

Ron Chrisley and Joel Parthemore

Synthetic Phenomenology

Exploiting Embodiment to Specify the Non-Conceptual Content of Visual Experience

Abstract: *Not all research in machine consciousness aims to instantiate phenomenal states in artefacts. For example, one can use artefacts that do not themselves have phenomenal states, merely to simulate or model organisms that do. Nevertheless, one might refer to all of these pursuits — instantiating, simulating or modelling phenomenal states in an artefact — as ‘synthetic phenomenality’.¹ But there is another way in which artificial agents (be they simulated or real) may play a crucial role in understanding or creating consciousness: ‘synthetic **phenomenology**’.² Explanations involving specific experiential events require a means of specifying the contents of experience; not all of them can be specified linguistically. One alternative, at least for the case of visual experience, is to use depictions that either evoke or refer to the content of the experience. Practical considerations concerning the generation and integration of such depictions argue in favour of a synthetic approach: the generation of depictions through the use of an embodied, perceiving and acting agent, either virtual or real. Synthetic phenomenology, then, is the attempt to use the states, interactions and capacities of an artificial agent for the purpose of specifying the contents of conscious experience. This paper takes the first steps toward seeing how one might use*

Correspondence: Ron Chrisley, COGS/Department of Informatics, University of Sussex, Brighton, BN1 9QH, UK. ronc@sussex.ac.uk

- [1] Thanks to Rob Clowes for suggesting the term ‘synthetic phenomenality’ to refer to this other work.
- [2] It was thought that the term ‘synthetic phenomenology’ had made its first appearance in a discussion between Axel Cleeremans, Antti Revonsuo, Tom Ziemke and the first author at a machine consciousness workshop in Birmingham in 2003, but it was later discovered that the term had been used in the title of a talk by Scott Jordan in 1998.

a robot to specify the non-conceptual content of the visual experience of an (hypothetical) organism that the robot models.

1. Specifying the Content of Experience

As with any science, a science of consciousness requires an ability to specify its *explananda* (facts, events etc. to be explained) and its *explanantia* (states, facts, events, properties, laws etc. that do the explaining). Conscious states (experiences) may be expected to play both of those roles. A science of consciousness, then, has a double need for a way to specify experiences. At least part of what is essential to most, if not all, experiences is their *content*. The content of an experience is the way the experience presents the world as being. How can we specify the content of particular experiences?

The standard way of specifying content is by use of ‘that’ clauses. For example, ‘Bob believes that the dog is running’ ascribes to Bob a belief, the content of which is the same as, and is therefore specified by, the phrase following the word ‘that’: i.e. ‘the dog is running’. We call this means of specifying content *linguistic expression* because the content is specified not by finding a piece of language that *refers* to the content in question (as does the specification ‘The content of the experience that Christina had 2.5 minutes ago’), but rather by finding words that *have* or *express* that very content.

There is reason to believe that within a restricted range of mental states, those with what we might call *conceptual* content, linguistic expression is expressively adequate. Given the close ties between language and concepts, reasoning, logical thought etc., it seems reasonable to suppose that for the mental states involved in such activity, there is a piece of language that expresses their content.

But even if all conceptual contents could be specified by linguistic expression, there are reasons for believing that conceptual content is inadequate to the task of characterizing all of the content of experience. That is to say, experience has not only conceptual but also non-conceptual content: content that does not meet the conceptual norms of rationality (McDowell, 1996), articulability (*ibid.*), recombability (Evans, 1982) and endogenous control (Prinz, 2002). There are some, however, who refuse to acknowledge a distinction between conceptual and non-conceptual content; some do so for substantive reasons, while some are engaged merely in a terminological dispute. Fortunately, settling such debates isn’t necessary for our purposes here; all parties can agree that there are contents of experience that, whether conceptual or not, cannot, whether in principle or merely in practice, be specified by linguistic

expression. This is enough to motivate the search for an alternative means of content specification. Those few who remain unconvinced can consider the following in an exploratory spirit; even if we don't, strictly speaking, *need* to find an alternative to linguistic expression, it will be illuminating to see if and how it can be done.

