The XML Approach to Synthetic Phenomenology

David Gamez[*]
[*]Department of Computer Science
University of Essex
Colchester
C04 3SQ
UK
`daogam@essex.ac.uk`

Abstract

One of the major challenges in synthetic phenomenology is to find a way of systematically describing artificial non-conceptual phenomenal states. This paper puts forward a solution to this problem that uses three different XML files to describe a machine's structure, internal states and phenomenology. The advantages of XML are that it can be read by both machines and humans, it is good at capturing hierarchical relationships between data and it can be automatically generated, analysed and archived. XML could also be a useful tool for other methods of representing non-conceptual mental content, such as content realization and ability instantiation. Furthermore, as scanning technologies develop, the XML approach could be applied to the neurophenomenology of humans, which would serve as a foundation for a more scientific psychology of both humans and machines and facilitate precise comparisons between the two. The XML approach outlined in this paper will be used to describe the synthetic phenomenology of Holland's and Troscianko's CRONOS robot that is currently under development at the University of Essex and the University of Bristol.

1 Introduction

Synthetic phenomenology is a recently emerging discipline that aims to describe the phenomenal states of artificial systems. This is essential for the monitoring and debugging of machine consciousness and it could also address concerns about the possibility of suffering in machines. This paper puts forward an approach to synthetic phenomenology that uses three different XML files to describe a machine's structure, internal states and phenomenology. The advantages of XML are that it can be read by both machines and humans, it is good at capturing hierarchical relationships between data, it can be automatically generated, analysed and archived, and it avoids many of the pitfalls and presuppositions of natural language. As scanning technologies develop it may also be possible to use XML in neurophenomenology, which would allow detailed comparisons between human and artificial systems.

The first part of this paper covers some of the limitations of natural language descriptions of the phenomenology of non-human systems. After setting out the advantages of an XML approach, the central section outlines one way in which XML representations could be used in synthetic phenomenology. This is not intended to be a final and definitive methodology, since there are no doubt better ways of applying XML to this area. However, by presenting one way in which it could be done I hope to make the case that XML could be a very useful tool for the phenomenology of artificial systems.

2. Problems with Describing the Phenomenology of Non-Human Systems

Phenomenology, especially in the work of Husserl and Heidegger, derives its significance from the claim that the phenomena we experience are as important and substantial as the physical world described by science, which is often portrayed as a secondary interpretation of the phenomena. In this way phenomenology sets itself up with an 'objective' field of phenomena that are assumed to be the same for everyone and can be unproblematically described in natural human language The problem with this approach is that these assumptions about common experience start to break down once phenomenology is applied to the experiences of infants, animals and robots. To illustrate this problem, I will consider a short extract from Wordsworth (2004), which contains a fairly straightforward description of daffodils in natural human language:

> When all at once I saw a crowd,
> A host, of golden daffodils,
> Beside the lake, beneath the trees,
> Fluttering and dancing in the breeze.

Most people have had the experience of daffodils fluttering and dancing in the breeze and when Wordsworth's description is read by humans, they can readily imagine a similar past experience and understand his words well enough. However, even this straightforward description presents problems since it is extremely vague and imprecise and each reader will imagine the daffodils differently. More serious problems start to arise when we try to use ordinary language to describe the experiences of an infant placed in front of a field of daffodils. As Chrisley (1995) points out, we cannot simply say that the infant sees a host of golden daffodils because the infant has a preobjective mode of thought, which is unable to locate the daffodils within a single unified framework. Adults understand daffodils as something objectively located in three dimensional space, whereas infants do not necessarily continue to believe in the existence of the daffodils when they are occluded. In the adult and infant the word "daffodils" refers to two different concepts and experiences. As Chrisley puts it: "The infant's concepts are not fully objective and are therefore non-conceptual. To ascribe conceptual content to the infant in this case would mischaracterize its cognitive life and would not allow prediction or explanation of the infant's behavior." (Chrisley, 1995: 145).

