

An Externalist Process-Oriented Framework for Artificial Consciousness

¹Riccardo Manzotti, ²Vincenzo Tagliasco

¹ULM University, via Carlo Bo, 8, 20143, Milan (Italy)

²University of Genoa, viale Francesco Causa, 16143, Genova (Italy)

¹riccardo.manzotti@iulm.it

²vincenzo@dist.unige.it

Abstract

In this paper, we present a view of conscious perception that supposes a processual unity between the activity in the agent and the perceived event in the external world. It is a kind of radical externalism (both vehicle and content) that suggests that the boundary of the agent's conscious mind are much larger than those of its body. We suggest a process-based approach as an explanation of ordinary perception and other variants of phenomenal experience such as illusions, memory, dreams, and mental imagery. This approach provides new insights into the problem of conscious representation in the brain and phenomenal consciousness. It is a form of anti-cranialism different from but related to other kinds of externalism. Eventually, we will sketch the relation between this model and the capability of developing new goals in an agent.

Unity between the environment and the agent

During the last ten years, interest in the scientific understanding of the nature of consciousness has been rekindled (Hameroff, Kaszniak et al., 1996; Jennings, 2000; Miller, 2005). To date, a satisfactory and accepted framework has not been achieved either because experimental data is scarce or because a misleading theoretical standpoint is assumed.

The effort for a scientific understanding of consciousness has been flanked by a related approach named *artificial consciousness* (sometimes *machine* or *synthetic consciousness*) aiming at reproducing the relevant features of consciousness using non biological components (Buttazzo, 2000; Holland, 2003; Manzotti, 2003; Adami, 2006; Chella and Manzotti, 2007; Adami, 2006; Bongard, Zykov et al., 2006). This new field has strong relationships with artificial intelligence and cognitive robotics. Most researchers focus on a weak form of artificial consciousness. They try to replicate forms of access-consciousness (similar to Baars' global workspace). On the contrary, we try to address the problem of phenomenal consciousness.

Implicit in most theories of conscious perception is the supposition that, although an external event and its representation in the brain are causally connected, they are

nevertheless separate. As against this, we outline a process oriented framework applicable to perception which is a foundation for the proposal that there is a unity between the "external world" and the "perceived world". The proposal is that the classic separation between subject and object must be reconceived so that the two, while maintaining their identities as different perspectives on a process, actually occur as a unity during perception. This leads us to sketch out a new view of consciousness, which can be summarised by saying that consciousness consists in the occurrence of a unity between the brain and the part of the world that is being attended. Here, we use the word 'unity' in the same sense in which we say that a magnetic field is a unity that can be described in terms of the different categories of a south and a north pole. We distinguish between ordinary visual perception, the actual presence of the object of perception in the external world, and other cases of conscious experience, when the object is no longer there or has never been such as memory, mental imagery, and dreams.

According to a traditional point of view, the object is separated from the subject. There is much merit in this, since the two are conceptually separable. Yet that presumed separation is less obvious than commonly supposed – an underlying processual unity might exist. In many instances, for example the touch of the skin or the taste of food, this continuity is manifest. For other senses such as vision or hearing, a fairly convincing case can be made for the occurrence of an underlying unity. Furthermore, the separation between subject and object leads to the long debated issues of mental representations and 'qualia' where the former usually have a functional role and the latter ought to be responsible for qualitative experience (access consciousness vs. phenomenal consciousness).

However, as soon as the separation between functional and phenomenal properties is assumed, a problem arises. Why should a functional structure have a phenomenal side? As far as we know, in the brain there is nothing like the content of our phenomenal experiences. "If consciousness is a brain process ... , how could my conscious experience of my grandmother have these features — such as the colour of her eyes — that no brain process could have?" (Place, 1956). When a subject tastes the flavour of a piece

of chocolate, in his/her brain there is nothing with the property of that flavour. On the contrary there are neural patterns with completely different properties. Why should the latter be experienced as the former? Nobody knows. Furthermore, nobody knows how phenomenal experience, supposedly emergent from neural patterns functionally linked with external objects, is related with the physical properties of the piece of chocolate.

