of course, need not be encodings of another representational scheme. Encodings can themselves be encoded, and so can the world. Thus, a variant of Cubic might encode objects in space directly via the oscillation scheme sketched above, and, of course, a symbolic encoding—a set of equations—could encode those oscillation patterns. Moreover, P-Cubic need not rely on reference and encoding to cognize a world that does not satisfy V1-V5. If we add to P-Cubic's arsenal some representational resources that are both non-propositional and non-perceptual, we open up the possibility of not merely referencing or encoding, but of *representing* the mechanics of objects in space in a way that abstracts away from color, size and shape. This, of course, is exactly what, e.g., vector diagrams of the objects on a pool table do. Such diagrams do not represent everything about the balls on the table: they represent only relative positions of centers of mass and their momenta. This, of course, is exactly the kind of abstract representation that is needed to get the mechanics right.[13] The example illustrates the

[12] It is important not to confuse linguistic description with encoding. Describing a picture is not the same as encoding it. But we can use language to specify encodings, as in the discussion of Gödel numbering in the footnote above.

[13] If you took a snap shot of a pool table with a "Newton Camera", it would show a bunch of coplanar arrows.

possibility of accurately representing the objective structure of things one cannot experience as having that structure.[14]

When we take all of this into account, the simple picture of being epistemologically limited or even trapped in the presuppositions of ones representational scheme (or language) dissolves away. Only creatures whose resources are limited to a single language, or a single representational scheme, or a single encoding scheme, or creatures with a variety of resources but with no overlap whatever in the targets of those resources, need be limited in the way empiricism and Kant imagined. Figuring out in advance—a priori, as it were—what is knowable by a complex creature with multiple representational and encoding schemes at its disposal, and a language to help manage them and to invent new ones, is probably a fool's errand. Limitations perhaps there are, but they are difficult to discern, and nothing of the kind imagined by Kant, let alone Hume, is in the cards.

## 4. Representational Specialization and World Making

We get a different perspective on essentially the same point by considering a further Kantian sounding consequence of the target specificity of representational systems, viz., that a given system can seem to create its targets. A speaker of English (or whatever) hears words, sentences and meanings that a non-speaker does not. One needs to be in the grip of a theory that draws a sharp distinction between perception and inference to follow Berkeley and say that the non-speaker has his ears and the use of them as well as the speaker, so the difference cannot be in what is *heard*. Inference, very likely of a non-propositional sort, there surely is in speech perception, for it is quite obvious that speakers hear things that non-speakers do not. And it is equally clear experimentally that non-speakers hear things—e.g., certain phonetic contrasts—that speakers do not (e.g., Dupoux, et. al., 2001). Reflection on these and similar experimental facts can make it seem plausible that words, sentences, and meanings are all somehow *created* by the knowledgeable listener. This sounds way too Kantian to the contemporary analytic ear, however. It seems preferable to say instead that the words are there in the speaker's mind, and the hearer (unconsciously) infers them from the auditory signal. But how different are these two, at bottom? The inference in question, after all, must be constrained by a set of forced options, an hypothesis space, if you like, that limits the possibilities to known words, or anyway to pronounceable syllables, in a known language.[15] The upshot seems to be an analogue of Cubic for language perception: for a system specialized for speech perception in a given language, everything that cannot be treated as legitimate speech elements will simply be treated as noise. For such a system, the whole auditory world is indeed a text. A familiar extension of the point is to argue that conceptual schemes create their own objects, or even that they create the objects of the perceptual schemes they are thought to regiment.[16]

There is a sense in which representational specialization, especially perceptual representation, does create the world we live in. Walk down a street in a city that speaks a

---

[14] This representation is not inaccurate because it does not represent color or shape, for its target is simply the relative positions and momenta. Inaccuracy in this instance would be error in position or in the magnitude or direction of a momentum vector.

[15] A foreign word or phrase here and there is ok, but fluent speech that changes language every word is impossible: the syntax simply won't allow it.

[16] See Haugeland, 1998, chs. 10 and 13. Dogs, Haugeland would claim, cannot see tennis balls *as such*. For a discussion, see, chapter # in this volume.

language you do not know. There is speech and there are signs everywhere, but it is all just *stuff*. No words heard or seen, at least not as such. But what *is* the *stuff*? What are the lines and shapes and colors and pitches and volumes and timbres? Wouldn't these disappear *as such* too if we took away the resources for representing them? A familiar argument for nativism starts here: some representational resources must be innate, otherwise none could be learned, for there would be no *experience-of* to learn from. It is not the nativist response I want to pursue, however, but the transcendental idealist response: no world to know without thought, no thought without representational resources, no representational resources without the a priori limitations imposed by the form of those resources. My point about this is simple: Even if all of this is true, it does not have the consequences it appears to have for a representationally pluralistic system, as we have seen. The source of the idea that representational specialization has serious implications for knowledge and world making does not derive from representational specialization itself, but some form of representational *monism* explicit or implicit in other doctrines. It is implicit in the doctrine that all of our ideas are copied from our impressions, and that our impressions can be catalogued in a short list of proper objects of perception: color, shape, size, texture, pitch, volume, timbre, etc., qualities that can be perceived *as such*. And it is explicit in the doctrine that all representation resolves somehow into language, so that the world turns out to be a set of propositions or a "text".

Think again of P-Cubic, and imagine that P-Cubic's propositional resources are limited to first order combinations of terms whose referents are given *as such* in perception. Even that, as we saw, does not limit P-Cubic to the presuppositions of its perceptual system. But it does limit P-Cubic's propositional targets to those expressible in a first order language with a limited empiricist lexicon. These limitations are evidently the product of two factors: representational specialization and a doctrine about the limitations of P-Cubic's lexicon. The latter assumption is not a matter of form, but of function. It is the assumption that the function of all primitive terms in P-Cubic's language of thought is to refer to "sensible qualities", an assumption that has no non-circular justification that I know of. For the justification of the limitation is always that the only knowable propositions are among those allowed by the restriction. You cannot argue to an ontology from a premise about limited representational resources, then argue for limiting those resources on epistemological grounds. Our resources do limit us, but they are what they are, not what some antecedent restriction on what is knowable implies that they ought to be. Moreover, as I hope the discussion above has made clear, it is no simple matter to say what limitations are actually faced by creatures with a plurality of schemes and a plurality of strategies for exploiting them.

References

Ayer, A. J. (1936) Language, Truth and Logic. London: Victor Gollancz Ltd.

Clapin, H (2002) Philosophy of Mental Representation, Oxford: Oxford University Press.

Cummins, Robert (1996) Representations, Targets and Attitudes. Cambridge, MA: MIT Press.

Dupoux, Emmanue; Pallier, Christophe; Sebastian, Nuria; Mehler, Jacques (2007) A destressing "deafness" in French? Journal of Memory and Language. Vol 36(3), Apr 1997, pp. 406-421

Haugeland, J. (1978), 'The Nature and Plausibility of Cognitivism', *Behavioral and Brain Sciences*, 1, 215-26.

Haugeland, J. (1991) “Representational Genera.” Reprinted in Haugeland, 1998a, 171-206. First published in W. Ramsey, S. Stich and D. Rumelhard, eds., Philosophy and Connectionist Theory. Hillsdale: Laurance Erlbaum, 61-89.

Haugeland, J. (1998) Having Thought. Cambridge, MA: Harvard University Press.

Hume, D. (1748). An Enquiry Concerning Human Understanding. Reprinted in Hume’s Enquires, J. A. Selby-Bigge, ed. Oxford: Clarendon Press, 1963.