These problems become even more difficult when the attempt is made to describe the phenomenology of a non-human animal, such as Nagel's famous bat (Nagel, 1974). When a bat flies over a field of daffodils it receives a complex pattern of returning ultrasound pulses, which are processed into phenomenal experiences that are likely to be very different from our own. Sentences like "the bat is experiencing a host of golden daffodils" are at best an extremely misleading description of the bat's phenomenology.

The same problems are encountered when attempting to describe the phenomenal experiences of artificial systems. Whilst we may have grounds for attributing phenomenal consciousness to some robots, we have almost no basis for believing that they will have the *human* phenomenal experience of yellow when daffodils are placed in front of them, or even that they will have the same experience of yellow as each other. Robots may also be built that have unconscious daffodil recognizers, so that they are only conscious of the abstract presence or absence of daffodils. Other robots might only be capable of processing stationary daffodils, leading to highly divergent phenomenal experiences that cannot be captured in ordinary language.

Natural language evolved to describe human experiences and so it is not surprising that it is very bad at describing the phenomenology of bats and robots. Synthetic phenomenology needs a better and more systematic way of describing the phenomenal states of artificial systems and the central claim of this paper is that XML representations are more appropriate for this task. After setting out the advantages of an XML approach, section 4 will demonstrate how it can be used to describe synthetic phenomenal experiences in a systematic manner.

3. Advantages of XML for Synthetic Phenomenology

The eXtensible Markup Language (XML) is a platform-independent way of structuring and organising data so that it can be easily shared between systems. XML is stored as plain text and it has a tightly structured format that enables the relationships between data items to be easily expressed. It is also possible to validate the structure of an XML file without any prior knowledge of its form. XML is starting to be used widely and there are a number of reasons why it would be suitable for synthetic phenomenology:[1]

1. XML is much more precise and highly structured than natural language, which allows it to describe complex nested hierarchies and represent the relationships between different pieces of information. This also enables easy cross referencing between different files.
2. XML can describe low level details of the system's hardware, but it can also abstract from them so that high level comparisons can be made between machines with different architectures and between humans and machines. Whilst two systems' lower levels might be different – perhaps using neurons or silicon - the higher levels are likely to be more similar, allowing direct comparisons between different systems once everything is encoded in XML.
3. XML can be written and read by both machines and humans. When doing simple small scale analyses it is useful to be able to manually read and edit an XML description of a machine's inner state. However, it is also very easy to automatically generate and analyse the state of a machine using XML, for example by writing programs

---

[1] A good XML tutorial can be found at: http://www.w3schools.com/ xml/default.asp

that look for phenomenal mental content using different theories of consciousness.

4. Its human and machine readability also make XML good for debugging consciousness. Once you have a highly structured representation of a machine's inner state and a methodology for analysing this for phenomenal consciousness, you can see how the machine's phenomenal states can be improved or increased.
5. XML is easy to archive, either by converting the XML files into a database format or by storing them directly. Sequences of mental content that are stored in this way can be examined later offline.
6. XML is a good foundation for the other techniques for representing non-conceptual mental content, such as those suggested by Chrisley (1995).
7. XML is very flexible. In addition to tags and data, XML can contain references to external files, pieces of code and equations. This enables it to include features that cannot be precisely described in human language.

Although these advantages also apply to some of the alternatives to XML, such as JSON, YAML and OGDL, the popularity of XML and the availability of good parsers in most programming languages make it the best choice for the approach to synthetic phenomenology that I am setting out in this paper.

4. The XML Approach

This section outlines one way in which XML could provide a systematic framework for describing the phenomenology of artificial systems. This approach works using three separate but interlinked XML representations:

1) *System.* A systematic description of the system and its sensors.
2) *Test Suite*. Identifies active elements within the system that are systematically correlated with outside events impinging on the sensors. This treats the machine as a complete unknown that is systematically probed by exposing it to stimuli and measuring changes in its internal state. During the generation of the test suite no attempt is made to say what the stimuli might be like for the machine, although human descriptions are included to help with later analysis.
3) *Mental Content.* If the test suite is constructed in enough detail, a good idea should be gained about the range of correlations between internal states of the machine and activation of the machine’s sensors by the outside world. However, at any point in time only a small proportion of the potentially active elements will be active and this set of currently active elements are recorded in a third XML representation of the machine’s mental content. This includes tags to indicate whether it is phenomenal mental content, which are filled in at a later stage by programs designed to analyse the system, test suite and mental content XML for signs of consciousness.

