

## **Taking Embodiment Seriously: Non-conceptual Content and Computation**

Ronald L. Chrisley  
School of Cognitive & Computing Sciences  
University of Sussex

### **Abstract**

The development and deployment of the notion of pre-objective or non-conceptual content for the purposes of intentional explanation of requires assistance from a practical and theoretical understanding of computational/robotic systems acting in real-time and real-space. In particular, the usual "that"-clause specification of content will not work for non-conceptual contents; some other means of specification is required, means that make use of the fact that contents are aspects of embodied and embedded systems. That is, the specification of non-conceptual content should use concepts and insights gained from android design and android epistemology.

### **1. Introduction**

A central idea in cognitive science is the that there is something that computers and brains do, but which other objects, be they rocks, trees, or bicycles, do not. There are several candidates for what that commonality between brains and computers is: they process information, they represent, they compute. It is hoped that by finding similarities between computers and brains, we might be able to use the techniques of computer science to help us demystify the mental activity that the brain supports.

However, leaving the postulated similarities at such an abstract level may limit the productivity of this approach. After noting the connections, if any, between what brains do and what computers in general do, it seems a good idea to look for more specific connections between human cognition and the activity of particular kinds of computers, computers that share at least some of the purposes and functions of the brain. Recent trends in cognitive science suggest that we have already come close to exhausting the insights that can be gained from comparing human cognition to the activity of disembodied, formal computational systems such as theorem provers and expert systems. In particular, it seems wise to enrich the study of human cognition with an emphasis on computers that control robots which perceive and act in a real-time, real-space environment. Indeed, the idea seems to go beyond the mere computer/brain analogy, but extends the correspondence to the robot/body, and further: it is often claimed that the respective environments of these systems should be included within our theoretical view if we are to achieve a deeper understanding of cognition. Although this class of computational/robotic systems surely includes more than what is typically meant by the term "android", it seems clear that all androids fall into this class.

Likewise, some philosophers of mind and language have rejected traditional, logic-based, formal notions of thought as too static and coarse-grained to account for the aspects of cognition that involve perception and action, or any kind of cognitive dynamics, such as learning, development, or concept formation. But rather than just stopping with the rejection

(an *eliminativist* position), these philosophers have suggested new notions of the content of mental life, notions of content which can allow for modes of intentional interaction that is not mediated by fully formed concepts. These notions of *non-conceptual content* vary from author to author, but they have in common an intended purpose: to permit the explanation of a greater range of intentional phenomena, to widen the scope of intentional explanations from the mere logical to the fully psychological.

There are (at least) two connections between these recent developments, in complementary directions.

First, some of the arguments that establish the need for non-conceptual content for the explanation of human cognition apply also to the case of artificial computational/robotic systems operating in real-time and real-space. That is, a proper understanding of android epistemology will require these new notions of non-conceptual content and non-conceptual mentality.

Conversely, the development and deployment of the notion of non-conceptual content requires assistance from a practical and theoretical understanding of computational/robotic systems acting in real-time and real-space. In particular, the usual "that"-clause specification of content will not work for non-conceptual contents; some other means of specification is required, means that make use of the fact that contents are aspects of embodied and embedded systems. That is, development and deployment of the notion of non-conceptual content will use concepts and insights gained from android design and android epistemology.

This chapter addresses this second connection between non-conceptual content and embodied, robotic computation. Section 2 explains what is meant by content, and how it is used in explaining the behaviour of intentional systems. It also makes clear the distinction between conceptual and non-conceptual content. Section 3 shows why standard means of content specification ("that"-clauses) will not work for non-conceptual content. Section 4 proposes some alternative means of specification, and points out that all of the plausible ones employ notions of the embodiment and embeddedness of an intentional system, notions which are made precise via a detailed understanding of the relationships between computation/robotic properties and intentional ones: an understanding, that is, of android epistemology.

## 2. Content-Based Explanation

The intentional explanatory strategy that dominates cognitive science typically understands psychological states in terms of attitudes (belief, desire, knowledge, intention, etc.) toward contents (that there is a door ahead, that  $2 + 2 = 4$ , etc.); such attitude/content pairs are appealed to in intentional psychological explanation. For example, one might explain why a robot opened a door (i.e., show that the robot's opening the door wasn't just an accident; if circumstances had changed slightly -- if the door were one foot over to the right, say -- the robot would still have opened the door) by claiming that the robot *intended to open the door*; one could explain the possession of this intention as being the result of the robot's *desire to be in the next room* and its *belief that opening the door will help one get into the next room*. I will call explanations that appeal to such attitude/content pairs *content-based explanations*.

This notion of attitude/content pairs has played a key role in artificial intelligence (AI), but the emphasis there has been on one attitude in particular — knowledge — since intelligent action is seen to require knowledge of some sort. Correspondingly, within the field of AI

there has been a great deal of interest in the nature of knowledge: how it can be manifested in programs or robots, and how it can be acquired or transferred. Although this line of inquiry is one approach to the goals of understanding the epistemology of artificial agents, it is also true that this particular emphasis on knowledge has resulted in a gap in current approaches to AI: not merely a (relative) lack of understanding of attitudes other than knowledge, but also, and perhaps to a greater extent, the absence of a proper understanding of the notion of content, and how it relates to the explanation and design of intelligent systems.

## 2.1 What is Content?

Content is the way the world is presented to a subject of experience. It is convenient to explicate what is meant by content by appealing to the notions of "that" clauses, information, and truth-conditions, but this is done by showing how content is different from those notions, not by identifying content with any of those three:

(1) "That" clauses, such as "that there is a door ahead" and "that  $2 + 2 = 4$ " do indeed specify contents, but it would be wrong to define content as that which is specified by such clauses. First, "that" clauses can only specify complete, propositional contents; yet expressions that sentences comprise (such as names and predicates) carry their own, sub-propositional contents. Second, and more importantly, it is argued in section 3 that there are propositional contents that cannot be specified using "that" clauses.

(2) Like information, content can be understood to be "carried" by states, representations, symbols, and expressions. But information is typically understood as something that cannot be false: a state can only carry the information that  $x$  is  $P$  if it actually is the case that  $x$  is  $P$ . This is notoriously not so with content: sentences and beliefs can be false, and the content of my experience as of an oasis in front of me can be an illusion.

(3) Having some kind of norm of correctness (truth-conditions, say) is necessary for the possibility of falsity, and is a characteristic feature of content. However, one cannot look to truth-conditions alone to take the place of content in psychological explanation, since, as Frege pointed out, there are propositions that have identical truth-conditions, yet we might assent to one and dissent from the other (e.g. "the Morning Star is Venus" and "the Evening Star is Venus", given that the Morning Star is the Evening Star); the propositions might have different cognitive significance, different content.

Content, then, is a way of taking the world to be, and has two essential features: characteristic norms of correctness and a characteristic pattern of cognitive significance.

## 2.2 Conceptual and Non-Conceptual Content

I join others (Crane, 1992; Cussins, 1990; Davies, 1990; Evans, 1982; Haugeland, 1991; Peacocke, 1992) in arguing that a distinction should be made between conceptual and non-conceptual contents.<sup>1</sup> For several reasons, much of the work in AI has concentrated by default on the case of conceptual content, but there is reason to believe that understanding non-conceptual content is essential to understanding (and therefore to designing) intentional systems in general; see (Cussins 90) .

<sup>1</sup>I should point out that none of the cited authors characterize non-conceptual content in exactly the same way, nor does my notion exactly agree with any one of theirs. But the differences are largely irrelevant for the purposes of this chapter.

