The Theory-Implementation-Theory Iterative Cycle: Through the Eye of an AIBO

Joel Parthemore and Ron Chrisley

PAICS Research Group
Centre for Research in Cognitive Science (COGS)
Department of Informatics
University of Sussex
Falmer, Brighton, UK

1 Introduction

Synthetic phenomenology is a term used by Ron Chrisley among others to refer to methods of specifying the contents of experience that go beyond the capacity of language to express, in an attempt to account for phenomena like change blindness. If I want to convey to you the visual experience of a certain agent, I might take a photo or series of photos; but, to be useful, the pictures would need to be altered in ways that do justice to the manner in which the agent is embodied and embedded in its environment. In the human case, not only does the visual field contain blind spots that are not part of the agent's visual experience; the agent's experience may also contain information that is not currently within its field of view.

This paper addresses the practical considerations of translating a theory like synthetic phenomenology into a working model: in this case, one implemented on an AIBO robot. How are general principles translated into lines of code; how does abstract code play out in real-time robotics; and in what ways does practice have consequences for theory?

By using an artifact that we can interact with and manipulate in a way that we cannot for practical or ethical reasons with living organisms, and without any assumption of the artifact having experience, we can use it to model experience. Then we can see how well the model we have created and the specification it implies allow us to predict aspects of the artifact's behavior or explain what it's already done. We can use the results to feed back in to refining the model.

In the process seemingly trivial theoretical issues may prove much more complicated when posed as implementation issues. Some parts of theory may prove untranslatable, either because of hardware limitations or time constraints.

So for example, the AIBO's "eye" is a single fixed camera located at the midpoint between its two apparent eyes. A human eye saccades without any need to move the head. But the only way the AIBO can perform an equivalent saccading is by moving its entire head. Likewise it is impossible to specify precisely where the AIBO should turn its head, or to determine at precisely what angle its currently oriented. This poses challenges for reconstructing its visual "experience".

The more general question addressed by this paper is: what are the benefits and drawbacks when an abstract philosophical theory is translated into a specific computational model? What is the potential for a continuous feedback loop between theory and implemented model?

2 Synthetic Phenomenology: A Primer

The usual methods for specifying the contents of experience are linguistic ones. That's fine for experiences – a conversation, for example – that are or at least seem to be propositionally structured. Most experiences, however, don't appear to have that kind of structure. A narrative account may still capture a great deal: it's often said of a fine piece of writing that you can see the scene or the events that the writer is describing as though they were right before your eyes. In opposition to that, there's the proverb that "a picture is worth a thousand words." A visual experience is not a picture and probably not even a set of pictures; but a set of pictures, suitably modified, may be more true to the visual experience than the narrative.

Synthetic phenomenology is the attempt to find non-linguistic means for specifying the content of experience not by re-creating that experience but by making a model of it (hence the "synthetic" part) that bears enough resemblance[1] to it as to be able to convey something of its content.

A model of experience implies if not requires a model of the experiencing agent. Since the agent one is modeling is presumably embodied, not to mention richly embedded in its environment, the agent model probably needs to be embodied, too (not just a computer simulation) – and interacting with its environment in as rich a manner as materials and thoughtful design permit. There's no need to claim that the agent model actually *is* an agent, nor that the experience models *are* experiences, though they conceivably could be.[2] Synthetic phenomenology

---

[1] I use the word guardedly: no visual likeness need be implied, only, at some level of abstraction, a structural isomorphism.

[2] Call it the Pinocchio Principle: at some point the model of something is no longer "just" a model but becomes an instance of it.

aims to be a robust theory. But in the tension between theory and practical usefulness, synthetic phenomenology emphasizes the usefulness: what matters in the end is *not* whether the models of experience are experiences themselves but what they are able to tell us about the experiences they model.

3 Blinded But Not Blind: Escaping the "Grand Illusion"

Synthetic phenomenology may be useful not just for conveying the content of e.g. a visual experience but for explaining what is *missing* from that content. Take the well-known case of change blindness. "Everyone knows" that, when one's attention is focused very intently on a task, one becomes "blind" to lots of things one would notice otherwise[3]; but most people consider themselves in more usual circumstances, all other things equal, to be competently attentive to changes in their visual field. Nevertheless there are quite dramatic changes that people, unless they have been primed, will consistently miss. When the change is subsequently pointed out to them, it becomes obvious. There's the classic experiment where subjects are being asked for directions, when a door is suddenly carried between them and the person asking for directions. The person asking for directions changes, replaced by someone of different build, wearing different clothes, and only half the subjects notice. **[6, p. 645]**[4]

There are websites[5] you can visit where you're shown one photograph; then there's a blank frame or a flash; then you're shown a second photograph, identical except for one *major* change. It takes a surprising number of times through the sequence for most people before the spot what the change is. Of course, once people *see* the change, they can't *not see* the change! The common thread through all these examples is that between one part of the conversation and the next, or between one image and the next, there is a brief interruption: a distraction, if you will.