The XML structures that could be used to contain the data for each of these stages will now be covered in more detail.

4.1 System

The system XML file describes the structure of the system, including sensors, actuators and internal components. This is needed to clarify the range of tests that could be applied to the system and to help with the identification of potential phenomenal states. Some extracts from an XML file describing a typical system are given below:

```
<system>
  <description>Robot</description>
  <sensor id="1">
    <type>light</type>
    <shape>rectangle</shape>
    <width>400</width>
    <length>300</length>
    <coordinate_system>Cartesian
               </coordinate_system>
    <wavelength_range>0.7–0.4
               </wavelength_range>
  </sensor>
  <!-- Add more sensors here -->

  <actuator id="1">
    <type>motor</type>
    <location>wheels</location>
  </actuator>
  <!-- Add more actuators here -->

  <neuron id="1">
    <position>2,3,3</position>
    <type>pyramidal</type>
    <algorithm>Leaky integrate and
                    fire</algorithm>
  </neuron>
  <!-- Add more neurons here -->

  <connection id="1">
    <presynaptic_neuron>1
             </presynaptic_neuron>
    <postsynpatic_neuron>3
             </postsynpatic_neuron>
    <synapse_type>excitatory
                    </synapse_type>
    <weight>0.9</weight>
    <delay>22</delay>
  </connection>
  <!-- Add more connections here -->

</system>
```

Brief explanations of some of the more important tags are as follows:

<sensor> A sensor sensitive to light, touch or sound, for example.
<**actuator**> An actuator, such as a motor or hydraulic piston.
<neuron>, <connection> In this system the internal states are held in neurons, whose parameters are specified here along with the connections between them. Other systems might use Bayesian networks or first order logic to hold their internal states.

4.2 Test Suite

A test suite is a systematic way of linking the presence of events and objects in the environment to changes in the machine's inner state. To generate a test suite the system is probed using a number of different tests and correlations between the stimulus and the machine’s state are recorded as a list of active elements. The behaviour of the machine is also treated as data that is correlated with its internal states. To avoid presuppositions about three dimensional space, the input to the machine is specified in terms of changes in the machine’s sensors and not as the presentation of three dimensional objects. With systems based on real or simulated neurons the test suite could be created by following the traditional approach of recording from neurons or groups of neurons. Systems along the lines of Franklin’s IDA (Franklin, 1998) could be tested by using a debugger to monitor which variables or memory locations change in response to environmental stimulation. This avoids problems raised by Searle (1980) about the difference between manipulating a symbol and understanding a symbol since no assumptions are made about the meaning of any of the system’s internal states.

A comprehensive test suite needs to be designed with care so that it can probe all possible sensitivities of the machine and specify them as precisely as possible. This could start with simple low level features, such as points, lines, and edges and work its way up to more abstract stimuli, such as faces and houses. All of these single modality tests would have to be combined with input from other modalities, such as audition, proprioception and sensation. They would also have to be carried out whilst the machine is engaged in different activities, such as looking to the left, moving forward, and so on, to take account of sensorimotor contingencies. Whilst this sounds like an enormous quantity of work, initial tests of this type are likely to be carried out on very simple machines and as the methodology develops it will be possible to automate the creation of the test suite by writing programs that examine the system XML file and generate a comprehensive series of tests. The tests could also be automated in many cases by simulating the input to the sensors. Some sample extracts from a test suite XML file are given below:

```
<test_suite>
  <test id="1">
    <human_description>Moving
          forward towards point of
          light</human_description>
    <sensor_input>
      <sensor>1</sensor>
      <type>light</type>
      <size>5,5</size>
      <location>55,44</location>
      <wavelength>0.55</wavelength>
      <file>Test1.dat</file>
    </sensor_input>
    <!-- Add more sensor inputs -->

    <actuator_output>
      <actuator>1</actuator>
      <type>motor</type>
      <direction>clockwise
                       </direction>
      <speed>5</speed>
    </actuator_output>
    <!-- Add more actuator outputs -->

    <active_element>
      <type>neuron population</type>
      <neuron id="27">
        <firing_rate>0.88
                     </firing_rate>
      </neuron>
      <!-- Add more neurons -->

    </active_element>
    <!-- Add more active elements -->

  </test>
  <!-- Add more tests -->

</test_suite>
```