The multiplication of entities is the result of the above mentioned assumption about the separation between perceiver and perceived or between subject and object. If the perceiver is physically separate from the external world, what is perceived must be something else – that is, a representation of the external object. This view, attributed to René Descartes, posits a separation between the “external world” and the “mental world” (for an historical overview see Manzotti, 2006). What is a representation, or better a re-presentation? Rather naively but worryingly, it is something that re-presents something else.

Up to now, the nature of the relation is a tantalizing mystery. Different solutions have been proposed: correlation, causation, lawlike causation, emergence, identity, supervenience. None has proved to be completely satisfactory. Once the separation between the subject and the object is accepted, both the representational and the phenomenal aspect of neural processes present awkward problems (Fodor, 1981; Chalmers, 1996; Lehar, 2003; Kim, 2005).

With respect to the conscious experience of the world, the brain carries the same burden as the XVII century soul. The brain is supposed to be able to interpret the incoming electric signal and produce a conscious experience of the world by means of some internal coding (a view labelled as “codicism” by Mark Bickhard Bickhard and Terveen, 1995). The quest for neural correlates of consciousness (NCC) is currently based on this undemonstrated framework (Koch, 2004) as well as most consciousness inspired models in AI (Baars, 1988; Aleksander, 2000; Hesslow, 2002; Baars and Franklin, 2003; Franklin, 2003; Kuipers, 2005; Shanahan, 2005). The conjectural neural coding in the brain should perform the miracle of changing the water of neural activity into the wine of phenomenal experience (the nice metaphor is due to Humprey (1992)).

Once the hypothesis of the separation between subject and object (brain and external world) is accepted, a dualistic view of perception cannot be avoided. In turn, dualistic approaches lead to the issue of re-presentation and to that of the nature of phenomenal experience. Given the shortcomings of these approaches, it is worthwhile and perhaps timely questioning such an hypothesis (Rorty 1979; Searle 1980; Fodor 1981; Millikan 1984; Dretske 1993; Dretske 1995; Tye 1996; Clark 1997; Bickhard 1999; Metzinger 2003). Is the subject really separate from what s/he perceives? Is the perception really separate from the perceived?

For instance, when a red patch is perceived, starting from the external object to the activity in the brain, an uninterrupted continuous physical process takes place.

Rays of light in terms of actual physical quanta are continually bombarding the retina, which in turn is being successively hyperpolarized (at the receptors) and then depolarised beyond, with the nerve impulses thus set up being transmitted continuously to the relevant processing-perceptual areas of the brain. This is a continuous process which occurs as long as the red patch is visually perceived – physically spanning from the patch and the brain. From a physical point of view, there is no separation between the external object and the activity inside the brain. There is a continuous chain of causally related physical phenomena. When the view of the patch is interrupted, even if momentarily, the link is broken and we cease to see the red patch. Thus, there is nothing strange in supposing that there is some kind of unity, embodied by a process, between subject and object. Rather, it is surprising, given the actual and easily demonstrable continuity, that this option has received so little attention. Moreover, the supposition of continuity leads to an alternative view of conscious experience that is worth looking into.

Internalism vs externalism

Currently, the majority of scientific and philosophical literature on consciousness is biased by a seldom challenged assumption – the separation between the subject and the object. Although it is obvious that the body of the subject is separate from the body of the object, it is by no means so obvious that the mind is confined by the same boundaries of the brain. “Where does the mind stop and the rest of the world begin? [...] Someone accepts the demarcations of skin and skull, and say that what is outside the body is outside the mind” (Clark and Chalmers, 1999, and more recently Wilson, 2004). Indeed, there are many phenomena which extend beyond the boundaries of the body (behaviors, actions, perceptions, ecological processes). The mind could be one of them.