The need for an alternative has been recognized before, e.g. Chrisley (1995). Peacocke (1995) offers scenarios, ways of filling out the space around a subject, as a means of specifying a kind of non-conceptual content: scenario content. Bermudez (2003) carefully considers the problem of specifying what he calls 'non-linguistic' content, and offers some strategies for doing so. Other work, although perhaps not conceived of by its authors as potential solutions to this problem, can nevertheless be considered as such. In particular the work of Lehar (2003), which offers cartoons and sketches as a way of specifying striking aspects of human experience, is very sympathetic to the general approach we offer below.

2. Depictions: Do I have to Draw you a Picture?

One obvious alternative to using linguistic expression as a means of content specification is to use non-linguistic, non-symbolic specifications.

2.1. *Evocative depictions*

Instead of attempting to specify the content of a visual experience in words, one might simply draw a picture or take a photograph of the world from the subject's perspective, and use that depiction as the specification of the content (or at least the non-conceptual content) of the experience. These depictions aim to present the world to the recipient of the depiction in such a way that, as a result, the recipient has an experience with content of the same type as the one being specified. Call these *evocative* depictions of experiential content.

Although coarse features of the content of a visual experience can be conveyed by the simplest of sketches or casual snapshots, an attempt at communicating anything beyond that requires the application of significant artistry and skill. Merely photographing, from the subject's vantage point, the scene the subject is seeing would be inadequate. Such an approach assumes a naïve literalism that focuses solely on the world's contribution to an experience, at the expense of failing to recognize the ways in which the subject also contributes to that experience. Perhaps such an approach would succeed if experiencers (and theorists of experience) never changed, if they always brought the same interests, memories, alertness, concerns, knowledge etc. to each

experience; if psychology were constant not only diachronically, across time, but synchronically, across subjects. But it is hard to image such invariance in the face of a changing world as being a kind of psychology at all. Although the world plays a crucial role in determining the content of perceptual subjective experience (that is why it is *perceptual*), so also does the subject (that is why it is *subjective*, or even why it is *experience*). Consider the experience of a person who is red-green colour blind looking at a scene that contains some red or green. A normal colour photograph of the scene would give some idea of what the person's visual experience is, but it would mischaracterize the colour experience. In order to capture the subject's contributions to experiencing an objective scene (and perhaps to 'subtract out' inappropriate subjective contributions from the theorist receiving the specification) our literalistic specifications will have to be altered in some way. This can be done either by altering the depiction itself (e.g. changing all the red and green to grey), or by augmenting the depiction with indicative components (facts) or explicit imperative components (instructions) that, together with the evoked experience, give the recipient knowledge of the specified non-conceptual content (e.g. supplementing a normal colour depiction with the symbolic annotation 'Red and green in this photo should be interpreted as being the same shade').

It is important to keep in mind that evocative depictions are constructed by one theorist for another, as a tool, so that experiential contents may be referred to canonically. But a depiction of this sort is not a depiction of the subject's experience; it is a depiction of the world, the object of the subject's experience, done in a way that evokes the way the subject is experiencing that world. Further, it is not being proposed that such depictions are 'in the head' of a subject that is having an experience with the specified content, still less that the subject is aware of or perceiving such a depiction. Even though evocative depictions exploit the relation between the non-experiential (a picture) and the experiential (the experience it evokes), they, like linguistic expression, do not require one to possess a correct theory of such a relation. Evocative depictions can succeed independently of one's theory of the relation between non-experiential and experiential states, assuming one has such a theory at all.

2.2. *Referential depictions*

There is another way of using depictions to specify the content of experience that is *not* independent of one's understanding of the relation between the experiential and the non-experiential. *Referential* depictions of experiential content do not specify the content of an