Some ways of representing the world (contents) are objective or near-objective, some are not. A way of representing some aspect of the world is objective if, e.g., it presents that aspect of the world as something that could exist while unperceived. Strawson (Strawson, 1959) maintained, as did Evans (Evans, 1982) after him, that at least in the case of thinking about spatio-temporal particulars, truly objective thought is manifested in the possession and maintenance of a unified conceptual framework within which the subject can locate, and thus relate to any other arbitrary object of thought, the bit of the world being thought about.

If this is a correct understanding of objective thought, then it has important implications for the understanding of pre- (or non-, or sub-) objective representation. Pre-objective representation involves contents that present the world, but not *as* the world, not as something that is or can be independent of the subject. An infant's early perceptual/motor interactions with its environment is a plausible example of the presence of a pre-objective mode of thought. The infant tracks an object (thus suggesting that there is some intentional relation between the infant and the object) when it is perceptually occurrent, yet when the object is occluded from view, the infant loses interest in the object (perhaps dropping it), and is in fact startled if the obstruction is removed to reveal the object. Since the very notion of an *object* essentially involves the notion of something that can exist even though occluded, the infant is not thinking of the object *as* an object. The contents of the infant's thoughts concerning the object do not present the object as something objective, as something that could exist while unperceived. According to the Strawsonian/Evansian line I am taking, then, the infant's lack of objectivity must be manifested in the lack a unified framework of thought: the infant is unable, in general, to locate objects in such a framework. I will call such pre-objective contents *non-conceptual* contents (*NCCs*); *conceptual* contents, on the other hand, are objective.<sup>2</sup>

It is a consequence of this way of understanding the conceptual/non-conceptual distinction that conceptual contents will necessarily be systematic, will meet Evans' Generality Constraint (Evans, 1982, p 104):

(GC) For any conceptual contents (ways of thinking of properties) **F** and **G**, and any conceptual contents (ways of thinking of objects) **a** and **b**, if a subject knows what it would be for **a** to be **F** and for **b** to be **G**, then it must know what it would be for **a** to be **G** and for **b** to **F**.

The constraint (GC) is a direct consequence of the necessity, for there to be objective, conceptual thought, of a unified framework within which to locate all properties and particulars.

Non-conceptual contents, on the other hand, are not systematic. In the case of NCC, the mode of thought is pre-objective; a unifying framework is not present, and there are, therefore, properties and particulars which cannot be related in the proper way. The idea of non-conceptual content, then, implies that one can represent the world with proto-concepts that do not universally recombine with all other possessed proto-concepts.

---

<sup>2</sup>The extent to which infants — even neo-nates — are unable to conceive of existence unperceived is a hotly debated topic, with received opinion currently swinging in favour of objectivity at a very early age, if not at birth. Such empirical details are irrelevant here, since the example of the infant still illustrates what is *meant* by NCC, even if it is at odds with what we know about the intentional capacities of infants. Furthermore, what little justification that has been offered for the stampede away from seeing infants' abilities as pre-objective seems to be based on some philosophical misunderstandings.

For example, consider an infant which cannot, as before, think of a particular object (a glass, say) as existing unseen, but *can* represent its mother as being behind, out of view (on the basis of hearing her voice or feeling her arm, say). The contents of such an infant will violate the Generality Constraint, since the infant may be able to think (something like) *glass in front of me* and *mother behind me* but not *glass behind me*. The infant's contents are not fully objective, and are therefore non-conceptual. To ascribe conceptual content to the infant in this case would mis-characterise its cognitive life, and would not allow prediction or explanation of the infant's behaviour.

An important constraint on content ascription is what I will call the *Possession Principle*: a subject may only entertain a content C if the subject possesses all of the concepts (if any) that C comprises. Thus, *glass in front of me* would not be a content that the infant could entertain, since the infant's failure to meet the Generality Constraint demonstrates that the infant does not possess one of the constituents of the content, the concept *glass*.

### 2.3 Other Ways a Content May Fail to be Conceptual

However, the manner in which contents fail to be objective is richly textured. A content may manifest its non-conceptuality by failing to respect any of a number of conceptuality constraints. A non-exhaustive and non-exclusive set of such constraints might include, in addition to (GC), the following properties of conceptual content:

(SP) it must have subject-predicate structure;

(RP) if the content is to be a way of thinking about an item, the subject must know which (in the sense required by what Evans calls "Russell's Principle") item is being thought about;

(MC) if the content is to be a way of thinking about an item, the subject must be able to think of the same item in a number of other ways;

(EU) if the content is to be a way of thinking about an item, the subject must be able to think of the item as existing unperceived, as something for which the qualitative/numeric distinction applies, or as something which can be re-identified;

(PE) if a subject is capable of taking the attitude of belief toward the content, then it must be able to entertain the possibility that its belief is false; it must have the concept of belief to have any beliefs.

Thus, non-conceptuality could be manifested in the failing to meet of any of these conceptuality constraints. It should be emphasised that these constraints are offered as examples only. It might be that some of them are not required of all conceptual contents; or, it might be that some of them are required of *all* contents, and thus failing to meet them is not a way for a content to be non-conceptual. Nevertheless, these examples will be of use in illustrating a proposed alternative means of content specification, in 4.1 below.

## 3. The Inadequacy of Standard Specifications for NCC

### 3.1 Pre-Conceptual Linguistic Use Specifications

In order for a theory of intentional action to be able to appeal to specific contents in its explanations, it must have a means of canonically specifying those contents, a means of specifying them according to their essential properties, such as their truth conditions or their

cognitive significance. For example, one can specify a content by the phrase "the content toward which subject A took the belief attitude exactly 10.3 seconds ago", but this would not be a canonical specification, since it does not pick out the content it does in virtue of the content's essential properties, but rather its accidental ones. A standard means of canonically specifying contents is what I will call the "linguistic use" means of specification: providing an expression in English (or other natural language), usually preceded by "that", with the same content as the one to be specified (e.g., "The content of my belief is *that the object on the table is a computer*" or "The last bit of register A0 being on means *that a message is in the input buffer*").

Of course, almost any proposed means of specifying content will use language in a more general, and conventional, sense of the word "use" than the one I am employing here. However the following criticism of linguistic use methods are not directed towards proposed forms of specification that use language in this broad sense. The expression "linguistic use" is meant to be a technical one: I mean to include under the term only those means of specification that pick out contents exclusively in the manner mentioned in the definition above: viz., by providing an expression in a natural language that has the very same content as the one that is to be specified. In this narrower, more technical sense, means of specification may use language in the general sense, and yet not be subject to the negative conclusions in what immediately follows.

Although the linguistic use means of specification might work well for conceptual contents, there are several reasons why one might think that it is not adequate for the specification of sub-objective, non-conceptual contents (NCCs).

First, a direct claim can be made: language is itself conceptual – all the contents involved in using language are fully objective – and therefore linguistic use specifications can only specify conceptual contents. Linguistic use specifications are what has been called elsewhere (Cussins, 1990, especially pp. 382 ff; Peacocke, 1986, especially p. 17) *conceptual* specifications of content: specifications that are made in such a way as to require a subject to possess the concepts used in the specification if that subject is to be able to take an attitude toward that content. Thus linguistic use specifications, employing conceptual language, will not be able to specify the contents of, say, an infant's or animal's psychology, since such specifications would require the infant or animal to possess concepts which it in fact lacks.