Some people have taken the change blindness data as conclusive evidence that human vision is based on a "grand illusion", that the richly detailed and fully colored experience that most people claim to have across their entire visual field[6] is some clever trick that evolution has foisted on us. Ron Chrisley has proposed an alternative explanation[2], one in which the visual experience is

---

[3] "Perhaps you have had the following experience: you are searching for an open seat in a crowded movie theater. After scanning for several minutes, you eventually spot one and sit down. The next day, your friends ask you why you ignored them at the theater. They were waving at you, and you looked right at them but did not see them." [5, p. 1059]

[4] Note that the sample size in these experiments is often quite small: just fifteen in the experiment being cited here.

[5] . . . For example, http://www.usd.edu/psyc301/Rensink.htm.

[6] . . . Not everyone. It might be interesting to compare the visual reports of eyeglass wearers (who may be forcibly more aware of the limitations of their peripheral vision) with people who do not wear any corrective lenses.

richly detailed and colored to the periphery, but not because of any illusion; rather, visual experience includes both current visual input and expectations based on what one would *expect* to see were one to redirect one's attention in some direction – expectations that are based in large part on past (recent) experience.

His explanation for the change blindness data focuses in on the "global flash" or simultaneous changes common to the examples above. The algorithm can be described like this.[7] Location $\boldsymbol{L}$ is defined here as a subregion of the visual field specified by a center point and radius (as opposed to e.g. a single pixel or the output of a single rod or cone).

1. An experiencing subject has an expectation-as-visual-experience that if she were to look at $\boldsymbol{L}$ she would see $\boldsymbol{X}$ if and only if:
    - (a) $\boldsymbol{L}$ is located within her current field of view or
    - (b) She has previously focused her attention on (foveated to) $\boldsymbol{L}$.
2. *If* the scene changes at $\boldsymbol{L}$ from $\boldsymbol{X}$ to $\boldsymbol{Y}$ and $\boldsymbol{L}$ is located anywhere within the current visual field *and* the change is sufficiently localized (i.e., not to count as global), *then* a change flag will be raised to indicate that a local change has been detected.
3. What happens next depends on the number of change flags raised.
    - (a) *If* a single change flag is raised at location $\boldsymbol{L}$, *then* foveal attention will be drawn to $\boldsymbol{L}$.
    - (b) *Otherwise* if several change flags have been raised in addition to the change flag at $\boldsymbol{L}$, such that the number of change flags is less than some threshold n, foveal attention will be drawn to the location of one of those flags according to some measure of salience (or perhaps a location will be chosen randomly).
    - (c) *Otherwise* if the number of change flags exceeds the threshold n, then all of the flags will be reset and ignored and the change treated as a global change (or possibly in some instances, not treated as a change at all).

4 Translating Theory into Practice

The algorithm offers a clear alternative to the "grand illusion" approach. At the same time, it's difficult to see how one would test it with human subjects: how, for starters, can one guarantee what's in the subject's current visual field (the array of sensations available to the visual processing centers of the brain)

---

[7] Note that this is not offered as part of a general theory of vision but rather as an explanation of the change blindness data. (Adapted from [2].)

as opposed to field of view (that portion of the world defined by the current horizontal and vertical sweep of the visual input)?[8] What would a "change flag" look like? How, other than other than by very indirect inference, can we determine what processing is going on?

4.1 Motivations

The same algorithm looks relatively quite easy to implement and test on a computational model, implemented on a robot. The combination of computational model and robot I'll refer to henceforth as the **"robot model"**.[9] It's possible to control for the differences between field of view and visual field. A change flag is simply a data structure. All of this can be implemented quickly using well-known and by-and-large traditional AI programming techniques. One can use one's theory both to explain the robot model's actions and to predict its future actions. Where the actions diverge from one's predictions, the theory can be adjusted and the robot model quickly modified in accordance. There's no need to assume that the robot model really is experiencing anything, and ethically, it's just as well if it's not.