experience by causing the recipient to have an experience with that content (although they may, at times, have such an effect). Rather, they aim to give the recipient discriminating knowledge of the content in question (Evans' 'knowledge which', distinct from 'knowledge that' and 'knowledge how' (Evans 1982)), knowledge that gives the recipient the ability to distinguish the specified experiential content from other experiential contents. (Symbolically augmented evocative depictions are referential in this sense, but they differ from the purely referential depictions being considered here in that they require the recipient to have an experience that is then referentially augmented to yield knowledge of the target content, whereas referential depictions do not.) Referential depictions assist the recipient in visualizing the structures that the associated theory says determine the specified experience. Doing so in the right way might be crucial for the success of the specification. For example, if one assumes, as we do in the case of the model described in Section 3, an expectation-based theory of perceptual experience (that the non-conceptual content of a visual experience is given by the set of expectations that one has relative to a set of relevant possible actions, such as eye movements), then one could attempt to specify a particular experiential content by compiling a list of the various actions a subject having that experience might perform, and the expected sensations that would result. In some sense, reference to the correct content would have been secured, but not in a way that is of use to the recipient. If instead one arranged designators of expected sensations spatially, where the location of a designator depended on the spatial properties of the potential action that generates the expectation (as we do below), then the recipient may be much more likely to know which set of expectations, and thus which content, is being specified. (Compare giving someone a list of 1s and 0s corresponding to the binary contents of a jpeg file, as opposed to giving them the picture that file encodes.)

2.3. *Enactive depictions*

We call the kind of specification used in the above example *enactive* depictions. Like referential depictions, enactive depictions are theory-mediated, and can thus be seen as a special case of that class. However, because they are associated with a particular kind of theory of experience — theories that take action-indexed expectations of sensory input to be constitutive of the content of experience — they have available to them a particular mode of conveying the abilities that manifest a content, and therefore the content itself. Since enactive depictions present these expectations in a spatially-indexed way, iso-

morphic to the spatial relations of their associated actions, any recipient of such a specification will themselves *come to have* a set of expectations that are isomorphic to the expectation set of a subject with the experiential content being specified. It follows that, according to the expectation-based theory of experience being assumed here, the recipient of such a specification will have an experience with a content that is structurally isomorphic to the content being specified. Knowledge of this fact, and acquaintance with their own experiences, allows the recipients to ‘enact’ the relevant content, and therefore to know which experience is being specified. In Section 4 it will be suggested that enactive depictions can be more interactive, generated on-the-fly in response to a recipient’s probing by use of an embodied, robotic system. This potentially permits specifications of content of substantially greater temporal and conceptual sophistication. But before such systems can be described, an example of the simpler case must be given (Section 3).

2.4. *Depictions: discussion*

The work of Igor Aleksander (e.g. Aleksander and Morton, this volume) is a landmark in the field of machine consciousness, and it, too, employs the notion of a depiction. In fact the notion of depiction employed in that work is very similar to the notion of an enactive depiction: sensations indexed by spatially-ordered actions. But the role that depictions play in Aleksander’s work is different from the role they play here. There, depictions are (one of the) mechanisms underlying conscious experience. In Aleksander’s work, no one, especially not a theorist, ever *looks at* a depiction.³ By contrast, depictions here are not mechanisms but communicative devices, used by theorists to specify experiences for each other. To be of use they must be *seen* by the theorists that use them. Despite this contrast, there is nothing incompatible about these two roles. In fact, an identifiably depictive mechanism underlying perception might facilitate a depictive specification of the experiences realized by that mechanism.

3. A Robot-Based Means of Specifying Experience

This section describes how a robotic system, SEER-3, is used to specify particular experiences. First, a simple discriminative theory of experience is assumed, without argument, for the purposes of illustration

[3] Sometimes Aleksander’s papers and lectures themselves depict a particular depiction a system is using; in such cases, he *is* engaging in synthetic phenomenology as understood in this paper. However, there is little or no discussion in his work of how this should or could be done; his contributions instead primarily concern synthetic phenomenality (cf the abstract of this paper). Thanks to Mike Beaton for helping to make this point clear.

only. Then the SEER-3 robot itself is described, showing how, in the light of the assumed discriminative theory of experience, it can function as a model of experience. With these pieces in place, it is shown how SEER-3 can be used to dynamically generate enactive depictions to specify to a theorist the experiential states being modelled.

3.1. *A discriminative theory of experience*

The approach to specifying experience we employ here relies on the use of a robotic system that models the experience of some hypothetical subject. It therefore requires a theory of consciousness that relates the states of the model to experiential states. This theory need not be a *constitutive* theory: it need not give necessary and sufficient conditions for a system to be an experiencer in general (that is, it need not solve the Hard Problem (Chalmers, 1995)). For our purposes, it can be assumed that the hypothetical subject is an experiencer, and that the robotic system models the aspects of that subject that are relevant to it being an experiencer (see Section 3.2). What is required is a *discriminative* theory of experience: a theory that determines, given the modelled facts on an occasion concerning a subject that is an experiencer, exactly which experience the subject has on that occasion.