In spite of the strong intuitions behind this line of thought, there are reasons why it might be more illuminating to establish the incompleteness of linguistic use specifications by a means other than one which relies on the principle that all language is conceptual; for one thing, many would want to deny that language is entirely conceptual. What I will do, then, is split linguistic use specifications into two types: *purely descriptive*, and *indexical*.. I will argue that specifications of neither type can specify NCCs. First, a means of specification, in order for it to be of use in a scientific theory, must specify NCCs canonically, which rules out descriptive linguistic use. Furthermore, content specifications must be context-independent, which rules out indexical linguistic use. Thus some means of specification other than linguistic use is required.<sup>3</sup>

<sup>3</sup>Peacocke (Peacocke, 1990) shows the insufficiency of standard specifications for a restricted class of NCCs, *perceptual demonstrative* contents; and develops (Peacocke, 1989; Peacocke, 1992) an alternative means of specification, *scenarios*, for these contents. But the goal in this chapter is to establish the insufficiency of linguistic use specifications for NCCs in general (or at least for a broader or distinct class of NCCs than does Peacocke); for these NCCs the scenario means of specification will not work (nor was it intended to).

### 3.2 The Inadequacy of Descriptive Linguistic Use

Cussins (Cussins, 1990) has brought together some insights from (Evans, 1984) and (Perry, 1979) that can serve as an argument against the possibility of using descriptive (i.e., non-indexical) language to specify non-conceptual contents.

Perry shows that there are contents, constitutively linked to perception and action<sup>4</sup> (e.g. the contents of one's "I" thoughts), that are not equivalent in terms of cognitive significance to any contents specified, in a purely descriptive manner, by the linguistic use method.

However, one *can* give a linguistic use specification of the contents of one's "I" thoughts, but only if one employs indexicals, as in the ascription: "RC believes: 'I am spilling sugar all over the supermarket floor'". One cannot use a non-indexical specification, as in "RC believes 'the person named RC is spilling sugar all over the supermarket floor'", since it specifies a content that is distinct from that of the first-person thought I would normally have in that situation, as can be seen by the differences in the two contents' connections to action (due to amnesia, I might think in the latter case "Well, the person named RC had better clean it up" and go on my way, whereas in the former case no amnesia could get me to think that it was anyone else's mess but mine). In order for the belief *the person named RC is spilling sugar all over the supermarket floor* to have any implications for my action, it must be supplemented by the belief *the person named RC = I (me)*. The belief *I am spilling sugar all over the supermarket floor* requires no such further identification; its connections to action are direct, un-mediated.

The application of Perry's insight to the case of NCC is direct: if any NCCs are directly connected to perception and action in the way that "I" contents are, then Perry's arguments establish that such contents cannot be specified by means of descriptive language use. One line of reasoning that leads one to conclude that all NCCs *are* constitutively connected to perception and action is the following. As observed before, NCCs are sub-objective in virtue of the fact that they do not enable the bit of the world being thought about to be integrated into a unified framework of particulars and their inter-relations. It is this lack of a framework which restricts sub-objective thought to contents that are essentially linked to perception and action. The idea employed here is that indexicality is the starting point; contents that (merely) have constitutive connections to action and perception are the basic case. It is only through the construction of a non-solipsistic conception of the world via some unified framework of particulars and relations that one's contents can display the kind of perception and action transcendence that is characteristic of descriptive modes of thinking. It is the very sub-objectivity of NCCs that allows the application of Perry's argument to their case: they cannot be specified by descriptive language use.

This understanding of NCC seems to agree with (at least one reading of) what Evans meant by non-conceptual content:

Let us begin by considering the spatial element in the non-conceptual content of perceptual information. What is involved in a subject's hearing a sound as coming from such-and-such a position in space?... When we hear a sound as coming from a certain direction, we do not have to think or calculate which

<sup>4</sup>Actually, Perry claims that it is the fact that a belief is a "locating belief" that makes its specification essentially indexical; I'm favoring here Cussins' analysis that it is a content's constitutive links to perception and action that requires non-descriptive specification for that content.

way to turn our heads (say) in order to look for the source of the sound. If we did have to do so, then it ought to be possible for two people to hear a sound as coming from the same direction (as 'having the same position in the auditory field'), and yet to be disposed to do quite different things in reacting to the sound, because of differences in their calculations. Since this does not appear to make sense, we must say that having spatially significant information consists at least partly in being disposed to do various things. (Evans, 1982, pp. 154-5)

This is very similar to Perry's way of characterising ways of thinking that are essentially linked to perception and action. Just as there is no "calculation" in the case of Evans' example of auditory content, there is no "calculation" in Perry's example of the first-person mode of thought: one knows, in an un-mediated manner, that such thoughts are directly related to one's own actions and perceptions. An identification with some descriptive mode of thought (e.g. *I (me) = the person named RC*) is not required for action.<sup>5</sup>

Another indication that NCCs are, like the contents of the indexicals "I" and "now", constitutively linked to perception and action is that if one attempts to specify such contents by means of linguistic use, then one tends to use indexicals in so doing.<sup>6</sup>

If what has been said is correct, then NCCs are indexical contents in Perry's sense, and therefore, like the content of conceptual first-person thoughts, cannot be specified by descriptive linguistic use. But there is reason to believe that non-conceptual contents, unlike the contents of conceptual first-person thoughts, cannot be specified by the alternative of indexical linguistic use, either.

This is so, if not because of the conceptuality of language, as discussed above, then for reasons related to the requirement that scientific theorising be context-independent (in a particular sense). That is, even if indexicals could, *per impossibile*, be used to specify NCCs via linguistic use (either because they do not have conceptual content, or because they can somehow linguistically specify a content that is not, strictly speaking, the content they carry, or because it is possible to devise new indexicals that introduce, in a non-systematic way, elements of the environment into the content being specified) such indexicals alone would be inadequate for the particular task at hand: a context-independent intentional science. Specifically, the function of the indexical is merely to call attention to other factors (subject, context, and their relation) so that a content may be specified. In such a case, all the individuating work is being done by those highlighted elements, not the indexical itself.<sup>7</sup>

<sup>5</sup>A word of caution: Evans' point should not be construed to be claiming that non-conceptual contents are somehow infallible, because of their direct connections to perception and action. The essential links can be inappropriate for the current situation, therefore yielding a false NCC: because of a reflection, the sound might be heard as coming from the right (with all the commensurate right-directed dispositions), when in fact the source of the sound is straight ahead.

<sup>6</sup>Another caution: though I am arguing that all NCC's are indexical, in that they are non-descriptively linked to perception and action, I am not claiming that the relation in the other direction is true (that all indexical contents are non-conceptual, or sub-objective); on the contrary, I think "I" has conceptual content (I might be wrong on this, as on anything else, but fortunately it would have no undesirable consequences to what I am arguing here if I were). Unlike (Cussins 1990, p. 391, n. 46), I do not feel justified in rejecting out of hand indexical linguistic use specifications of contents for a scientific psychology. Some indexical contents (the first-person, the present-tense) seem to be conceptual (at least enough to avoid the problems of context-dependence), and are thereby specifiable by indexical linguistic use. Thus, I am required to provide an argument (which I do) for the claim that indexical linguistic use specifications will not work for the case of non-conceptual contents (although *embedded* indexical specifications might succeed; see section 4.2).