There are other advantages. The limits to what one can do ethically with human subjects are quite closely circumscribed. The ethical constraints are relaxed with other primates – there are plenty of experiments where monkeys are killed at the moment of a particular sensory stimulus and their brains carefully sliced and dyed – but animal testing is highly and increasingly controversial; and when one is studying *experience*, the results obtained with non-human animals may be of limited interest, since there are limited means to interrogate the subjects about their experiences.[10] On the other hand, one can imagine asking an updated version of the AIBO model described in this paper: "did you see a change?" or "what changed?" or "what would you expect to see if you looked *there*?", and getting an intelligible response in English or logically structured pseudo-English or the like. Of course at the point at which one interprets or simply suspects the robot model of having experiences – not just having models of experiences – then the ethical situation changes entirely.[11] However for the simple models currently being implemented, there seems no reason to attribute any actual agency to the robot model and hence no concerns (or at least far fewer concerns) about what

---

[8] In addition to the field of view and the visual field, there is a third field we may wish to distinguish: the field of (conscious) visual experience.

[9] The model described through the rest of this paper is not actually implemented on the robot but on a control laptop or desktop computer, as the processing power on board the AIBO robot we are using is too limited. The laptop or desktop computer communicates with the AIBO by wireless connection.

[10] What this risks giving you, as Blay Whitby has pointed out, is a very anthropecentrically biased view of experience. [7]

[11] Considerable media attention was given recently to a Robot Ethics Charter being drawn up in South Korea, that besides addressing the ethical applications of robots e.g. as companions to the elderly also addresses the potential requirements for robot rights as autonomous agents. One such article can be found at http://news.bbc.co.uk/1/hi/technology/6425927.stm.

may ethically be done to it. (What may be done *with* it may be a different matter!)

4.2 Limitations

As with any exercise in translation, when one is translating a theory into a particular computational model implemented on a certain piece of hardware, one is limited by what can be expressed within that model (indeed, some ideas may not be expressible within *any* computational model!) and by what is possible on the hardware (much of which may not be fully obvious in advance). It may be tempting to limit the theory to what is possible within the model, possibly even without realizing one is doing it. On the other hand, these limitations can potentially be useful in helping set the boundaries of the project, since time and resources are finite.

What sounds easy in theory, goes the proverb, is often much more difficult in practice. Again, this is both a limitation and an advantage: a limitation in that it can bring progress to a halt while an implementation issue that *seemed* trivial gets resolved; an advantage in that it can force one to reconsider each and every aspect of theory.

5 Through the Eye of an AIBO

The current robot model is controlled through a command-line interface that allows one to specify such parameters as whether the model should do change detection or simply construct a composite image representing its visual experience by looking around randomly; whether or not a foveal/parafoveal distinction should be made; whether or not a "blind spot" should be incorporated; and how long the AIBO should look around (i.e., how many saccades it should make). Once wireless contact has been established with the robot, four display windows are opened: one for the raw camera output (the "field of view"), one for the modified camera image (the "visual field"), one for the composite image (the "visual experience") and one for the various stages of change detection (to show, step by step, how the model goes about deciding whether localized change has occurred).

Expressed in its most modest form, the goal of the project is to go some way toward specifying the non-linguistic content of the visual experience of an agent embodied in a manner similar to the AIBO robot and embedded in a similar environment. At the same time, it would be nice to be saying *something* about the content of human visual experience, or at least mammalian visual experience.

<insert picture of computer display here>

5.1 Challenges

This presents immediate problems: the AIBO has only one "eye", a camera located at the midpoint between its apparent eyes, which instead are LED arrays used for communication. The camera cannot be re-directed without moving the entire head, so there is no clear analogy to a saccade. Instead we have decided (perhaps arbitrarily) to count as a saccade a movement of the head along its two degrees of freedom (keeping all other parts of the robot in the same position). Because there is only one "eye", there is no counterpart to binocular vision, and no obvious way to simulate or approximate it.