These points can be made clear by considering the particular theory to be used in conjunction with SEER-3: the sensory-motor expectation theory of experience mentioned in Section 2.3. According to the simplest version of this theory (which is loosely inspired by the work of O'Regan and Noë (2001) but is not meant to be faithful to that work), (part of) the (non-conceptual) content of the visual experience of a subject at a time t is a spatially distributed conjunction of:

1. a subset of the sensory information being received at t ('foveal' sensations);
2. the foveal⁴ sensations the agent would expect to have were it to perform at t an action a , drawn from a set S of actions, spatially displaced from the location of the sensations in (1) in a way isomorphic to the spatial properties of the action a ;
3. the possible classification of some spatial regions as ones in which change has been recently detected (see Section 3.2).

The theory implies that the content of the visual experience of a humanoid subject would be sensations in a two-dimensional spatial

[4] Note that the foveal region in the model described in Section 3.2 is much larger relative to the visual field than in, e.g., the human case.

array, with the current foveal sensations at the centre; the visual sensations that the subject would expect to have, were it to look up and to the left, would be located up and to the left in the array; the visual sensations that the subject would expect to have, were it to look to the right, would be located to the right in the array; and so on. (Talk of a sensation S being spatially located at L is metaphorical shorthand; what is thereby referred to is a content C — an abstract object with no location — that presents the world as being S -like at L .) There are, of course, problems with this simple theory. For one thing, it does not allow enough of a gap between experience and the world; subjects can never be in error concerning the region of the world on which they are currently fixing their gaze. But this is of no consequence for present purposes; recall that we put this particular theory forward only because we require some theory or other in order to provide an example of (enactive) depictive specification of experience. Although there is no space to do so here, the theory can be modified easily to avoid this problem anyway.

Furthermore, the theory assumed here, as simple as it is, has some advantages. For one thing, it can be used to resolve so-called ‘Grand Illusion’ problems (Noë, 2001): the fact that, for most people at least, visual experience seems to be coloured to the periphery; the absence of blind spots in experience; the stability of experience despite constant saccading and head movement, etc. It also allows a relatively simple explanation of change blindness (Rensick *et al*, 2000), one that does not force us to abandon a representationalist approach to understanding experience (*contra* Noë, 2006). These claims cannot be defended at length here, but the examples of enactive depiction in 4.3 go some way toward giving them plausibility. This is another advantage to depiction: it completes the data-theory-model-data cycle, allowing us to evaluate a theory more directly by experiencing the implications of that theory. Making the consequences of the theory available to us in such an embodied, interactive and intuitive way might facilitate the kinds of conceptual change required for us to solve the Hard Problem of consciousness (cf. a similar discussion concerning intelligence in the final paragraph of Chrisley, 2003).

3.2. *The SEER-3 robotic model*

That our purpose here is demonstration of a technique, rather than arguing for a particular theory or model, allows us to invert the normal modelling relation. Rather than choosing a subject of experience and attempting to construct a robotic system that models it, we can start with a convenient robotic system and suppose that there were some

experiencing subject for which the robotic system is an adequate model. In the case of SEER-3, the robotic platform used is an off-the-shelf commercial zoomorphic robot (the Sony AIBO ERS-7), suitably programmed to implement a basic model of visual experience. The robot has a single, fixed-position video camera mounted at the tip of its nose and, in these demonstrations, only the head of the robot moves. Thus the head plays the role of a large, saccading eye in an otherwise stationary subject with monocular vision.