<sup>7</sup>One might wonder: how is it that indexical linguistic use specifications seem to work in some cases, even though no systematic way of specifying various aspects of the context is at hand? The reply: in the case of indexical linguistic use



Thus, in order to specify the content, one would need more than mere linguistic expressions; one would also need an environment related to those expressions in order to allow those expressions to function, and thus carry content. Thus, even *indexical* linguistic use cannot be used to specify NCCs.

This conclusion is also supported by the following line of reasoning: given 1) NCCs are indexical; 2) the suggestion that the content of all linguistic indexicals can be reduced to "I" and "now" plus some descriptive component; 3) there might be organisms which entertain non-descriptively specifiable NCCs but do not possess the first-person mode of presentation (the content of "I"); then it follows that there will be no indexical linguistic use specification for the NCCs of such organisms.

These arguments agree with the conclusion of other writers (e.g., Peacocke, 1981, p 191): in specifying contents that are constitutively linked to perception and action, such as particular first-person modes of presentation, we cannot employ the content in question, but must refer to it instead. The task, then, is to find ways of referring to such modes of presentation that identify them not only uniquely, but canonically, as discussed above.<sup>8</sup>

#### **4. Alternatives to Standard Means of Specification**

To be frank, I don't yet have a fully worked-out alternative means of content specification. What the rest of this chapter will do, however, is describe some possibilities that are currently under consideration, and explain why they are at least plausible candidates for an alternative. This will not only serve to explicate many of the issues that have been discussed, and to pacify those "what else could there be?"-type worries that may be nagging some readers; it will also be seen along the way why one must take embodiment seriously if one is to be able to specify non-conceptual contents for an intentional science.

##### **4.1 Conceptual Subtraction**

One idea is: perhaps linguistic use fails only as a matter of technicality; perhaps some modification of it can overcome its limitations, while using the same, fundamentally non-embodied, approach. It seems such a modification would have to be something like the conceptual subtraction (CS) means of content specification. As the name might suggest, this method is similar in spirit to a pure conceptual specification, such as linguistic use. Nevertheless, and the above arguments against linguistic use specifications of NCCs notwithstanding, it seems possible that the CS method *can* specify non-conceptual contents, because it is distinct from pure conceptual specification in a crucial way. The CS method is an attempt to stay as close to our practice of linguistic use specification while throwing out the restrictions that make linguistic use inadequate. The problem with such purely conceptual specifications, as we have seen, is that they cannot specify sub-conceptual contents. Any attempts at specifying the content of, say, particular pre-objective experiences of an infant, would over-ascribe, in that such specifications would imply that the infant

---

specifications of conceptual contents (such as those that specify the mature first-person mode of presentation with "I"), the objective, context-independent nature of conceptual thought permits specification with only one extra contextual parameter: the person grasping the content in the case of the first-person, and the time of the grasping, in the case of the present tense.

<sup>8</sup>Of course, one should evaluate other proposed alternatives to the standard means of linguistic use, few though they may be, before concluding that another alternative is needed, but there is no space here for such a survey. Suffice it to say that other alternatives (possible worlds, possession conditions, proto-propositional specifications) are insufficient either because they ignore cognitive significance (are purely extensional), or they only are faithful to the case of cognitive significance in the same cases (i.e., conceptual, or non-conceptual with shared environment) that linguistic use is, which I have already argued is insufficient for a scientific psychology.

possesses abilities that it does not, in fact, possess. That is, they would violate the Possession Principle (cf. section 2). For example, specifying the content of the infant's belief as "that there is a glass in front of me" would imply that the infant possessed the concept of a glass, with its attendant concepts, not only *drinking*, *manufacture*, and *glass*, but also *object* and *location*; it would also imply that the infant's thinking about the glass adhered to the generality constraint, and supported the ability to think of the glass as something that could exist unperceived. This would invite the theorist to make false predictions about, and would disallow correct explanations of, an infant's behaviour.

The idea behind CS specification is to proceed with a conceptual, linguistic use specification, but also to tag the implications of that specification to which one does not wish to be committed; that is, to start with the conceptual content, and then subtract out properties of the conceptual content (such as "meets the generality constraint" or "supports the idea of existence unperceived") which the content to be ascribed does not possess. Then there will be no over-ascription of abilities, and therefore no false prediction or inaccessible explanation.<sup>9</sup>

In order to be able to employ the CS method, one must first capture all the different implications an ascription of a conceptual content carries with it. This will involve both a cataloguing of the general requirements for all conceptual contents and all concepts, such as the Generality Constraint, and a listing of the particular requirements for each individual concept. One might end up with a list like the one in section 2 for the general conceptual requirements, with the addition of something like:

Particular conceptual requirements:

*bachelor:*

- (1) organism must also possess the concept *unmarried*
- (2) organism must also possess the concept *male*

...

*drinking glass*

- (1) organism must also possess the concept *liquid*
- (2) organism must also possess the concept *drinking*

...

Etc.

It is very important to note that the above, as well as the list made in section 2, is only meant to serve as a toy example of the enumeration of constraints for the CS method, and not as a specific proposal for what these constraints should be. Also, the capturing of the commitments need not proceed via enumeration; the catalogue will undoubtedly employ quantification. It might also be recursive, in that conceptual requirements might themselves have further conceptual requirements, such that a content might meet some of requirements for, say, (RP), and not others.

<sup>9</sup>The basic idea of the CS means of specification seems to have been independently reached by Colin Allen, see (Allen, 1992).

Once this enumeration of conceptual requirements is in place, specifications of sub-conceptual contents in ascriptions would be possible:

The content of the infant's belief is *that there is a drinking glass [-GC, -EU, -1] within reach*

where the qualifiers within brackets after a concept indicate in what ways that part of the content fails to be conceptual.

In order for any alternative specification to succeed, several conditions must be met. Any particular application of the method must indicate at least one content, at most one content, and, as discussed before, it must indicate the content canonically.

In the case of the CS method, the first condition prompts one to wonder: how can one be sure that a subtracted content is actually a content at all? One suggestion<sup>10</sup> for a criterion for an abstract entity to be a content is that it be able to help rationalise a subject's behaviour by serving as a premise in practical reasoning. Thus, the need to meet this first condition highlights the fact that the CS approach only has meaning within the context of inference rules that relate such subtracted contents. A "logic" of subtracted contents is required, one that will capture the *a priori* relations between, e.g., the content "glass[-EU, -2] at location1[-EU, -MC]" entertained at time  $t_1$  and the content "glass[-RP, -2] at location2[-EU]" entertained at time  $t_2$ , where location<sub>1</sub> and location<sub>2</sub> are ego-centric specifications of places, such that they are co-referential, given the turning action performed between times  $t_1$  and  $t_2$ . The inference from the first to the second content, in as much as it is correct, will have to fall under some inference rule in this "logic" of subtracted contents.

The second condition puts further constraints on the CS method. For it seems possible that there may be any number of ways that a concept could fail to incur some particular conceptual commitment. For example, it seems that any number of contents meet the condition "just like the concept *glass*, but does not meet the Generality Constraint". So it seems that one's catalogue of conceptual commitments is going to have to be sophisticated indeed if one is to be able to specify a content uniquely.

But perhaps this just shows that the second condition is, strictly speaking, too strict. Of course, there is something to the idea that content specifications are useful only when there is some restriction on the contents that they specify. But this need not imply that specifications are of use in psychological explanation only when a unique content is specified. One might be able to specify only some restricted set of contents, those that share a particular set P of properties, as opposed to specifying a unique content. But if the explanation to be given need only appeal to the fact that the content possesses the properties in P, and if it can be made intelligible that the non-intentional characterisation of the system to be explained could instantiate some content with the properties in P, then perhaps no further individuation is required. In fact, anyone who thinks that many of our ascriptions of conceptual content are, strictly speaking, inaccurate will have to appeal to some consideration such as this in order to make sense of the fact that such ascriptions are as successful as they are.