With regard to the nature of the mind, two very broad standpoints must be considered: internalism and externalism. The former states that our consciousness is identical (or correlated) to the processes, events or states of affairs going on *inside* the boundary of our body (or brain). The latter affirms that our consciousness might depend partially or totally on the events, processes or state of affairs *outside* our head or even *outside* our body.

Most current approaches to the problem of consciousness lean towards the internalist viewpoint (Crick, 1994; Edelman and Tononi, 2000; Metzinger, 2000; Rees, Kreiman et al., 2002; Crick and Koch, 2003; Koch, 2004). However, this approach raises several conundrums. If the mind is entirely located or dependent on events or states of affairs located inside the cranium, how can they represent events taking place in the external world? Consciousness appears to have properties which differ from anything taking place inside the cranium (Place, 1956). Spurred on by common sense, literature has revealed a very strong impulse to “etherealize” or “cranialize” consciousness (Honderich, 2000). The internalist perspective has

consistently led to dualism and still promotes a physicalist version of dualism by endowing the brain (or a brain subset, the Neural Correlate of Consciousness) with the same role as the dualistic subject. Koch's recent book (Koch, 2004, p.87), endorses an unbiased internalist view with respect to consciousness and the brain: "The entire brain is sufficient for consciousness – it determines conscious sensations day in and day out. Identifying all of the brain with the NCC [Neural Correlate of Consciousness], however is not useful because likely a subset of brain matter will do."

On the other hand, many authors, like ourselves, have questioned the separation between subject and object – between representation and represented. They are looking for a different framework in which subject and object are two different perspectives on the same physical phenomenon. Their views could be labeled as some kind of externalism (Hurley, 2001; Hurley, 2006).

According to Mark Rowlands, there are two variants of externalism: *content* externalism and *vehicle* externalism. The former corresponds to the "idea that the semantic content of mental states that have it is often dependent on factors [...] that are external to the subject of that content" (Rowlands, 2003, p. 5). The latter is more radical and affirms that "the structures and mechanisms that allow a creature to possess or undergo various mental states and processes are often structure and mechanisms that extend beyond the skin of that creature" (Rowlands, 2003, p.6).

In the following paragraphs, we will present a version of vehicle externalism (Manzotti and Tagliascio, 2001; Manzotti, 2003; Manzotti, 2005; Manzotti, 2006b; Manzotti, 2006a).

The Enlarged Mind: An Externalist Framework

In order to support vehicle externalism as a framework for artificial consciousness, we outline a framework in which the separation between the conscious perception of the world and the perceived physical world is not reconsidered. The rationale is the following. The agent is conscious of that parts of the environment that produce effects thanks to the agents's body and neural structure. Objects, patterns, wholes are singled out by the agent's body. Consciousness is the way in which the environment is entangled in the behavioral history of the agent.

The rainbow is perhaps the best example in which there is no separation between the observed object/event and the observer. The rainbow is not a physical material object, but rather a process that needs the interaction with the agent's body. If no observer were there, would the rainbow take place? No, it would not, because the light rays would continue their travel in space without interacting and, eventually, they would spread in the surrounding environment. Their opportunity to take place as a rainbow. On the contrary, if an observer were there, the converging rays of light would have hit his/her photoreceptors and a fast but complex chain of physical processes would have

continued from the retina to the cortical areas up to a point where the process corresponding to the rainbow would reach its end.

The rainbow is not a thing: it is a *process*, in which there is an entanglement between a physical condition and the agent's body. The light rays do not constitute a distinctive unity (the rainbow) unless and until they are embedded in a process. The occurrence of the rainbow depends not only on the presence of the physical conditions given above and the observer, but on a causal continuity between the twos. This approach suggest a kind of direct realism based on the identity between the physical process embedding the perceived object and phenomenal experience itself.

We elsewhere proposed to call this process – which is constitutive of what there is and what we perceive – an *onphene*, derived from the Greek words *ontos* (what there is) and *phenomenon* (what appears) (Manzotti and Tagliascio, 2001; Manzotti, 2003; Manzotti, 2006a; Manzotti, 2006b). It refers to a process in which the traditional distinction between cause and effect (perceiver and perceived, representation and represented, subject and object) is missing.