Some of the more important XML tags are as follows:

<**test**> A test that is applied to the machine to probe its responses to a particular stimuli. Tests that do not activate any elements do not need to be included.
<**human_description**> Description of the stimulus by humans, which may be useful as part of the process of describing the phenomenology of the machine.
<**sensor_input**> Input is defined in sensory rather than world coordinates. This is to avoid the presupposition of three dimensional space that might be made if we talked about presenting a round object at a distance of three metres, for example.
<**actuator_output**> Any actions carried out by the machine whilst the stimulus is being presented.
<**active_element**> The part of the machine's inner state that is activated by the test. In a neural system

this could be a single neuron or a population of neurons with a particular distribution of firing rates. In a more traditional computer system this could be a list of memory locations that are altered by the stimulus. Active elements are defined in relation to the test stimuli that activated them and have no meaning outside of this context.

4.3 Mental Content

Only a small proportion of the elements inside the machine that respond to stimuli are likely to be active at any point in time. The currently active elements are stored in the mental content XML file, along with the active connections between them. This mental content is capable of influencing actions and could be involved in planning. For example, if a machine has a group of simulated neurons that selectively respond to images of houses, then these neurons could initiate motor patterns that cause the sound "house" to be emitted. The house-sensitive neurons could also become activated when the machine was offline, leading to an experience analogous to imagining or dreaming about a house. Some of this mental content may be conscious and a tag has been included to record whether this is the case. The contents of this tag are filled in at a later point when the system, test suite and mental content XML files are examined according to a particular theory of consciousness (see next section). Sample extracts from a mental content XML file are given below:

```
<mental_content id="66">
  <time>4010551056</time>
  <active_element>
    <id>2</id>
    <intensity>0.7</intensity>
    <phenomenal>yes</phenomenal>
  </active_element>
   <!-- Add more active elements -->

  <active_connection id="3">
    <type>synchronisation</type>
    <from>1</from>
    <to>2</to>
  </active_connection>
  <!-- Add more active connections -->

</mental_content>
```

Some of the more important tags are as follows:

<active_element> Reference to one of the active elements defined in the test suite along with some of its current properties.
<active_connection> An active connection could be synchronisation between firing neurons, active processing by the CPU or simultaneous broadcast along a radio link. Since active connections are not necessarily topologically bound they are defined separately from the static connections in the system file.
<phenomenal> Records whether this active element is phenomenal mental content. The contents of this tag are filled in by examining the system, test suite and mental content XML files for signs of phenomenal consciousness.

4.4 Phenomenal Mental Content

The final stage in the description of the phenomenology of the machine is the identification of the parts of the mental content that are likely to be phenomenally conscious. This is done by analysing the system, test suite and mental content XML files using a theory of consciousness. It is highly likely that different theories of consciousness will make different predictions about the phenomenal mental content of the machine, which provides a good way of discriminating between them by comparing their different predictions with first person reports about phenomenal states.[2] This process of identifying the phenomenal mental content will now be illustrated using Tononi’s φ, Aleksander’s axioms and Metzinger’s constraints.

4.4.1 Tononi's φ

According to Tononi (2004) consciousness is linked to a system's capacity to integrate information. This is precisely quantified by Tononi as the number φ, which is the amount of effective information that can be exchanged across the minimum bipartition of a complex, where a complex is the subset of elements with $\varphi > 0$ and no inclusive subset of higher φ. Whilst there is not space to go into the details here, the system, test suite and mental content XML representations outlined in this paper would make it easy to calculate the amount of φ and pinpoint the active elements with high φ that are likely to be phenomenally conscious. It would even be possible to add a φ tag to the active elements within the mental content XML file.