With respect to the third condition, the CS method of specification will inherit the advantages of purely conceptual specifications: the ability to specify content in terms of its essential properties. If the first two conditions can be met or dispensed with, it seems that one can

<sup>10</sup>Thanks to David Charles for this suggestion.

only question the canonicity of CS specifications if one is willing to question the canonicity of linguistic use specifications as well.

Another advantage of this close relation to conceptual specification is the ability to unify the conceptual and non-conceptual aspects of content within the same formalism.

But there are several obstacles to the successful deployment of this method. One possible worry is that the commitments to be subtracted must be atomistic: it must be the case that if one subtracts a commitment, one is not logically forced to subtract out other commitments. Or at least if there are such holistic inter-relations, they should be explicitly captured in a syntax of some kind. For example, if it is impossible to fail to meet the Generality Constraint without also failing to meet Russell's Principle, then either these should be rejected as candidates for commitments to be subtracted, or one must rule out, formally, the possibility of C[-GC] and C[-RP] for all concepts C.

This worry seems unfounded, however. As long as the commitments P referred to in a specification are sufficient to meet the three conditions above, it doesn't seem necessary to refer to other commitments, even if they are holistically related to those in P. This view might have to be abandoned once one starts to develop a logic for subtracted contents, since one might want to guarantee, e.g., that distinct specifications imply distinct contents. But note that this is not guaranteed even for linguistic use.

But there are other worries. Perhaps there is no canonical, finitely-specifiable list of conceptual requirements, either in general, or for particular contents. Another possible difficulty is that the method might not be general enough; there might be non-conceptual contents that are not expressible as subtractions of conditions from conceptual ones. The problem is not that there might non-conceptual contents that are subtractions of concepts other than those which we, as human theorists, possess; the fact that we do not possess these conceptual contents is not in itself an argument against the idea that they could be specified as logical functions of the concepts which we do in fact possess. Rather, the worry is that there might be non-conceptual contents that are not subtractions of any conceptual contents, be they in our possession or not. Without an argument against such a possibility, it would be excessively teleological to assume that all non-conceptual contents must be able to be expressed as subtractions from the conceptual contents into which some of them develop.

Finally, there is a general problem for non-embedded means of content specification, including both linguistic use and conceptual subtraction: the externalism of content. A general externalist claim is that the intentional nature of the cognitive phenomena to be explained requires that the specifications of the contents involved must make reference to the environment of the subject. This is not only because intentional properties do not in general supervene on the states of the organism alone<sup>11</sup>, but also because intentional phenomena can only be specified, explained, and understood in terms of their directedness toward the external world and the potential to interact with it. For example, it seems very likely that a means of specification must include some way of representing the spatial environment of the subject if it is to be able to express and explain spatial NCCs and their inter-relations, as used in the construction of cognitive maps.

---

<sup>11</sup>Martin Davies (Davies, 1991) gives some compelling examples of non-conceptual, perceptual contents that do not supervene on the internal state of the organism experiencing that content; the important differences are in the way the organism is embedded in the environment, and thus the environment can be expected to play a major role, beyond the one of specifying truth-conditions, in the specification of non-conceptual contents.

One might question this conclusion by noting that conceptual content is intentional, yet we specify such contents via non-embodied means (linguistic use). One reason why we can get by with non-embodied specifications for conceptual content, but not for non-conceptual content, will be given in 4.2, below.

But there is also reason to believe that we *can't* in general get by with non-embodied specifications, even in the purely conceptual case. If externalist positions such as those expressed in (Burge, 1982) and (Putnam, 1975) are correct, then there is no way that a sentence of English language on Earth could specify even the conceptual contents entertained by our Twins on Twin-Earth. So, *a fortiori*, linguistic use could not specify Twin-Earth NCCs.<sup>12</sup> In order for the CS method to avoid this limitation, it must be the case that one can subtract from an Earthly conceptual content to yield a Twin-Earthly NCC. This seems possible only if the NCCs of Earth and Twin-Earth are the same, i.e., if externalist arguments apply only at the conceptual level. Yet there are those (e.g., Davies, 1991) who would maintain that even (some) non-conceptual contents are external. If so, we have yet another reason to reject non-embedded means of specification of NCCs. Perhaps, then, it is time we turned to embedded alternatives.

## 4.2 Embedded Indexicals

One such alternative means of specification is suggested by the discussion at the end of section 3. There, descriptive linguistic use was rejected as a means of NCC specification because it cannot accommodate contents that have constitutive connections to perception and action. And indexical linguistic use was rejected because it alone could not specify content, but rather it must be supplemented by an environment within which language can function.

But we actually do specify contents via indexical linguistic use; we might explain RC's behaviour by saying "He started cleaning up the mess because he realised *that he himself was the one making a mess*", which includes a content specification by means of indexical linguistic use (the reflexive "he himself"). So either we don't need to appeal to an environment when specifying contents via indexical linguistic use, or appeal to such an environment is possible, and effortless.

Well, "no" and a qualified "yes". No: we *do* need to appeal to the environment with such specifications. One has to know which subject is in question in order to fully grasp the significance of their first-person thoughts. This is clearer in the case of demonstratives: "He thought *<that> is a doorway*" will explain why a subject ran into a false stage door if one understands, *inter alia*, that the *<that>* refers to the false door. Just as attributions such as "The content of the agent's visual perception was *this*" are effective, if at all, only when the speaking theorist and hearing theorist share the same environment (a condition that cannot, in general, be expected to be fulfilled in the practising of cognitive science<sup>13</sup>), the specification of non-conceptual contents (indexically or otherwise) will have to recreate this context by invoking some detailed description of the agent's environment.

<sup>12</sup>Note that arguments merely to the effect that symbols must be grounded, that there must be *some* environment in order for an agent's states to have any content at all (Harnad, 1990), do not in themselves argue against non-embodied specifications. It is only when one claims that external conditions partially *individuate* a content that one can put forward this kind of externalist argument for embodied specifications of content. The argument that symbols must be grounded is compatible with an internalist individuation of contents, even though it demands that such contents can only exist in the context of an environment toward which they are directed.

<sup>13</sup>But see section 4.4.2.

But if so, doesn't this just show that we already take embodiment seriously, and make implicit appeals to a subject's environment, when we employ conventional indexical linguistic use specifications? Yes, but: the simplicity of the task of world-involving in such cases is a consequence of the systematicity of conceptual, linguistically specifiable contents involved. Such contents have conceptually elegant rules of world-involvement: e.g., "a use of the first-person mode of thought refers to the person who is using it". Once that conceptual simplicity is absent, as in the case of non-systematic NCCs, the rules for world-involvement become fragmented, non-systematic, and ad hoc, and thus demand more effort for their specification, as well as the specification of the environment in which they function.

Can the specification of such non-systematic indexicals proceed by means of linguistic use? Can there really be a term that has associated with it the world-involving function appropriate for an NCC? Technically speaking, I suppose so, if we can specify such contents at all. For once one had some theoretical grasp of the function in question, one could simply introduce a term that had that function as its indexical function. But the point is that one would have to have some way of theoretically grasping that function in the first place, since the function will not be one with which we are already familiar in our everyday use of language. Even if, strictly speaking, indexical linguistic use is possible, its possibility is contingent upon that for which I am arguing: an alternative to linguistic use specifications.