The traditional problems of consciousness can thus be reconsidered once an externalist and process-based standpoint is adopted. The world in which each subject is living is no longer a private bubble of phenomenal experiences concocted by the brain. Each subject lives in and experiences the real world – the two being different descriptions of the same process. As conscious agents, we are part of a physical flow of processes possible thanks to our physical structure. We venture to suggest that these processes have the same properties of our own experiences as well as the same properties of the external world. Thereby we suggest that these processes are identical with conscious experience.

According to the view presented here, the mind is identical with everything the subject is conscious of – everything being a process. Furthermore, the existence of what the mind is conscious of depends on the occurrence of those processes that are identical with the mind itself.

We consider the physical process that begins in the external world and ends in the brain as a unity. Such process occurs thanks to the brain, to the body and to the surrounding environment. The rainbow is an excellent example of an process, in which the act of observation, the observer and the observed entity cannot be split. All occur jointly. They are different ways to look at the same process. But the example of the rainbow, though a very compelling one, is not unique in leading to this conclusion. We propose that most perceived objects (if not all) have a structure analogous to that of the rainbow. The relevance of this argument lies in the fact that the brain is not self-sufficient with respect to mental events. The brain could thus be envisaged as the end part of a larger network of physical processes.

This is a view which could be considered a kind of radical externalism. The mind is literally and physically identical with a collection of processes spanning in time and space

beyond the boundaries of the brain and the body. It is also a realist standpoint since it assumes that the experience of the world regards the world itself and not a mental representation of it.

Difficult cases revisited

Any realist theory faces the apparently insuperable counterexamples of situations like dreams, hallucinations, illusions or afterimages. All situations in which there seems to be no external object to be perceived. We would like to address quickly some of these situations in order to show that a process-externalist-realist approach is going to overcome most of the problems.

After-images

If we have a red after-image as a result of a flashbulb going off, the spot we 'see' in front of the photographer's face looks red, even though there is no such spot. There are different possible interpretations of this fact. A simple one suggests that my feeling of red is due to the instantiation of the same kind of neural activity which is normally elicited by red objects. Thus, the phenomenal redness would be independent from the external world. Yet, another explanation is possible. Due to the intensity of the flashbulb, my photoreceptors do not work in the usual way. Their sensitivity to certain wavelength has been dramatically modified and reduced. As a result we continue to be causally in relation with the flashing of the flashbulb. For instance, if the flash was so intense that for a few seconds we are completely blind, the result is that my brain is no longer causally connected with the surrounding environment. From a causal point of view, the state of my receptor is connected with the flash and not with the current scene in front of my eyes. That's why we have a feeling of red or white instead of having the qualia related with what is in front of my eyes.

There is also another kind of after-images that is worthwhile to be examined. Imagine to stare at a cross in the middle of four squares of four different bright and saturated colours – let's say red, green, blue and yellow. After ten seconds someone substitutes the four coloured squares with a white surface. You will still see four squares for a while. They will have complementary colours of the original one. A naïve scientist could explain the phenomenon by saying that some neural activity elicited by the original colours is still going on and producing a phenomenal quale. We suggest an alternative explanation. For the sake of simplicity, let us consider only one of the squares – the red one. Staring at the cross you let an area of your retina being stimulated by a red light for ten seconds. As a result there was an adaptation by those receptors that are more sensitive to red light. When the white surface was shown, that area of the retina reacted in a different way to the white light than the other areas, namely that area was less sensitive to red components than to green and blue ones. In that part of the visual field you were almost blind

to red and not to green and blue. Since in the white light there are all components of the spectrum, you saw a yellow square, which is the sum of green and blue. In sum, the after-image was not inside the head. The after-image was that part of the world that you were not blind to – namely the green and blue components of an area in the white surface.