The AIBO has, of course, no optic nerve, and no equivalent to one. Nonetheless we have chosen to simulate a blind spot similar to the one produced by the optic nerve in humans (and any other animal with a recognizable eye). The AIBO has no equivalent to rods or cones and therefore no difference in concentration of rods and cones across the visual field. Each pixel in its 416 x 320 field of view has more or less precisely the same level of color sensitivity. Nonetheless we have chosen to make a foveal/parafoveal distinction, *as if* the AIBO had a greater concentration of cones toward the center of its visual field, and its peripheral areas contained only a scattering of cones, with a predominance of rods. Such a distinction is necessary if the AIBO's visual field is not to be equivalent to its field of view – as it is not, in all animals that have recognizable eyes. At the same time, making these compromises – abstracting, if you will, from certain awkward features of the AIBO hardware – leaves us open to charges that we are equivocating about just whose visual experience we *are* modeling: is it, indeed, an agent embodied in the manner of the AIBO? ... Or is it a human agent, a mammalian agent, the AIBO itself, or something else?

5.2 Lessons

The robot is not a simulation in a computer program but an "actual" robot – and not just any robot but one mass-produced for the consumer market, with compromises made for the sake of affordability. The software toolkit we are using to interact with the AIBO is beta software, actively in development, with known efficiency issues when used in synchronous mode. (In asynchronous mode, commands are fed to the AIBO and data such as camera images is received from the AIBO asynchronously. That's fine if, for example, the AIBO is moving around, interacting with its environment, and all we want is streaming feedback of what it's "seeing" moment by moment: a "rough feel" for rather than any precise statement of the content of its visual experience. Unfortunately what we're looking for is something *like* a precise statement.)

In asynchronous mode, it is easy to stream thirty frames a second from the AIBO's camera. In synchronous mode, where the control program is telling the robot to look at a certain location, it's difficult to do more than one frame a second, and even then, the AIBO will occasionally (and seemingly randomly) become unresponsive for up to three or four minutes at a time. In the process of looking around it will suddenly freeze; after a period of time it comes back, as if no time had intervened.

We wanted to assemble a version of the camera image at each location – modified for the foveal/parafoveal distinction and blind spot – into a composite image representing the content of the AIBO's visual experience as a combination of current visual field input and (recent) past experience. The difficulty is that the individual images map onto an elliptical plane – the inside surface of a sphere [12] – whereas the composites are necessarily (given our current display hardware) in the Euclidean plane. As a consequence, the individual images in the composite don't quite line up. One gets a feel for the curvedness that the projection should have, but it's awkward. (On the other hand, perhaps that very "awkwardness" is more in keeping with visual experience, which isn't necessarily tidy.) A non-trivial 2-D graphics transformation would clean this up.

Another difficulty with lining up the individual images is that the AIBO's head never turns precisely where the control program tells it to. Asking it its current location helps but isn't entirely accurate either.

Because of these limitations, change detection is only possible when the AIBO is looking at one spot for a period of several seconds, so that it can compare its current visual frame with a previous one of the same location. That seriously limits the opportunities for testing the change blindness algorithm described above. To further complicate matters, the AIBO's head never stays entirely still; it drifts, which can introduce the *appearance* of change when no change has taken place.

In response to these limitations, we have incorporated a simulation of boredom into the model. When the AIBO gets "bored", it begins looking around, thereby assembling a composite view constituting its visual experience of the current

---

[12] Even that is an approximation, since the camera is taking a picture not of a specific ***point*** but of an area.

setting. When it stops being bored, it focuses on one spot until it either notices change (and saccades – reorients its head – toward the change[13]) or becomes bored again.

5.3 Consequences for Theory

Translating a theory into an implemented model is always a process of discovering implicit assumptions and being forced to reflect upon not-fully-explicit assumptions, both in the theory and in the implementation. So for example, the current robot model assumes a stable world, where the only thing that is moving, most of the time, is the AIBO's head. That limits how much the model can inform the theory about change recognition and change blindness. The model is embedded in its environment in only a very weak sense. On the other hand, allowing the AIBO to move around and interact with its environment, and having the environment likewise interact with the AIBO, would make the model much more complex. Instead of a single 2-D composite image, would one need a collection of 2-D composites or a 3-D composite (or something of fractal dimension between 2- and 3-D)?

The current model has no representation of objects, only of individual pixels in its visual field and visual experience. Without some understanding of objects, motion detection is probably not possible; instead of motion detection, one has a much more austere notion of change detection. One simple, non-conceptual way of representing objects would be as areas of stability within a changing visual field that move as a contiguous whole.

