Review of General Psychology
2008, Vol. 12, No. 2, 137–146

1089-2680/08/$12.00 DOI: 10.1037/1089-2680.12.2.137

Understanding Understanding: An Evolutionary Autonomous Agent Approach

Donald S. Borrett
Toronto East General Hospital

Hon C. Kwan
University of Toronto

Two poles of understanding define the hermeneutic circle: gestalt understanding, in which the experience of a text or piece of music is apprehended as a unity, and conceptual understanding, in which the work is broken down into more determinate parts. In our sensory-motor interaction with the world, the environment is composed of discrete objects, but there is also an omnipresent gestalt background of nonrepresentational practices that confer meaning on these objects. It is argued that neuroscience can provide an explanation of how a physical system instantiates these types of understanding. A naturalized version of temporality, that extended temporal horizon that frames the flux of sensible experience and confers meaning on it, is equated with the dynamical system concept of temporal hierarchical organization. With this naturalized concept of temporality, it is demonstrated how the two poles of understanding that define the hermeneutic circle can emerge in an evolutionary autonomous agent as those dynamics best suited to maintain optimal grip in that particular agent.

Keywords: evolutionary autonomous agents, dynamical systems theory, hermeneutic circle, phenomenology, temporality

In the title of his book *Understanding Understanding*, Heinz von Foerster (2002) captures the scientific essence of second order cybernetics. Its goal is not simply to externally understand how an organism effectively interacts with its environment, but to outline how the organism's nervous system can instantiate the faculty of understanding that allows it to interact the way it does. In the humanities (here understood as a philosophical, rhetorical, historical, or aesthetic analysis), however, the understanding of an experience is analyzed through the reasons for, and implications of, behavior. The exercise of understanding understanding in relation to the humanities, then, requires an explanation of how the physical operation of the nervous system results in a faculty of understanding in which the interpretation of experiences plays a causal role in governing behavior.

D. S. Borrett, Division of Neurology, Toronto East General Hospital, Toronto, ON, Canada; H. C. Kwan, Department of Physiology, University of Toronto, Toronto, ON, Canada.

Correspondence concerning this article should be addressed to D. S. Borrett, Division of Neurology, Toronto East General Hospital, Toronto, ON, Canada. E-mail: d.borrett@utoronto.ca

As a first-order approximation, we can say that there are two fundamentally different ways that we make sense of the world. In a conceptual or predicative understanding, the world is broken down into determinate parts and then reassembled according to a rational framework. This framework has predictive and conceptual value and is the basis of all scientific understanding. The second way of organizing experience is through a gestalt or prepredicative understanding. In this case, our experience is apprehended as a unity (albeit one that can be further subdivided into parts) that is the basis from which the parts acquire meaning (Taylor, 2005). In the humanities, this relationship between the parts and the whole is referred to as the hermeneutic circle and has been addressed by such authors as Schleiermacher (2000); Heidegger (1962), and Gadamer (1994). A literary or philosophical text, although assembled from individual words, is understood through a gestalt process that confers new significance to individual words or passages depending upon the context of the reader, allusions to other texts, prior historical processes, and many other things. Likewise, a piece of music, although composed of individual notes, is generally apprehended as a whole which then permits the understanding of individual notes or passages. In all of our immediate sensory-motor interactions with the world, our environment is composed of discrete objects but there is an omnipresent gestalt background of nonrepresentational cultural practices that confer meaning to these objects based on our experience. This article is concerned with the question of how an artificial system can instantiate both conceptual *and* gestalt understanding and how their instantiation can interact in a physiologically plausible manner to result in the hermeneutic circle.