From the point of view of phenomenal experience there are two kinds of after-images: those that correspond to a persistent perception and those that correspond to changes in the normal behaviour of receptors. In the former case, the triggering event continues to be the origin of the process ending in the brain. In the latter case, we perceive something that is normally present in the environment and that is hidden to us by the normal behaviour of receptor.

Non veridical perception

We see the stick in the water as if it were bent. Of course it is not. The quale we have does not have the right representational content. However, from the point of view presented here, when we see the stick out of the water there is one kind of process; when we see the stick bent in the water there is a different process. It is no surprise that there is a different perception and a different phenomenal experience. It is not simply a matter of looking different. The two situations correspond to two different processes. Once again, simply dropping a substance oriented view of properties allows a reshaping of traditional problems.

Phosphene

If the visual cortex of a sighted subject is stimulated either by a strong mechanical blow or by direct electric stimulation, the subject reports having a visual phenomenal experience of sparkles of light called phosphenes. Where is the real light yet? Phosphenes seems to pose a real threat to the approach presented here.

A different explanation is the following. What takes places in the brain as a result of the stimulation of a visual area by a non visual stimulus (pressure on the eyeball, electricity, bumping) is historically related with a very long past history of visual stimuli. As a result it maintains a causal continuity with them. On the other hand, what would be the phenomenal perceived content if the eyes were disconnected from visual stimuli from the very beginning and yet subjected to pressure? According to explanations analogous to the law of specific nerve energies, perceived phenomenal contents should be visual phosphenes. On the other hand, according to the hypothesis of the continuity with the external world, the phenomenal content should be of tactile nature. The eyes should work as poor tactile receptors. Furthermore, if they were eventually exposed to light, by some technical or surgical means, they should elicit "tactile" phosphenes.

Interestingly, although the literature is rather poor on cases like them (for a review see (Senden, 1932), there is the famous case reported by William Cheselden of a born blind patient that, after an operation that partially restored his

sight, reported the first visual experiences as having a tactile phenomenal quality: “When he first saw [...] that he thought all Objects whatever touch’d his Eyes, (as he express’d it) as what he felt, did his Skin”. Of course, more empirical data is necessary to draw a final conclusion.

Teleologically open systems and externalism

Is it possible to design and implement an architecture exploiting an externalist-oriented view? Or, more modestly, is it possible to derive some practical constraints from an externalist view? The common ground could be a teleological entanglement between the agent and the environment. The continuity outlined by the externalist view could be achieved by a system teleologically open.

We should begin with some words of caution concerning the scope and limits of this final section. The basic idea under development is that a conscious agent is a physical system that entangles the environment in its teleological history. As a result, a conscious agent is capable of deriving from its environment not just new categories and representations but also new goals. It is thus instructive to compare artificial system deprived of consciousness with supposed conscious biological beings such as mammals. The hypothesis we make is that the latter will have a much higher degree of teleological openness – that is, the capability of acquiring new goals. Acquiring new goals should be the glue that keeps together newer and newer processes originating in the environment.

Current implementations of artificial systems focus on implementation of intelligent algorithms to achieve a fixed goal (or a fixed set of goals). Conscious subjects are capable of developing unpredictable and unexpected new goals.

Artificial systems are frequently designed with a fixed set of goals. Designers focus their efforts to find “how” those goals can be achieved. Learning is usually defined as a modification in agents’ behavior: a modification driven by a goal. Various learning paradigms focus mostly on this modification of behavior. Supervised and unsupervised learning as well as reinforcement learning are valid examples (Sutton and Barto, 1998, p.3): they are based on fixed goals. For instance, in a reinforcement learning based agent “the reward function [corresponding to the goal] must necessarily be unalterable by the agent” (Sutton and Barto, 1998, p.8). On the contrary many biological systems are capable of developing partially or totally unpredictable goals. There is evidence that such capability is greater in humans and mammals, which also the species are less controversially assumed to have some form of consciousness.