4.4.2 Aleksander's Axioms

Aleksander (2003) put forward five axioms as a set of mechanisms that are thought minimally necessary to underpin consciousness. These are depiction, imagination, attention, planning and emotion. Although these axioms are not necessarily sufficient for consciousness, they are a good starting point for deciding whether a machine might be capable of conscious states and the XML approach offers a good way of analysing a system for their presence. For example, the test suite XML of an agent that

---

[2] There may also be ways of indirectly testing the predictions made by different theories of consciousness.

was capable of depiction would contain active elements linked to external stimuli, and an agent would be experiencing imagination when its mental content XML contained active elements that were linked in the test suite to different stimuli from the ones that are currently present. For example, an active element might be linked to apple stimuli in the test suite and yet be part of the agent’s mental content when only bananas are in its field of view. One way of identifying the axiom of attention would be follow Damasio (1999) and Metzinger (2003) and look for active connections between active elements linked to the agent's self model and active elements associated with external content. Emotion could be discovered by looking for active elements associated with certain body states.[3]

4.4.3 Metzinger's Constraints

Metzinger (2003) set out eleven constraints that mental content must conform to if it is to be conscious. There is not space to go into the constraints in detail here, but the three most important, which are used to define a minimal notion of consciousness, are the activation of a coherent global model of reality (constraint 3) within a virtual window of presence (constraint 2) both of which are transparent (constraint 7). A system whose mental content conformed to these constraints would have a phenomenal experience of “the presence of one unified world, homogenous and frozen into an internal Now, as it were.” (Metzinger, 2003: 169).

The identification of which parts of the mental content conform to Metzinger's constraints is easier than it seems because Metzinger provides very detailed descriptions of the informational, representational, computational and functional characteristics of the constraints along with some likely neural correlates. All of this can be fairly easily extracted once detailed and systematic XML representations have been created for the system. For example, the presence of constraint 3 (integration within a global model of reality) could be established by looking at the active connections between active elements or possibly using Tononi's methodology. Some of the other constraints, such as transparency, may come for free on systems whose internal states do not have any sensors that could make them objects of representations.

4.5 A Description of the Synthetic Phenomenology?

Given the history of phenomenology, we might expect the final outcome of synthetic phenomenology to be a natural language description. Even if we cannot achieve this at present, it might be thought that this should be the final goal of the procedures outlined in this paper. Viewed from this perspective, the system, test suite and mental content XML would only be the preparatory stages for a traditional phenomenological account of the experiences of COG, CRONOS or IDA.

However, the problems discussed in section 2 make it unlikely that we are ever going achieve fluid natural language descriptions of non-human systems. Instead, it might be much better to treat the XML representations as the best description that we are going to get of the phenomenology of an artificial system. This has the great advantage that it is possible to see what you cannot say. We don't have adequate words in human language to describe a system that can only experience vertical lines, but we can represent such a system accurately using XML, and by looking at the XML we can start to understand how much and how little we can imagine what it is like to be such a system.

The XML descriptions also offer a good starting point for other ways of describing the phenomenology of artificial systems. The suggestions made by Chrisley (1995) about conceptual subtraction, content realization, ability instantiation and self instantiation could all be implemented automatically once the XML formats have been defined. XML would also enable precise comparisons with humans that have deficiencies in the same areas as a machine, and we could use the first person descriptions of these patients to help us imagine what it is like to be such a system. As scanning technology improves, the application of this approach to normal and brain damaged patients will become easier. Research by Kamitani and Tong (2005) on neurophenomenology using combinations of voxels suggests that it might even be possible to start this work today.