It is important to stress that we are dealing with understanding or intelligibility and not the simple attribution of knowledge to another agent. Introspective psychology, however, will not be enough. The psychologist seeks protocols derived from by established conceptual frameworks, focuses on objects that can be isolated in the laboratory, and attempts to establish empirical results that can be linked to neural correlates. As a result, introspective psychology already assumes a conceptual framework that requires explanation. To provide a causal explanation for the origin of conceptual understanding and a physical explanation for the relationship that exists between conceptual and gestalt understanding, it is necessary to start at a level in which no assumptions are made concerning the structure or nature of intelligibility. This more fundamental level is the sensory-motor interaction of an embodied agent situated in a particular environment. Although the branch of philosophy known as phenomenology is interested in the same mental content as psychology, it makes no assumptions regarding its underlying structure but rather explores the broader meaning of the content and places it in ordinary human experience such as temporality, intersubjectivity and language (Varela & Shear, 1999). The phenomenologist grounds all experience in the most fundamental level of an organism interacting with its environment and assumes that conceptual thought is derived from this level of experience. It is for this reason that our approach will be phenomenological. The first part of the project describes the phenomenology of understanding that will need to be accommodated in the physical implementation. The second part naturalizes the phenomenological data in such a way that phenomenological accuracy and biological plausibility are maintained.

The working hypothesis of the naturalization of phenomenology was stated by Varela (1999): “Phenomenological accounts of the structure of experience and their counterparts in cognitive science relate to each other through reciprocal constraints” (p. 305). The emergence of well-defined biological attributes must be consistent with a phenomenological description that stays close to lived experience. A detailed description of the nature of this project and its relationship to the philosophy of mind and the biology of life has recently been summarized by Thompson (2007). Since this naturalized phenomenology is the means to instantiate a physical model of understanding, it requires tools that accurately capture features of our experiences evident on phenomenological analysis. The use of any formalism that deals with determinate perceptual representations or motor programs will not be useful in this endeavor for three reasons. First, if the physical model uses determinate representations in its implementation, it assumes that which needs to be explained. Second, if perceptual experience is viewed as the sum of determinate entities, then the relationship of the gestalt experience to the more determinate experience will remain problematic. If gestalt experience is simply the sum of the individual components, then it contains no information beyond that of those individual components. Third, aside from the conceptual difficulties of constructing a whole out of determinate parts, phenomenology suggests that all perceptual experience is indeterminate (Kelly, 2005).

The use of dynamical systems theory, however, provides a framework that can capture the relevant phenomenological aspects of experience. The dynamical systems approach is increasingly being utilized both in modeling cognition (Chemero, 2000; Clark, 1997; Freeman, 1999; Kelso, 1995; Port & van Gelder, 1995; Thelen & Smith, 1994; van Dujin, Keijzer, & Franken, 2006; Varela, Thompson, & Rosch, 1995) as well as in the development of autonomous agents (Beer, 1997, 2003; Blynel & Floreano, 2002; Der, Hesse & Martius, 2006; Nolfi & Floreano, 2000; Nolfi & Marocco, 2001; Paine & Tani, 2005). With dynamical systems theory, all perceptual experiences can be placed in a spectrum of indeterminancy extending from less determinate to more determinate. In fact, a clear separation between the two types of understanding is not possible: they are end-points on a spectrum and not separate worlds (Taylor, 2005).

Evolutionary autonomous agents (EAAs) are robots or robot simulations which evolve autonomously following a fitness function. Being embodied situated agents, they provide simple systems in which the structure and dynamics of a neural network controller can be described and followed as the agent interacts with its environment during assigned tasks that relate to human cognition (Ruppin, 2002). The motivation for using autonomous agents to explore the naturalization of phenomenology has been nicely captured by Dennett (2007) in a response to a number of commentaries on his notion of heterophenomenology. Dennett made the claim that phenomenological analysis of consciousness as lived conscious experience invariably leaves out "all the grubby details of implementation to some later investigation that is not even outlined" (p. 268). This criticism of phenomenology is not new and certainly the lack of practical applications of phenomenological principles is an obvious hindrance to its acceptance in the scientific arena. He points out that "the goal of implementing features in a real working model constrains and provokes the imagination of the theorist" (p. 268) and suggests a yardstick by which to judge the relevance of any concept, phenomenological or otherwise, used in the debate about human consciousness. If there is any doubt about the scientific relevance or accuracy of any proposal, he suggests, translate it into robot-talk: "One of the great virtues of robots is that all the causation involved is garden-variety causation, well-behaved at the microlevel, however startling the macrolevel products are" (p. 249). If we are to make claims concerning the indeterminacy of perceptual experience or the emergence of conceptual imagination from a more fundamental sensory-motor level, scientific acceptance demands the invocation of nothing more than what present science has in its arsenal of techniques. With regard to the neurosciences, this arsenal involves causally-mechanistic local interaction rules between neurons in the human brain or neuron-like elements in autonomous agents. Although it is overly optimistic at this point in time to expect a robotic model to capture all elements of a phenomenologically based program to explain the whole/part dichotomy, it is reasonable, as Dennett suggested, to expect the outline of how such a program could evolve.