First, to develop new goals is important since the environment cannot be completely predicted at design time. Therefore a truly adaptive system must be able to add new goals, not only to modify its behavior in order to perform optimally on the basis of some fixed criteria.

The behavior of behavior-based artificial agents depends on experience and goals defined elsewhere at design time (Arkin, 1999). Motivation-based agents begin to show the capability of developing new goals (Manzotti and Tagliasco, 2005). In complex biological systems, behavior depends on experience and goals; yet, goals are not fixed. Goals are the result of the interaction between the subject and its environment. In many complex biological systems, it is possible to distinguish between phylogenetic aspects and ontogenetic ones, nature versus nurture (Gould, 1977; Elman, Bates et al., 2001; Ridley, 2004).

What is a goal? An agent’s goal is an event whose occurrence is more probable thanks to the agent’s structure (cognitive and bodily). Goals are embedded in causal structures that link the past with the future, the environment with the agent.

It is possible to classify artificial agents accordingly to their degree of teleological plasticity: fixed control architectures, learning architectures and goal generating architectures. In the first case, the system has no capability of modifying how it does what it does. In the second case, the system is capable of modifying its behavior to fulfill some a priori target. The system is capable of modifying *how* it behaves. In the third case, the system is capable of modifying not only *how* it does what it does, but also *what* it does.

System with a fixed control architecture have a fixed causal structure. There is no ontogenesis whatsoever. Notwithstanding the behavioral complexity of the system, everything happens because it has been previously coded within the system structure. A mechanical device and a complex software agent are not different in this respect: both are pre-programmed in what they must achieve and how they must achieve it. Nothing in their structure is caused by their experiences. Suitable examples of this category are Tolam’s artificial sow bug, Braitenberg’s thinking vehicles (Braitenberg, 1984), Brooks’ artificial insects and recent entertainment robots like Sony AIBO and Honda’s humanoid ASIMO (2002).

A different level of dependency with the environment is provided by architectures that can learn *how* to perform a task. Behavior-based robots can be classified in this category. Systems based on artificial neural networks are well-known examples of this kind of architectures. These systems determine how to get a given result once they have been provided with a specific goal. The goal can be given either as a series of examples of correct behavior (supervised learning) or as a simple evaluation of the global performance of the system (reinforcement learning) (Sutton and Barto, 1998). In both cases some kind of learning is applied. These systems lack the capability of creating new goals. By controlling its motors a behavior-based robot can learn how to navigate avoiding static and dynamic obstacles. However the goal behind this task is defined by the *a priori* design of the system. There are several examples of this kind of learning agent: Babybot at LIRA-Lab (Metta, Manzotti et al., 2000), Cog at MIT (Brooks, Breazeal et al., 1999).

An agent, which learns both *how* to perform a given task and *what* task, corresponds to a teleologically open architecture. This is the case for most, if not all, mammals; it is true for primates and for human beings. They are systems capable of developing new goals that do not belong to their genetic background. For their development, these systems depend more on the environment than the previous two categories. A system belonging to the first category does not depend on the environment for what it does or for how it does what it does. A system belonging to the second category does depend on the environment for how it does what it does, but not for what it does. A system belonging to the third and last category depends on the environment both for what and for how it does what it does.

We suggest that the kind of environmental entanglement achieved by a teleologically open architecture is the same exploited by human beings when they perceive consciously. It is a framework that could be used to deal with phenomenal consciousness in the field of artificial consciousness. If it would prove right, it would not require any kind of biological neural activity like those implicitly assumed by most of the NCC biased literature. Consciousness would be identical with the right kind of causal entanglement between an agent and its environment.