The challenge of NCC specification via indexical linguistic use will not primarily be a matter of choosing the right (non-systematic, non-linguistic) indexicals, but mainly a matter of specifying, in the appropriate ways, the subject, its context, and the relations between them. A large part of the work in developing a means of specification for NCCs will be formalising the practice of highlighting certain aspects of the subject/environment system so that a particular, non-conceptual way of representing that situation is indicated. And even the task of choosing the right indexicals will require some sophisticated way of relating the subject to its environment.

But embeddedness requires embodiment; because EI specifications must be embedded, they must take embodiment seriously. In order for one to be able to make reference to the relations between a subject and its environment, one must think of the subject as having a position in that environment. Also, in order to understand how the highlighted environmental factors play a role in fixing the content, one must have some understanding of (at least) the perceptual and motor capabilities of the system. It is for these reasons, and because one must specify the non-systematic indexical functions involved in grasping various NCCs, that EI specification must make reference to the underlying, non-intentional characterisation of a system.

### **4.3 Content Realisation**

The last two means of specification to be considered here, content realisation (CR) and ability instantiation (AbI), are both, unlike those before, non-conceptual specifications in that they do not express, but rather refer to, the content to be specified, and therefore employ concepts without requiring the organism to possess those concepts in order to entertain the content so specified. In the case of CR, this reference is achieved by mentioning a set of perceptual, computational, and/or robotic states and/or abilities that realise the possession of that content in a particular case or set of cases.

As mentioned before, specifications must indicate at least one content, at most one content, and must indicate contents canonically. These three conditions, involving as they do the

notion of indication, are primarily epistemic constraints. The non-conceptual nature of CR specifications, however, will give at least the first two of these conditions a metaphysical bite.

#### 4.3.1. At Least One Content: Realisation

By requiring that the referenced states indicate at least one content, the first condition entails that the states mentioned in a CR specification must realise a content; they must be sufficient for the possession of a content. One could imagine a weaker form of state- or ability-based specification, in which the states would not have to realise the content they specify, but would instead merely suggest to the theorist the content to be specified, with no accompanying metaphysical claim that those states specify the content because they realise it. But if the metaphysical relationship is abandoned, what relationship is to be put in its place? How is one to know if the states on offer will succeed in suggesting the content in mind? In the absence of answers to these questions, any alternative to the metaphysical approach is precisely the kind of specification that I am trying to avoid: one that succeeds, when and if it does, without appeal to any principle (or at least not any articulated principle). A scientific psychology requires more rigour than such a means could currently provide; it seems that such rigour could only be provided, if ever, by a means of specification informed by a theory of "suggestion" itself, i.e., a near-complete scientific psychology. Scientific psychology would have to be completed before it could begin.

Note that there is nothing in the CR approach that precludes an externalist individuation of content. It might be true that individualistic properties alone do not determine some or all contents (although not necessarily for the reasons given in, e.g., Putnam, 1975 and Burge, 1982), but this just means that the states used to specify such contents will themselves have to be externalistically individuated. This is not in itself a difficulty (*pace* Fodor, 1981), since there are several examples in cognitive science of such an embedded notion of state or ability.

It is important to the proper understanding of CR specification that one note that although sufficiency of the states for the specified content is required, necessity of the former for the latter is not. That is, the specification does not have to provide or even invite a reduction of the content it specifies. To specify a content by mentioning one realisation of that content is not to indicate the physical type that constitutes possession of that content in general. Indeed, CR specifications do not even require that a reduction of the content to a non-intentional vocabulary is possible (which is just as well, since there are good reasons to believe that such reductions are *not* possible). Conversely, the fact that there might be infinitely many other physical configurations, that fall under no nomologically-governable physical type, that are also realisations of the given content, counts not one whit against the ability of the particular realisation mentioned to pick out, clearly and distinctly, the content in question. Consider how one might indicate to someone a particular economic phenomenon by describing a particular manifestation of that phenomenon, in terms of a particular currency, set of countries, etc.

#### 4.3.2. At Most One Content: Holism

In order for the states referenced in a CR specification to specify a content, they must not only realise the content in question; the second condition mentioned above, together with a simplistic notion of realisation-based specification entails that the specifying abilities must realise *only* that content. This would appear to be at odds with the fact that content is

holistic: contents come in groups, so any abilities that are sufficient for one content are going to be sufficient for others as well. Thus, it seems that CR specifications might have difficulty respecting the second condition.

However, if one assumes some kind of structural modularity in the system being used to specify the contents<sup>14</sup>, then CR specifications may be able to respect the second condition, may be able to specify only one content, despite holism. Although it might be true that a system acquires a set of contents (call them A, B, C & D) as one, it might be that there are isolable components (p, q, r & s) of a state realising those contents such that each component is counterfactually related to only one of the set of contents. So having component p is not sufficient for possessing content A, since q, r, & s are required as well, but it is the p component part of the total state that is "responsible" for A. If p were to change to p', then A would no longer be present, yet if q, r, or s were to change to q', r', or s', A would still be present. This, then, would allow one to single out p from its supporting context of q, r, & s, in order to specify A uniquely.<sup>15</sup>

### 4.3.3. Canonicity

Even if CR specifications indicate one and only one content, it must be ensured that they do so canonically. That is, they must avoid, e.g., being like the linguistic use specification mentioned before: "the content toward which the subject took the belief attitude exactly 10.3 seconds ago". One might think that CR specifications cannot specify contents canonically, since canonical specifications must invoke the essential properties of the content, while CR specifications proceed by mentioning a particular realisation of that content, which might be thought to be only contingently related to the content.

But a requirement for such strict necessity seems to too stringent. Consider linguistic use specifications. Although there *might* be a necessary connection between a content and a *word* that expresses it, the relation between the content and the *sounds* or *marks* that instantiate the word that has that content is contingent. So if such marks are sufficient for canonical specification, then it seems possible that other entities contingently related to the content, such as one of its realisations, could also be sufficient for canonical specification.

To make good this analogy between marks and particular realising states, there needs to more to CR specifications than mere realising states, just as there is more to linguistic use specifications than mere marks or sounds. In the case of linguistic use, there is a practical capacity, on the part of the theorist, to relate these arbitrary marks to the contents they express. This practical capacity is part of being a member of a linguistic community, and is acquired through exposure to the norms that the community applies to the sounds and marks that the language comprises. So it would seem that CR specifications, if they are to be canonical, must rely on some practical capacity for relating particular realising states to their general forms, the forms which are essential to any state that realises the content in question.

<sup>14</sup>Note that one does not have to make this modularity assumption for the system being explained, only the system that is being used to specify the contents of the system being explained.

<sup>15</sup>This modularity requirement might seem to be at odds with what has been said elsewhere (Cussins, 1990; Chrisley 1993), that NCC's are perhaps best suited, not to a classical computational architecture, but rather to the distributed representations of PDP models. If that claim is correct, and if one takes the modularity required for CR specification to be at odds with PDP representations, then it would appear that CR specification cannot appeal to modularity to isolate unique NCC's in the holistic web. However, this appearance is misleading. Modularity of the sort necessary for CR specification is not incompatible with distributed representation, as complete, propositional NCC's might be localizable, even though their constituents are highly distributed over the locale of their comprising propositional content.



And this capacity might have to be acquired through familiarity and practical interaction with the system in question.<sup>16</sup> Note that one would not have to develop such a capacity for every system to be explained, but only for the system or systems that one wishes to use for the purposes of CR specification of content.