5. Discussion

One of the first issues that must be clarified about the XML approach to synthetic phenomenology is that it makes no presuppositions about whether any particular machine is the sort of system that is capable of supporting conscious states. Robots, stones and human beings are all systems that are capable of internal states; all three can be analysed using the XML approach that I have set out here and it will be an empirical outcome of this approach if it turns out that the mental content of a stone is always devoid

---

[3] The identification of planning in an agent's XML descriptions would require a fully temporalised version of the XML approach, which is not covered here.

of phenomenal states. This *empirical* outcome must be distinguished from the *a priori* question about whether certain types of non-human system are capable of supporting conscious states, since it is possible that the XML approach will make predictions about consciousness in systems that we consider highly unlikely to be capable of consciousness – the economy of Bolivia, for example. This *a priori* question is tackled by the ordinal probability scale, set out in Gamez (2005), which evaluates the likelihood that a machine can support phenomenal states by systematically comparing its architecture with the human brain.

It has been suggested that this XML approach to synthetic phenomenology ignores behavioural criteria of consciousness, such as reports that a system might make about its mental contents. If this was thought to be important, then it would be easy to include the actuator outputs in the mental content XML file, so that the external behaviour of the system could be included in the analysis of its consciousness on a moment to moment basis. However, the problem with behavioural criteria for consciousness is that apparently conscious behaviour can be generated by systems that we are reluctant to attribute consciousness to (such as the population of China communicating with radios and satellites), which is why an internal architecture approach has been favoured here.

As this methodology develops there are likely to be a large number of ambiguities about what constitutes an element, how to handle overlapping elements, how to define active connections, the best way to analyse mental content for phenomenal states, and so on. Although these might initially appear to be weaknesses of the method, they are actually strengths because they indicate that synthetic phenomenology has the potential to become a paradigmatic science that can move forward by asking questions and resolving ambiguities such as these. At the moment synthetic phenomenology is so unclear that even its lack of clarity is unclear to it and tightening up the methodology through XML representations would make it capable of asking and answering precise questions and enable it to move forward in a sustainable manner. Different ways of resolving the ambiguities will make testable predictions about the phenomenal states of a machine or organism and as neural scanning becomes better we will actually be able to test these predictions on human beings and eliminate inaccurate methods. In the early stages it is likely that different theories will generate conflicting XML representations. However, this will at least make differences explicit; whereas at present our descriptions of inner states are so woolly and imprecise that disagreement or comparison between methods is rarely an issue.

For reasons of brevity and clarity this paper has set aside questions about the temporal nature of phenomenal experience. One solution to this would be to break the stimuli up into sequences of frames and separate the test suite and mental content into a list of associated XML files. Another temporal problem is that active elements may change as they develop and so it may not be possible to generate a single test suite that is valid for all time. This type of system will have to be retested at regular intervals or have its adaptivity frozen whilst the description of its synthetic phenomenology is taking place.

6. Previous Work

The approach that I have set out in this paper is closest to some of the techniques for representing non-conceptual content discussed by Chrisley (1995). These include content realization, in which content is referred to by listing "perceptual, computational, and/or robotic states and/or abilities that realize the possession of that content" (Chrisley, 1995: 156), ability instantiation, which involves the creation or demonstration of a system that instantiates the abilities involved in entertaining the concept, and two forms of self instantiation, in which the content is referred to by pointing to states of oneself or the environment that are linked to the presence of the content in oneself. Whilst all of these techniques are promising ways of referring to non-conceptual content, it will be very difficult to apply them in practice without a precise way of representing and organizing the computational, and/or robotic states and/or abilities. It is here that XML would be a useful tool since it could represent the structure of the systems that are being analysed along with their inner states when they are exposed to stimuli from the environment. Within the precise framework offered by XML the specification of non-conceptual mental content using Chrisley's techniques would be made considerably easier.