References

- Adami, C. (2006). "What Do Robots Dreams Of?" *Science* 314 (5802): 1093-1094.
- Aleksander, I. (2000). *How to Build a Mind*. London, Weidenfeld & Nicolson.
- Arkin, R. C. (1999). *Behavior-Based Robotics*. Cambridge (Mass), MIT Press.
- Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge, Cambridge University Press.
- Baars, B. J. and S. Franklin (2003). "How conscious experience and working memory interact. ends in Cognitive Sciences, Volume 7, Number 4, April 2003." *Trends in Cognitive Sciences* 7 (4): 166-172.
- Bickhard, M. and L. Terveen (1995). *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*. Amsterdam, Elsevier Scientific.
- Bongard, J., v. Zykov, et al. (2006). "Resilient Machines Through Continuous Self-Modeling." *Science* 314 (5802): 1118-1121.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. Cambridge (Mass), MIT Press.
- Brooks, R. A., C. Breazeal, et al. (1999). "The Cog Project: Building a Humanoid Robot". in *Computation for Metaphors, Analogy, and Agents*. Nehaniv. Berlin, Springer-Verlag. 1562.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. New York, Oxford Univ. Press.
- Clark, A. and D. Chalmers (1999). "The Extended Mind." *Analysis* 58 (1): 10-23.
- Elman, J. L., E. A. Bates, et al. (2001). *Rethinking Innateness: A Connectionist Perspective on Development*. Cambridge (Mass), MIT Press.
- Fodor, J. A. (1981). *Representations: philosophical essays on the foundations of cognitive science*. Cambridge (Mass), MIT Press.
- Franklin, S. (2003). "IDA: A Conscious Artefact?" in *Machine Consciousness*. O. Holland. Exeter (UK), Imprint Academic.
- Gould, S. J. (1977). *Ontogeny and Phylogeny*. Cambridge (Mass), Harvard University Press.
- Hesslow, G. (2002). "Conscious thought as simulation of behaviour and perception." *Trends in Cognitive Sciences* 6 (6): 242-247.
- Honderich, T. (2000). "Consciousness as existence again." *Theoria* 95: 94-109.
- Hurley, S. L. (2001). "Perception and Action: Alternative Views." *Synthese* 129: 3-40.
- Hurley, S. L. (2006). "Varieties of externalism". in *The extended mind*. R. Menary. Aldershot, Ashgate publishing.
- Kim, J. (2005). *Physicalism, or Something Near Enough*. Princeton, Princeton University Press.
- Koch, C. (2004). *The Quest for Consciousness: A Neurobiological Approach*. Englewood (Col), Roberts & Company Publishers.
- Kuipers, B. (2005). "Consciousness: drinking from the firehose of experience". in National Conference on Artificial Intelligence (AAAI-05).
- Lehar, S. (2003). "Gestalt Isomorphism and the Primacy of Subjective Conscious Experience: A Gestalt Bubble Model." *Behavioral and Brain Sciences* 26 (4).
- Manzotti, R. (2006). "Consciousness and existence as a process." *Mind and Matter* 4 (1): 7-43.
- Manzotti, R. and V. Tagliascio (2005). "From "behaviour-based" robots to "motivations-based" robots." *Robotics and Autonomous Systems* 51 (2-3): 175-190.
- Metta, G., R. Manzotti, et al. (2000). "Development: is it the right way towards humanoid robotics?" in ISA-6, Venezia, IOS Press.
- Place, U. T. (1956). "Is consciousness a brain process?" *British Journal of Psychology* 47: 44-50.
- Ridley, M. (2004). *Nature via nurture*. Great Britain, Harper.
- Rowlands, M. (2003). *Externalism. Putting Mind and World Back Together Again*. Chesham, Acumen Publishing Limited.
- Senden, V. (1932). *Raum und Gestaltauffassung bei operierten Blindgeborenen vor und nach der Operation (conception of space and gestalt in congenital blind children before and after surgery)*. Leipzig, Verlag Jd Ambrosus Barth.
- Shanahan, M. P. (2005). "Global Access, Embodiment, and the Conscious Subject." *Journal of Consciousness Studies* 12 (12): 46-66.
- Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning*. Cambridge (Mass), MIT Press.
- Wilson, R. A. (2004). *Boundaries of the Mind The Individual in the Fragile Sciences. Cognition*. Cambridge (Mass), Cambridge University Press.