The difficulty with a robotics program that does not address phenomenology has been aptly expressed by di Paolo (2002): "In a nutshell, my claim is that autonomous agents still lack one fundamental property of what makes real animals cognitive agents. This property is that of intentional agency . . . A robot failing in its performance does not show any signs of preoccupation—failure or success do not affect its structure in any dangerous way, nor does any form of concern accompany its actions simply because the desired goal is not desired by the robot but by the designer." He goes on to say that to serve as an accurate model of a living system, a robot must structure its activity and environment into a space of meaning which is defined as that which distinguishes between what is relevant or irrelevant to the robot itself. This issue of meaningfulness to the agent itself will structure the arguments concerning the relationship between the neurosciences and the humanities. Humanistic models of understanding assume that our conscious experience does influence behavior in some fashion. In discussing Hamlet's obsession with killing Claudius, for example, it is not sufficient to observe Hamlet's stereotyped behavior and infer that the protagonist is stuck in a fixed point attractor. Rather, we need to describe how Hamlet wrestled with his obsession by using reason either to support or refute his revenge. It is only through the active participation of a character's conscious thought to a behavioral outcome that a tragedy can arise. Even if manufactured with the physical ability to frown and produce tears and programmed with a self-destructive tendency, Asimo the Honda robot will never be a tragic figure. If conscious experience were an epiphenomenon or simply not relevant to the understanding of ourselves, then the entire exercise of wondering why an author chose a certain phrase, how our present situation influences a reading of a text written in a different cultural context, or why a musician chooses a particular interpretation of a composer's work is negligible, and we would have to ask with Quine (1969) ". . .why all this creative reconstruction, all this make-believe?" (p. 269)

In the humanities, however, to understand others is to attribute a mode of understanding to them that exists in entities that are external to those others. To address the whole/part problem of the hermeneutic circle, this involves more than attributing atomistic beliefs and desires to others since these faculties are conceptual (Taylor, 1993). What is needed is the attribution of a stream of conscious experience to others, since the conscious experience will incorporate both types of understanding needed to complete the hermeneutic circle.

It is, therefore, the modeling of conscious experience and how this experience influences motor behavior that is of interest in attempting to bridge the humanities and neuroscience. Since meaningfulness occurs within this conscious experience, successful modeling of this experience in an autonomous agent will address di Paolo's concern with contemporary robotics and its relationship to living systems. The issue of artificial consciousness, as opposed to artificial intelligence, is receiving increasing attention and criteria for synthetic consciousness in artificial agents have been proposed (Baars, 1997; Chella & Manzotti, 2007; Franklin, 1995; Freeman, 1997; Haikonen, 2003; Sun, 2001). For a system to instantiate aspects of our experience that capture essential elements needed to explain those activities implicit in the humanities, including the hermeneutic circle, a role for a conscious stream in modifying behavior is required. This requires an explicit statement of what constitutes a conscious experience in the model and this statement must be both phenomenologically accurate and biologically plausible. As opposed to the approach that looks at the fully developed human brain and attempts to localize the contents of consciousness in its structure (Metzinger, 2000), this approach attempts to identify fundamental aspects of phenomenal consciousness and to determine what is needed for a physical system to instantiate these characteristics.