The hope is that, as the model matures, it will become possible to compare its response to various change recognition and change blindness scenarios with human responses to the same situations. If the robot model and the human subject respond in the same way, so much the better for the theory. Of course, the model is *not* an agent, and an agent embodied in the same manner as an AIBO would not be, in any usual meaning of the word, human. Comparing its responses with human responses might still be misleading. Of course one could take the view that the change blindness algorithm and the robot model itself are just means to an end: non-linguistic methods of specifying the content of experience, which is the particular insight of synthetic phenomenology. The change blindness algorithm is the theory; synthetic phenomenology is the approach. The meta-theory is asking whether synthetic phenomenology provides a *useful* and *accurate* means of specifying non-linguistically the content of experience (or of certain kinds of experience), toward practical objectives like explaining the change recognition / change blindness data.

---

[13] It barks as well, to announce that change has been detected.

6 Application to Other Domains: Conceptual Modeling

Synthetic phenomenology and non-conceptual content are among the research interests for the secondary author on this paper. The primary author's interests are more toward theories of concepts, ways of modeling concepts, and conceptual knowledge representation formalisms (e.g., how might one specify the conceptual content of a virtual reality environment in one consistent and coherent manner). The robot model and the robot itself are simply means toward this end.

6.1 Application of Synthetic Phenomenology

Synthetic phenomenology has applications, hopefully, far beyond the toy world of the AIBO robot model. One possible application is specifying the content of concepts. On the one hand, it seems difficult if not impossible *not* to use conceptual language to talk about concepts: what a concept is in general, or what a particular concept is. On the other, using concepts to define or otherwise specify concepts leads to all the usual self-referential paradoxes like Grelling's Paradox. What's to prevent a concept like the concept of all possible concepts, which must by definition contain itself, and contain itself containing itself, and so on in an infinite regression? Worse, one could imagine dividing all concepts into self-referential concepts (like the concept of concept: i.e., the concept CONCEPT) and non-self-referential concepts. Then one has, seemingly, the concept of all self-referential concepts and the concept of all non-self-referential concepts. The difficulty is: is the concept of all non-self-referential concepts itself a self-referential or non-self-referential concept? Providing a non-linguistic and furthermore non-conceptual means of specifying the content of concepts offers the possibility to avoid those pitfalls without the excessive measure of trying to banish the paradoxes through some artificial division of concepts into concepts, meta-concepts, meta-meta-concepts and so on, in the spirit of Russell and Whitehead's Theory of Types.[14] If concepts can be specified using a logical "language" that invites a visual metaphor, in the way that logic, number theory and geometry are all different ways of expressing the same mathematical notions, then the opportunity arises both to explain how concepts can, from a functional viewpoint, be defined by other concepts, and from an operational viewpoint be specified in completely non-conceptual terms.[15] It also raises the possibility of applications like e.g. tools for helping people build external models of their conceptual domains (as I described in [4]).

6.2 Application of This (Cyclic) Approach

The iterative theory-model-theory approach suggested by this project and this paper is common to a well-established tradition of hands-on philosophy, as ex-

---

[14] For an excellent presentation on the limitations of the Theory of Types, at least when it is applied outside of set theory, see [3, p. 42].

[15] My thanks to Mike Beaton and Tom Froese for independently raising this idea for me.

emplified in cognitive science. It has been described in very similar-sounding terms by Dustin Stokes, in discussing the Drawbots project at the University of Sussex.**[1]** For a theory of concepts, it suggests a way to be precise in an area where precision is *very* difficult to achieve, since, regardless of how we *specify* the contents of concepts, our experiences of them, like our experiences of everything else, are conceptualized[16]. On the one hand, the design of a knowledge representation formalism as an instrument of representing knowledge – and arguably, nearly *any* project in AI research – assumes a (possibly entirely implicit) theory of concepts: what a concept is, how concepts relate to language, how concepts relate to agents and referents, whether concepts exist independently of or are inter-defined by one another, and so on. On the other, it provides an opportunity for discovering our implicit assumptions and re-addressing our often not fully explicit ones. "What is a concept?" is a deceptively simple-looking question. Through an iterative theory-model-theory process, one might hope to move closer to an understanding of "the nature of what is being represented".

One project along these lines is the CYC project, summarized neatly through the projects website: http://www.cyc.com. Whether the CYC project started out with an explicit theory of concepts is not fully clear; what is clear is that the practical demands of the project – capturing the essence of "common-sense" reasoning toward tools for improved data mining, text analysis, and more intelligent search engines for the Web – have made explicit a theory of concepts, one that makes a lot of controversial assumptions: most particularly, that knowledge is mainly conceptual, that all conceptual knowledge can be propositionally expressed, and that all propositions can be captured in a form of higher-order predicate logic. Concepts are specified conceptually: a difficulty that an approach from synthetic phenomenology might hope to avoid.