However, the more that one models CR specification on the case of linguistic use specification, the more one runs the risk of limiting CR to conceptual contents. It could very well be that the requirement of a public, practical capacity to understand others is what restricts linguistic use to objective contents. If so, one might worry that the requirement for a practical capacity to understand the canonical realising system in CR specifications might likewise limit such specifications to the conceptual case. Although this worry cannot be dispelled entirely here, it should be pointed out that there is a stronger link between a content and one of its realising states ("intrinsic intentionality") than the relation between a content and the arbitrary properties of one of the symbols that convention and practice have associated with that content ("derived intentionality"). In fact, once the general parameters of the specifying system have been determined, there might be a necessary relationship between a state, given that it is bounded by those parameters, and the content it realises. Therefore, canonical specification may be possible without relying on practical capacities that might restrict one to conceptual contents. But perhaps this is just optimism; at present the issue is unresolved.

#### **4.4 Ability Instantiation**

The limitations of both non-embodied (linguistic use, CS) means of specification, as well as of those that merely *mention* embodiment (EI, CR), might have a common cause. Specifically, it might be that the only way canonically to specify NCCs is via an explicit demonstration or actual instantiation of the idealised robotic and computational abilities involved in entertaining that content. Attempts to specify an NCC in a linguistic use manner fail to indicate (to any theorist seeking to understand the agent) the correct content and therefore leave certain connections to perception, action, and other contents inexplicable; perhaps merely *mentioning* the abilities must also fail, for similar reasons. Practical, canonical specification of an NCC might require the actual instantiated presence of an ability, rather than the conceptual idea of that ability. There are two possible ways that the abilities could be instantiated: external to the theorist, in some apparatus (external ability instantiation, or EAI); or within the theorist/environment system itself (self-instantiation, or SI).

##### **4.4.1. External Ability Instantiation**

There are practical reasons why actual instantiation of specifying abilities, as opposed to mere reference to them, might be required. The demands for embeddedness in content specification call for a means of specification that not only allows one to represent explicitly the spatio-temporal relations between the system being modelled and its environment; it must also itself be a concrete system that persists through time and possesses computational abilities. As mentioned before (in the discussion of the possibility of embedded indexical, EI, specifications), the context of the subject will have to be reconstructed if we are to be able to specify (at least some) NCCs; this should be achieved, at least in part, via a judicious choice of the syntactic properties of the specification formalism itself. It should have an

---

<sup>16</sup>See section 4.4.2 for further discussion of how this might be possible.

active, computational format, rather than the static format of axioms and theorems on a printed page. The complexity of a fine-grained static formalism with spatio-temporal 'parameters' (e.g. "Believes[robot3, time1, *that is an obstacle*, the chair at location (x,y,z)]) would be prohibitive; instead, a computer simulation (of the interactions of the various axioms and 'given' conditions of the theory) would make explanation and prediction more tractable than if one were to use non-instantiated, referential, 'manual' analysis. In order to understand which content is involved in a particular situation, a theorist could look at how an instantiated system (e.g., an android, or a simulation of one) responds to different counterfactual contingencies, could monitor the evolution of the current state forward or backward in time, etc. This would not be feasible with a non-instantiated means of specification.

But there is a more theoretical reason why an external, active instantiation of some abilities that realise a non-conceptual content might be necessary for the canonical specification of that content. EAI specification can be seen as a response to the worries, just expressed at the end of 4.3, concerning the canonicity of content realisation (CR) specifications. Perhaps one can specify contents, as CR aims to do, in terms of the states and abilities which realise them, but the pre-requisite practical ability, on the part of the theorist, to move from abilities to contents might require an active presence of those abilities, with which a theorist can interact, not just reference to them. The observable temporally-extended action of the computational formalism (and its interaction with its environment) might be the only way canonically to specify certain NCCs, given their resistance to specification by standard means. Some phenomena may only be explicable through the use of *models*; perhaps only via models in which, e.g., actual time and space are used to represent the temporal and spatial aspects of the modelled system, as opposed to formalisms that represent those aspects with something else: a written variable or spacing on a page. It seems likely that in order to be able to specify NCCs and their inter-relationships, one will have to choose representations for them in such a way that there is a non-arbitrary relationship between the syntactic properties of the representations and the contents to which they refer: the syntactic properties will assist directly in specifying the content.

This approach (and others presented here, inasmuch as they are concerned with the question "what is a *canonical* specification?") places an emphasis on the theorist's own embodiment, with the notion of a theorist's *psychology* that such embodiment implies. Canonical specification cannot proceed independently of the cognitive make-up and limitations of the theorist using that specification; rather, what counts as a sufficient specification or, more generally, explanation, will depend on the conditions under which the *theorist's* abilities to grasp contents may be exercised. One tentative proposal is that our psychologies as theorists are such that we will only have canonical NCC specification when we employ actual instantiations of that content.

#### 4.4.2. Self-instantiation

There are two different proposals to be considered in this section, although both are similar in several ways, including their speculative nature and science-fiction feel. With that by way of a disclaimer...

The first form of self-instantiation continues the realisation thread under consideration in sections 4.3 and 4.4.1. There a worry was expressed: that the practical capacity required to move from realisations of contents to the contents themselves will have to be similar to the practical capacity for language to such an extent that only linguistic, conceptual contents can

be so specified. Despite the observation that the relationship between realisations and contents is less arbitrary than that between sounds or marks and linguistic contents, the discussion in 4.4.1 suggests that the practical capacity may have to be an almost social kind of interaction with the specifying system if it is to provide canonical specifications of NCCs. This fuels the worry.

Perhaps the instantiation that provides canonical specification should be something more intimately known, thus avoiding the need for interactive, social capacities. Perhaps the instantiation should be the theorist itself.

The scenario I have in mind is not the use of some private "inner pointing", against which Wittgenstein railed. Rather, imagine a (possibly not-too-futuristic, given recent advances in imaging techniques) situation in which the theorist learns the relation between publicly observable states/abilities and contents for his or her own particular case. The theorist's non-intentional state at any given time will be directly observable, and the theorist will have a privileged (though not necessarily infallible) acquaintance with the corresponding intentional state. This combination of extro- and intro-spection may permit the development of a practical capacity for the theorist to mention his or her own physical states to specify contents that would otherwise be ineffable. If one further assumes that the realisations of theorists' contents do not differ dramatically from each other, there will be the possibility of theoretical, scientific communication concerning NCCs.

This means of specification is not, strictly speaking, one of state/ability *instantiation*, since the specifications themselves could very well be references to or mentions of the states of the theorist that realise the content in question. But since the development of this capacity for such states canonically to specify contents canonically requires a period in which the theorist actually instantiates (perhaps some "basis" subset of) the contents to be specified, it seems appropriate to mention this method under the "instantiation" rubric. Any of the instantiation methods could have an initial period of specification via instantiation, during which technical terms are introduced to refer to the contents so specified. But use of such terms for specification would still be a case of instantiation specification, since the norms of use of such terms is governed by their means of introduction.

However, the second form of self-instantiation specification is more directly a case of instantiation. It also has more of the feel of the "inner pointing" which Wittgenstein argued against, yet with a grounded twist that might allow it to avoid coming under the purview of his private language arguments. The idea here is to cut out the middle man, by altering the environment of the theorist (to which one wishes to communicate a content) such that the theorist actually takes an attitude toward that content.