Other related work includes the description of the synthetic phenomenology of Khepera robots by Holland and Goodman (2003) and Stenning, et. al. (2005). In these experiments the internal model of the Khepera is held in a neural network, which stores a linked series of concepts combining sensory and motor information. The synthetic phenomenology of the Khepera is carried out by plotting a graphical representation of the sequence of sensations and movements stored in the neural network. The problem with this approach is that the Khepera is likely to have no notion of colour and a very limited idea about space and so this graphical representation is unlikely to be anything like the Khepera's actual 'mental' content. Another problem is that the graphical representation contains the complete internal model, whereas only a small part of this would be active at any point in time. It is also hard to see how this representation of an internal model could be systematically analysed for signs of consciousness. The XML approach could help with these problems since it offers a highly structured way of representing the current mental content of the Khepera, which could be compared with other robots and systematically analysed for signs of consciousness .

7. Conclusion

This paper has briefly outlined an XML approach to synthetic phenomenology in which XML plays a key role in the description of the conscious and unconscious states of the machine. This has many advantages and could help to circumvent many of the problems associated with the representation of nonconceptual mental content. By describing mental content this concretely it also forces us to face challenging theoretical and methodological questions, which will eventually open up the possibility of a systematic science of synthetic phenomenology that can pose and answer precise questions about the phenomenology of artificial systems.

The XML extracts included in this paper are intended as simple examples to illustrate the main ideas and a great deal more work is needed to turn these starting points into a usable method. Some of this development will be done as part of the work on the CRONOS robot at Essex and Bristol. In the longer term it may be possible to develop a single XML standard for both synthetic and neuro- phenomenology, which would facilitate precise comparisons between humans, animals and machines and enable us to automatically examine all three for signs of consciousness.

Acknowledgments

Many thanks to Owen Holland for feedback about this paper. Thank you also to the EPSRC for funding this work (grant number GR/S47946/01).

References

Igor Aleksander and Barry Dunmall. Axioms and Tests for the Presence of Minimal Consciousness in Agents. In Owen Holland (ed.), *Machine Consciousness*, Exeter: Imprint Academic, 2003.

R. J. Chrisley. Taking Embodiment Seriously: Nonconceptual Content and Robotics. In Kenneth M. Ford, Clark Glymour, & Patrick J. Hayes (eds), *Android Epistemology*, Menlo Park/ Cambridge/ London: AAAI Press/ The MIT Press , 1995.

Antonio, R. Damasio. *The Feeling of What Happens*. New York, San Diego and London: Harcourt Brace & Company, 1999.

S. Franklin, A. Kelemen and L. McCauley. IDA: a cognitive agent architecture. *IEEE International Conference on Systems, Man, and Cybernetics*, 3: 2646-2651, 1998.

David Gamez. An Ordinal Probability Scale for Synthetic Phenomenology. In R. Chrisley, R. Clowes and S. Torrance (eds.), *Proceedings of the AISB05 Symposium on Next Generation approaches to Machine Consciousness: Imagination, Development, Intersubjectivity, and Embodiment* 85-94, 2005.

Owen Holland and Rod Goodman. Robots With Internal Models. In Owen Holland (ed.), *Machine Consciousness*, Exeter: Imprint Academic, 2003.

Y. Kamitani and F. Tong. Decoding the visual and subjective contents of the human brain. *Nature Neuroscience* 8:(5) 679-685, 2005.

Thomas Metzinger. *Being No One*. Cambridge Massachusetts: The MIT Press, 2003.

Thomas Nagel. What is it like to be a bat? *The Philosophical Review* 83: 435-456, 1974.

J. Searle. Minds, Brains and Programs. *Behavioral and Brain Sciences*, 3: 417-57, 1980.

J. Stening, H. Jacobsson and T. Ziemke. Imagination and Abstraction of Sensorimotor Flow: Towards a Robot Model. In R. Chrisley, R. Clowes and S. Torrance (eds.): *Proceedings of the AISB05 Symposium on Next Generation approaches to Machine Consciousness: Imagination, Development, Intersubjectivity, and Embodiment* 50-58, 2005.

G. Tononi. An Information Integration Theory of Consciousness. *BMC Neuroscience* 5:42, 2004.

William Wordsworth. I Wandered Lonely as a Cloud. In Stephen Gill (ed.), *Selected Poems*, London: Penguin, 2004.