The hardware and software together comprise a model with the following components:

1. a visual processing component that transforms the raw camera signal by introducing a blind spot, reducing acuity outside the foveal area, etc.;
2. an attentional mechanism that re-directs the robot's gaze, either as a result of habituation or upon command from another component (e.g. the change detection mechanism, below);
3. an expectation-maintenance mechanism that uses currently received sensations (in the fovea only) to modify the robot's expectation of what sensations it would receive were it to make this or that head movement. A consequence of the foveal restriction is that unnoticed changes outside the foveal region (see below) will not result in changes of any expectations, leading to the possibility of change blindness (see Section 3.3);
4. a change detection mechanism that:
 - a. activates the attentional mechanism when local change in the current visual sensory field (foveal or non-foveal) is detected, and flags the fact; if change is detected everywhere in the field (global change), such as when the lights suddenly go off or come on (or, in our case, during a blink), then no re-direction or flagging occurs;
 - b. for the case of a currently foveated visual region that has recently been flagged as above (and only if it has been flagged), checks to see if current sensations violate expected sensations for that region; if so, the fact that such a change has been detected is flagged (in a way that is meant to model the third aspect of the content of visual experience mentioned in section 3.1).

Component 3, the expectation-maintenance mechanism, could be implemented in a simple recurrent neural network; work by the primary author on the CNM system (e.g. Chrisley, 1990) is an example of such an implementation. A possible advantage of such an approach

is that the automatic generalization features of neural networks would result in the extrapolation and interpolation of expectations to actions never before performed in the current context, resulting in an experienced visual field that spans the entire action space from the outset. The SEER-3 demonstrations reported here do not employ such an implementation for the expectation-maintenance mechanism; rather, a kind of look-up array is used. The array, corresponding in extent to possible 'eye' coordinates, is initially undefined, signifying the absence of any expectations. After performing a given action, such as the fixation of gaze x degrees to the right and y degrees up, the robot will modify its expectations for any action that would result in the robot receiving sensations from any point within the foveal radius r of (x, y) . The expectations for any location within that circle will be changed to be whatever sensations it is now receiving at that location. This will in general alter at least some of the expectations for all changes of gaze fixation to any point less than $2r$ from the current point of fixation.

This approach, in conjunction with the expectation-based discriminative theory assumed in Section 3.1, results in the field of visual experience starting from an initial, foveated region and expanding as more of the visual environment is explored. Although many would think this to be an unlikely feature of the visual experience of any actual organism, it is, by definition, a feature of the experience of the hypothetical subject the SEER-3 system models. Despite the developmental differences between the neural network and look-up table implementations, the steady-state extent of the modelled experiential visual field will be the same. Specifically, the modelled visual field will be a superset of the field of current visual sensation, delimited by where the 'eye' can saccade to.

The model being used is very restricted:

- as mentioned before, it only aims to model some of the non-conceptual aspects of visual experience;
- the organism it models, unlike real organisms, is incapable of anticipating change;
- the modelled organism is incapable of any movement other than the single 'eye'; or rather, if it were to make any such movement, its expectations (and thus its purported experience) would have little relationship to the sensations it would actually receive;
- in the model, there is only one movement that will result in a change of gaze fixation to a given location; contrast this with the indefinite number of ways you can fixate on a point by

moving your eyes and simultaneously moving your head in a compensating direction.

Despite these limitations, the model is adequate for the purposes of this paper, since it is rich enough to allow concrete applications of the discriminative theory.

3.3. Using SEER-3 to specify experiences

According to the discriminative theory presented in 3.1, the set of visual expectations an agent has plus its current visual input determines the content of some subset of its visual experience. It follows that a robot can model the having of a visual experience by having analogous expectations and input. Independently of whether the theory or model is correct, we can conclude, as per Section 2.2, that conveying the robot's expectations and its current visual input in the right way (although there may be more than one right way, even for a single recipient) can serve as a specification of the content of the modelled experience. The hypothesis we are investigating is that one right way of doing this is via enactive depictions (cf. Section 2.3). In the case of SEER-3, such depictions are constructed as follows. Suppose the expectation to be depicted is that the robot would receive, at the point of fixation, a sensation of hue h were it to fix its gaze y degrees up from and x degrees to the right of the origin of the axes of head movement. A depiction of this expectation would consist in placing a mark of hue h at the location y degrees up from and x degrees to the right of the centre of the depictive frame.



Figure 1 (see colour plate at back). An enactive depiction of the content of a visual experience modelled by the SEER-3 robot. A colour at a given location indicates SEER-3's expectation to receive that colour as input were it to look in that direction. The absence of any expectations of colour for a location is indicated by the colour grey. Regions for which change has been detected are indicated in bright red.