Crucial to this analysis of consciousness is the temporality of experience (Clark, 1998; Gallagher, 1998). Varela (1999) explicitly stated that "given the importance of the topic of the experience of temporality, let it be clear that I consider this an acid test of the entire neurophenomenological enterprise" (p. 267). Temporality refers to the experience of time not just as a series of "now" experiences but rather as a temporally extended window. The temporal horizons envelope that flux of experience that constitutes our immediate sensible experience and are the condition for the possibility of its meaning. In gestalt understanding, these temporal horizons are indistinct and experienced more as potential experience. In conceptual understanding, these horizons are experienced more determinately and are of more immediate significance. It has already been suggested that such temporal horizons can be naturalized and identified with the dynamical system concept of temporal hierarchical organization (THO), which is a measure of the type and distribution of time scales in a robot's neural network dynamics (Borrett et al., 2006). In this analysis, changes in the nature of these multiple time scales that occurred with breakdown in the skillfully coping robot were consistent with the changes that we experience in our temporal horizon under similar conditions. Breakdown during skillful coping was associated with a shift to shorter time scale components in the network dynamics and hence with an experience which is more immediate and determinate. Although the importance of a naturalized version of temporality in the conceptualization of meaningfulness was underlined in this study, how the agent itself could access this information was left open. Access to this temporal hierarchical organization, however, may be the means by which the agent itself experiences a horizon with which its sensible experiential flux is framed and in which meaningfulness emerges. A model of conscious experience can thus be suggested based on Merleau-Ponty's notion of the sensible flux (the visible) framed by its temporal horizons (the invisible) (1968). In the model the visible translates into a trajectory in the phase space defined by the dynamical neural network and the invisible translates into a measure of the THO that feeds back onto the network itself.

Phenomenological Background

In "Seeing Things in Merleau-Ponty," Sean Kelly (2005) discusses Merleau-Ponty's understanding of what it is to see objects as three-dimensional entities despite only seeing them in perspectival presentations. He argues that even though we cannot sensibly see the back of an object presented in particular perspectival presentation, the back of the object is still experienced perceptually because it is incorporated into our motor intentionality. Since the back of an object, in his view, is not a perceptual absence, its indeterminancy is a positive phenomenon. Kelly elaborates this thesis concerning the indeterminancy of perceptual experience, compares it to other approaches that do not consider the indeterminate as perceptually real and suggests that perception is indeterminate because it is essentially normative rather than descriptive. With this perspective, an object is an event to be optimally achieved rather than simply an entity to be described. Determinate sense data describe the world. We interact with our environment not to describe it but to get an optimal grip on it (Merleau-Ponty, 1962). Dreyfus (2005) was the first to underline the importance of this fundamental principle in artificial intelligence and to demonstrate its relevance to cognitive science. The perceptual field is not the simple sum of individual sensible components, but is the contextual gestalt that emerges under the defining normative principle. If a more determinate experience is needed to maintain optimal grip, the relationship between the originary perceptual gestalt and the more determinate experience is still subsumed under the normative principal; as such, the two models of understanding that we are trying to reconcile here are always already imbricated with each other.

The perceptual field is initially completely unstructured but becomes more determinate as the organism interacts with the environment in response to its needs: "A need, unlike a desire, is originally given as a pure restlessness; as the consciousness of one's undirected activity. It begins with the sense of a lack in oneself, without any sense of what would remove that lack; it begins with the sense of an indeterminate lack of something-or-other . . . Our quest of discovery is . . . initially directed not to get what we want but to discover what we want to get" (Todes, 2001, p. 177). Only after the perceptual field becomes structured through the resolution of needs can we ever come to have a desire for an object, adopt a belief about an object, or form an object-centered concept. Although the resolution of the tension produced by meeting a need may have been our fundamental normative principle as developing children, as adults we also seem to interact with our environment to optimize some kind of balance. If we are looking at a picture, we stand at a distance that balances richness and resolution (Kelly, 2005). If we are exploring our environment, our movements display a balance between constraint and spontaneity. This balance between conflicting elements is needed if we are to remain maximally adaptive and yet skillful in our interactions with the environment. A shift in balance toward constraint adversely affects our ability to learn and adjust to new situations. A shift toward spontaneity adversely influences skillful behavior in predictable situations. This optimization of balance between conflicting tendencies that were initially structured by the meeting of needs seems to capture the sense implied the phrase, "optimizing grip."