This can be seen as playing the same role for EI specification, in terms of grappling with the constraints placed on canonicity by the nature of the theorist's cognitive abilities, that EAI specification and the first form of SI specification played for CR specification. For example, in a typical EI (embedded indexical) specification, one might say "the infant sees the wall like *<this>*", followed by a description of the infant's environment, the its position and orientation within that environment, and its sensori-motor abilities. The analogous move to that made before, then, is to claim that this referential approach is not sufficient for canonical specification, a more instantiated approach is required. The move would claim that any success for the EI method would be due to the theorist being able to *imagine* the situation from the infant's point of view. But our imaginations are notoriously limited; why not

actually have an externally-prompted experience with the same content as the one to be specified?

Clearly, it would be too awkward (at present) for one to manipulate a theorist's environment to the extent necessary for such specifications. But we cannot rule out the possibility that technology (e.g., virtual reality) could be of assistance here, if or when it is developed.

Nevertheless, there are obvious potential difficulties: could an adult theorist, no matter how his or her environment is manipulated, really see the world the way an infant does? Even if one believes, as is surely the case, that adults entertain a wide range of non-conceptual contents, is it plausible that they are the same contents that are entertained by an infant? A bat?

Rather, it seems that if canonical specification requires *that* close of a link between theorist and subject, then we are severely limited in our capacities to understand each other from a scientific viewpoint. I choose to interpret this as a strike against such a strong notion of the requirements for canonical specification, rather than against the prospects for a scientific psychology.

## **5. Embodiment and Computation**

No matter which (if any) of the types of alternative specification actually turn out to be successful, it seems clear that NCC specification requires appeal to the spatio-temporal relations between the system being modelled and its environment (embeddedness); for this (recall the end of 4.2) and other reasons, then, such specification also requires either reference to, or the instantiation of, (some of) a content-exercising system's intrinsic non-intentional properties (embodiment).

It is natural to look to computation and robotics to provide ways of characterising and thinking about the functionally relevant aspects of the system's embodiment and its environment. But there are two thoughts that might give one pause.

First, computational phenomena are themselves arguably intentional. Computational states are typically representational, they are about things, they carry their own form of (sub-personal) content. So one might wonder how computational notions could provide the characterisations of non-intentional states required for NCC specifications. For indeed the embodiment and environment of the system must be characterised as non-intentional (or at least non-contentful) if an infinite regress of content specifications is to be avoided. But computational analyses specialise at coming up with elucidating, un-interpreted (if not downright non-intentional) ways of characterising intentional systems. Of course, computational states *are* intentional, are about something; but viewing them as, say, Turing Machine quadruples is to highlight their merely causal properties, and to ignore their semantics. Perhaps the value of this kind of analysis has been over-emphasised, or misunderstood; I certainly don't think that a complete understanding of computation will be primarily formal and non-intentional. Nevertheless, such characterisations do have their place, and they might be ideal candidates for capturing the embeddedness and embodiment of systems for the purpose of content specification.

On the other hand, some might think of computation as a world-independent, abstract notion, not the kind of thing that could square well with the requirements of embodiment and embeddedness at all. All that can be said here is that there are reasons, discussed elsewhere,

for rejecting this disembodied, asemantical view of computation (see, e.g., Smith, 1991). In fact, one could put the force of the issue the other way: given that content specifications must be embodied and embedded, if cognitive science is going to understand representational content in terms of computation, we had better develop our computational notions accordingly, rejecting the formal for the embodied and embedded.

One way that non-conceptual content and computation relate, then, can be captured with the motto: "do not ask what your formalism can do for your robot; ask what your robot can do for your formalism". That is to say, it seems that in order to specify contents and their inter-relations, a means of content specification for an NCC-involving cognitive science will require concepts and insights from a theory of computation (especially robotic and perceptual computation). Further, such a formalism might require not only concepts and insights, but instances of computational phenomena.<sup>17</sup>

## **6. Conclusion**

The existence of non-conceptual content (NCC) places several demands on any cognitive science theory that wishes to address the full range of human cognitive behaviour. I have argued that the way to answer these demands is to take embodiment seriously, by establishing a close connection between NCC and computational/robotic abilities. I argued that we need an alternate means of content specification that can, unlike the standard method of linguistic use ("that" clauses), canonically specify NCCs. I suggested that a worked-out means of specifying computational and robotic abilities might go a long way to meeting these requirements, but have yet to produce a fully worked-out means of specifying NCCs. The demands that must be met before this can be done are considerable, but they should not discourage: an emphasis on NCC not only constrains, but also liberates, in that it allows psychologists to direct their energies toward explaining cognitive phenomena which have to be ignored from within a conceptualist approach, since the phenomena essentially involve contents which are non-conceptual: cognitive phylogeny, conceptual development, perception, learning, and action.

## **References**

- Allen, C. (1992). Mental content. *Brit. J. Phil. Sci.* **43**.
- Burge, T. (1982). Other bodies. In Woodfield, A. (Ed.) *Thought and object: Essays on intentionality* (pp 97-120). Oxford: Clarendon Press.
- Chrisley, R. (1993). Connectionism, cognitive maps & the development of objectivity. *Artificial intelligence review* **7**:329-354.
- Crane, T. (1992). The non-conceptual content of experience. In Crane, T. (ed.) *The contents of experience*. Cambridge: Cambridge University Press.
- Cussins, A. (1990). The connectionist construction of concepts. In M. Boden (Ed.), *The philosophy of artificial intelligence* (pp 368-440). Oxford: Oxford University Press.
- Davies, M. (1990). Thinking persons and cognitive science. *AI and society*.
- Davies, M. (1991). Externalism and perceptual content. *Proceedings of the Aristotelian society*.
- Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.

<sup>17</sup>Thus, my discussion here does not concentrate on computation via the claim that the mind is computation (although that equation might be a consequence of the concerns here); rather, the emphasis is that computation is a (and perhaps the only) formal means of specifying otherwise ineffable contents.

- Fodor, J. (1981). Methodological solipsism considered as a research strategy in cognitive science. In Haugeland, J. (ed.) *Mind design* (pp 307-338). Cambridge: MIT Press.
- Harnad, S. (1990). The symbol grounding problem. *Physica D* **42**:335-346.
- Haugeland J. (1991). Representational genera. In Ramsey, Rumelhart, & Stich (Eds.), *Philosophy and connectionist theory* (pp 61-90). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Peacocke, C. (1981). Demonstrative thought and psychological explanation. *Synthese* **49**:87-217.
- Peacocke, C. (1986). *Thoughts: An essay on content..* Oxford: Blackwell.
- Peacocke, C. (1989). *Transcendental arguments in the theory of content..* Oxford: Oxford University Press.
- Peacocke, C. (1990). Analogue content. In Perry, Almog, & Wettstein (Eds.), *Themes from Kaplan*. New York: Oxford University Press.
- Peacocke, C. (1992) Scenarios, contents & perception. In Crane, T. (ed.) *The contents of experience*. Cambridge: Cambridge University Press.
- Perry, J. (1979). The problem of the essential indexical. *Nous* **13**:3-21.
- Putnam, H. (1975). The meaning of "meaning". In Putnam, H., *Philosophical papers, Vol. II: Mind, language, and reality* (pp 215-271). Cambridge University Press.
- Smith, B. C. (1991). On the threshold of belief. In Kirsh, D. (Ed.) *Foundations of artificial intelligence*. Cambridge: MIT Press.
- Strawson, P. (1959). *Individuals*. London: Methuen.