An example of a SEER-3 depiction produced by this process appears in Figure 1. (This is a grey-scale rendering; a colour version of this depiction appears as the second colour plate on page 181 at end of this issue.) The depiction is of the model's set of expectations after 191 saccades. As per the model described in Section 3.2, the pixel colour value at any point indicates a set of expectations based on recent previous experiences. For example, a blue pixel value at a location (e.g. the back of the chair) indicates not only an expectation to see blue at the point of fixation if the robot were to fixate on the location corresponding to the back of the chair, but also an expectation to see blue down and to the left of the point of fixation if the robot were to move its gaze to the location corresponding to the window.

There are two exceptions to the interpretation of colour just mentioned. First, the absence of any expectations for a location is indicated in grey. A grey pixel is therefore ambiguous: it could mean either an expectation to see grey, or the lack of any expectation at all for that location. Context should enable the theorist to resolve this ambiguity. Second, regions for which change has been detected (cf. the change detection mechanism in Section 3.2) are indicated in bright red; thus an ambiguity also exists for bright red. This latter ambiguity only results from displaying all kinds of content in one depictive frame; instead, one could have distinct layers of depiction for each kind of content.

A striking feature of this depiction is its fragmented nature: contours and object boundaries are not respected. This reveals the modelled experience to be non-conceptual, in that the relevant expectations are created and maintained in a way that is not guided by any concepts of, e.g., **straight-line**, **chair**, **hand**, etc.

As mentioned at the end of Section 3.1, we can use this depiction to evaluate some features of the theory and model employed in its generation. Specifically, we can see that an organism modelled by SEER-3 need not suffer from a 'Grand Illusion'. For example, the depiction makes clear how, on this model, visual experience can be coloured to the periphery of one's current visual input (and, in the case of the hypothetical, modelled subject, beyond!), despite the subject not currently receiving colour information from outside the fovea. Similar points can be made concerning stability and blind spots.

SEER-3 can also model a form of change blindness.

- In normal situations, local changes within the current visual field (foveal or non-foveal) will re-direct attention and cause a comparison of expectations with current input; if there is a

mismatch for a region, ‘detected change’ will become part of the content of visual experience.

- When a local change occurs at the same time as a global change (such as a flash of light), no change will be detected and no redirection of attention will occur; even if the gaze later moves to the region of change, and expectations are thus updated to reflect the change, the previous expectations and current input will not be compared, and thus no ‘detected change’ component will be present in experience. For SEER-3, a change in experience need not be an experience of change.

Of course, there are many aspects of the model and theory that do not fit with real, mammalian vision. So the foregoing are not offered as explanations of, e.g., change blindness in humans, but rather illustrations of how depictions can facilitate the evaluation of model and theory. However, insofar as SEER-3 is a representationalist approach, it does suggest that change blindness in itself does not require a non-representationalist explanation.

4. Future Work

There are many ways in which the work presented here could be developed.

SEER-3 involves a particular theory and model of experience. It would be helpful to determine for which other theories and models the means of specification offered here can be adapted, and for which a distinct means of specification must be found.

Implicit in the above is a notion of evaluating different means of content specification. Developing criteria of success (e.g. the extent to which they allow one to predict the behaviour of the experiencer) and ways of measuring this (e.g. can recipients of SEER-3 depictions perform better at predicting the behaviour of a model than recipients of standard verbal specifications of content?) would also be an important step.

Other kinds of experience and experiential content should be addressed. This includes not just an extension to non-visual, or even non-sensory experiential content, but also more sophisticated aspects of experience already possible within the visual modality, such as conceptual, affective and temporal content.

- **Conceptual content:** Much will hinge here on one’s theory of concepts. If to have an experience involving application of the concept **chair** just is to have a certain set of low-level expectations (e.g. ones that, unlike those depicted in Figure 1, respect

the boundaries of chairs), then specifications of conceptual content might be continuous with the specifications provided here. But if conceptual content involves a different kind of expectation altogether (e.g. the expectation that one will see a *chair*, rather than the expectation that one will receive this or that low-level sensory input), then a distinct approach might be needed. In such a case, the use of distinct depictive layers, touched on briefly in Section 3.3, might be employed.