7 Conclusions

Synthetic phenomenology, as described by Ron Chrisley, provides non-linguistic means to specify the content of experience, or at least of certain kinds of experience. Sometimes a picture really is worth (more than!) a thousand words. The change blindness account he has proposed seeks to account for the empirical evidence on change blindness while avoiding the "Grand Illusion" argument. Visual experience (and by extension, other modal forms of experience) is not an illusion foisted on us by evolution or anything else; rather, it is simply the case that visual experience is more than current visual input. We see what we expect to see, based largely on past experience.

---

[16] ...Though the conceptualization may be a matter of degree, as Ron Chrisley has pointed out (personal communication), it is, it would seem, impossible to have an experience that is *entirely* non-conceptualized. (By way of contrast, Mike Beaton has argued [personal communication] that the degree to which something is an experience is the degree to which it is conceptualized.)

The robot model described in this paper presents an implementation of synthetic phenomenology and an opportunity to translate change blindness theory into practice. Using a robot in place of a human subject offers many limitations (not least that a robot is *not* a human![17]) but also lots of advantages, particularly ethical ones: though there are plenty of ethical restrictions on what one can do *with* a robot, there are relatively few, it would seem, on what one can do *to* a robot. A robot model such as this paper describes also gives the opportunity to make explicit and "real" notions like "localized change flag" that otherwise have no more reality than "center of gravity" or "infinity": highly useful concepts both, but you won't find them physically instantiated anywhere, whereas in a computational model, you can point and say, *that* is the localized change flag!

Both a limitation and an advantage of the approach taken is that it places significant limits on what is possible. The AIBO is a robot mass-produced for the consumer market, designed to be relatively affordable, with plenty of compromises made in its design. Beyond that, the AIBO is not only not a human, it is not, in any recognizable sense, an experiencing agent. The AIBO has no counterpart to binocular vision, no counterpart to an optic-nerve blind spot, no obvious counterpart to a saccade. On the other hand, highlighting the differences in this way may provide a means toward better understanding them. As the model matures, it may have much to say about the strengths and weaknesses of one or another account of the change blindness data – just so long as we don't ignore all the ways in which an AIBO is nothing like a person.

Synthetic phenomenology has applications far beyond the toy world of the AIBO. We cannot help experiencing concepts, as we experience all things, *conceptually*; but specifying the contents of concepts conceptually invites an infinite regress and self-referential paradox. Synthetic phenomenology offers the possibility of an operational account of concepts that does not itself rely on conceptual language. Likewise the iterative theory-model-theory approach is part of a well-established tradition of "hands-on" philosophy, and combining it with a synthetic phenomenology approach to establishing a theory of concepts can reasonably be expected to be quite fruitful.

References

[1] Jon Bird and Dustin Stokes. What's alife got to do with art? Presentation to the Centre for Computational Neuroscience and Robotics (CCNR)

---

[17] "The AIBO can turn its head 182°, considerably further than most people can; but its field of view is considerably smaller (56.9° horizontally, as opposed to nearly 180° in humans – including peripheral vision – and a full 360° in some birds). With saccading (and no head movements), a typical person can see somewhere close to 240°. The AIBO, with its combined head movements/saccades, can see approximately the same range. (Of course with head movements and saccades, and with no other body movements, the typical person can achieve 360°.)" [2]

Workshop on Computational Biology, Robotics, and Related Work held at the University of Sussex, UK (13 April 2007), April 2007.

[2] Ron Chrisley and Joel Parthemore. Synthetic phenomenology: Exploiting embodiment to specify the non-conceptual content of visual experience. *Journal of Consciousness Studies*, 2007 (forthcoming).

[3] Douglas R. Hofstadter. *Gödel, Escher, Bach: An Eternal Golden Braid.* Penguin Books, 20th anniversary edition edition, 2000.

[4] Joel Parthemore and Jon Taylor. Charley: A linguistic formalism applied to writing environment design. *Intelligent Tutoring Media*, 3(2/3):85–92, May/August 1992.

[5] Daniel J. Simons and Christopher F. Chabris. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28:1059–1074, 1999.

[6] Daniel J. Simons and Daniel T. Levin. Failing to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review*, 5(4):644–649, 1998.

[7] Blay Whitby. Why the turing test is ai's biggest blind alley. Technical Report CSRP 568, University of Sussex, Falmer, Brighton, UK, 1997.