- **Affective content:** Use of layers might be of use here as well. If each action has associated with it an expected desirability of outcome (good/bad, painful/pleasurable, etc), then these can be treated as low-level ‘inputs’ resulting from the action, and thus mapped and depicted in a manner analogous to what has already been done for visual input.
- **Temporal content:** One can distinguish the temporal aspects of the depiction from the temporal aspects of the content. On an expectation-based view of content, there are two temporal dimensions to content, corresponding to the two temporal dimensions of expectations: the time of the having of the expectation and the time that the expectation is about (Chrisley, 2002). Both of these might be accommodated by using ‘movies’ rather than static depictions such as the one in Figure 1, thus aligning the temporality of the depiction and the content depicted. But the temporality of the depiction can be used to capture non-temporal dimensions of content as well. In a way, this is already the case with the depiction in Figure 1; in order to know which content is being specified, the recipient must spend time scanning the image, saccading here and there over it. In general, the use of immersive virtual reality techniques might be an effective way to exploit the temporality of the recipient’s experience to handle the multiple dimensions of experiential content (Chrisley, 1995).

We close with a moral of this research for those working in machine consciousness. Showing others the experience one’s robot (simulated or real) is meant to model at any given time may well have to involve much more than merely displaying the sensory input of the robot at that time. Understanding what else is involved might require one to be more explicit about the discriminative theory of experience one’s model presupposes.

Acknowledgements

This paper has benefited from the comments of many people at a number of talks over the past three years or so; the authors would like to thank everyone who has offered suggestions, including the people at the following conferences and seminars: Toward a Science of Consciousness, Tucson, 2004 and 2006; Workshop on Machine Consciousness, Antwerp, 2004; Artificial Intelligence and the Simulation of Behaviour, Hertfordshire, 2005; Brain Inspired Cognitive Systems, Lesvos, 2006; Embodied and Situated Cognition, Torun, 2006; E-Intentionality Seminars, Sussex, 2006 and 2007; Institute of Neuroinformatics, University/ETH Zürich, 2007.

References

- Aleksander, I. and Morton, H. (2007), *Journal of Consciousness Studies* (this issue).
- Bermudez, J. (2003), *Thinking Without Words* (New York: Oxford University Press).
- Chalmers, D.J. (1995), 'Facing up to the problem of consciousness', *Journal of Consciousness Studies*, **2** (3), pp. 200–19.
- Chrisley, R. (1990), 'Cognitive map construction and use: A parallel distributed processing approach', in *The Proceedings of the 1990 Connectionist Models Summer School*, ed. D.S. Touretzky, G. Hinton and T. Sejnowski (Los Angeles: Morgan Kaufmann).
- Chrisley, R. (1995), 'Taking embodiment seriously: Non-conceptual content and robotics', in *Android Epistemology*, ed. K. Ford, C. Glymour and P.J. Hayes (Cambridge, MA: MIT Press).
- Chrisley, R. (2002), 'Some foundational issues concerning anticipatory systems', *International Journal of Computing Anticipatory Systems*, **11**, pp. 3–18.
- Chrisley, R. (2003), 'Embodied artificial intelligence', *Artificial Intelligence*, **149** (1), pp. 131–50.
- Evans, G. (1982), *The Varieties of Reference*, ed. John McDowell (New York: Oxford University Press).
- Lehar, S. (2003), 'A Cartoon Epistemology', available online at <http://cns-alumni.bu.edu/~slehar/cartoonepist/cartoonepist.html>
- McDowell, J. (1996), *Mind and World* (Cambridge, MA: Harvard University Press).
- Noë, A. (2001), 'Is the visual world a grand illusion?', *Journal of Consciousness Studies*, **9** (5–6), pp. 1–12.
- Noë, A. (2006), *Action in Perception* (Cambridge, MA: MIT Press).
- O'Regan, K. and Noë, A. (2001), 'A sensorimotor account of visual perception and consciousness', *Behavioral and Brain Sciences*, **24** (5), pp. 939–73.
- Peacocke, C. (1995), *A Study of Concepts* (Cambridge, MA: MIT Press).
- Prinz, J. (2002), *Furnishing the Mind: Concepts and Their Perceptual Basis* (Cambridge, MA: MIT Press).
- Rensick, R., O'Regan, K. and Clark, J. (2000), 'On the failure to detect changes in scenes across brief interruptions', *Visual Cognition*, **7** (1–3), pp. 127–45.