Evolutionary Autonomous Agents

The importance of the use of genetic algorithms in evolutionary robotic design has been underlined previously (Nolfi & Floreano, 2000). Genetic algorithms allow a robot to evolve under the constraint of a fitness function without any explicit guidance from the programmer. Any algorithm that directly teaches the robot the nature of a particular object presents the object to the robot determinately and hence will not be able to contribute to any program that wishes to discover the causal origin of determinate conceptualizations in the robot's dynamics. It is clear, however, that a fitness function could serve as the robotic equivalent of the phenomenological concept of need. The EAA does not initially know what it should do but acquires the knowledge of what should be done only after it successfully accomplished the task as defined by its fitness. Because this initial activity is not object directed, an object does not result in the resolution of the need but rather a particular behavior in that situation was responsible for the resolution. Context is incorporated into the motor intentionality without the need for itemizing entities in the sensible field. Since genetic algorithms use optimization of fitness, the behavior that eventually evolves is the one that optimally resolves the need.

Although genetic algorithms provide a mechanism by which an optimization schema can be introduced as the agent is evolving and results in a normative model of perception, this does not answer the question of why, in humans with skillful coping, there seems to be an optimal balance between richness and resolution. It is observed that robots can "overevolve" and these robots tend to have fixed responses that do not adjust for changes in the environment. This overevolution typically occurs in a simple environment in which the task is straightforward. For the organism to have the ability to adjust its behavior and hence expand its experience of the world, the environment in which a robot evolves has to be complicated enough to necessitate the need for preserved adaptability. Without the agent being exposed to enough contextual baggage to necessitate the incorporation of flexibility in its functioning, the development of this flexibility will be hampered. In addition, perturbations need to be introduced during the learning paradigms to allow the development of the dynamics that will become associated with conceptualization. Only by experiencing breakdown during the course of its evolution can an agent develop the cognitive structure is be able to adapt to new situations.

A demonstration of this approach already exists (Borrett et al., 2006). EAAs were evolved that could successfully perform a two-step motor task. An obstacle was introduced into the arena after the agents had become proficient at the task and, as expected, the vast majority of agents could not accomplish the task with presence of the obstacle. Several agents were found that were able to negotiate around the obstacle despite never having previously experienced an obstacle. In these agents, the mechanisms to negotiate an obstacle had obviously emerged spontaneously. If these agents are mated and evolved further in more complicated environments with additional types of obstacles, one would expect the development of a population of agents that would not only function well during time of predictability but also have the ability to adjust for unexpected disturbances. It was observed that the dynamics of the agents differed depending on whether the environment was predicable (skillful coping) or whether the agent experienced the obstruction and had to move around it. In the latter case, the dynamics of the controller was characterized by increased percentage of short time scale components in its temporal hierarchical organization. This shift to higher frequencies was not simply the result of hitting an obstacle and sudden arrest of movement. Rather, in continuing around the obstacle, the agent utilized a dynamics characterized by shorter time scale components. This result could simply be dismissed as a switch from an automatic to a reactive strategy, and from a third-person point of view this is reasonable. The point, however, was that if the THO is identified with the phenomenological concept of temporality and if breakdown is one means by which our gestalt experience associated with skillful coping changes into conceptual experience, then the simple EAA paradigm could provide a naturalized explanation of the differences in the experience associated with skilful coping and conceptual thought. In conceptual thought associated with breakdown, the temporal horizon that frames the sensible flux becomes more determinate.

Genetic algorithms are still supervised learning paradigms; although they incorporate essential features that are needed to instantiate a normative model of sensory-motor interaction, they require an external observer to choose those agents that best meet the fitness requirements to proceed to the next generation in the evolutionary procedure. This being the case, the EAA paradigm is optimally suited for phylogenetic modeling with the fitness function being equated with survival. For ontogenetic descriptions, however, (i.e., for the agent to function autonomously without external supervision) it needs to know when things are going well and when it has optimal grip. By identifying optimal grip with the balance between constraint and spontaneity and by equating this balance with a particular distribution of time scales in the dynamics of the robot controller, a mechanism is suggested by which the agent has access to this information.

Any time signal can be characterized by its distribution of frequency components or time scales (Sprott, 2003) and this distribution has been identified with the THO. In properly functioning systems, a power law $(1/f)^{\alpha}$ distribution is commonly observed. Systems as diverse as stock market fluctuations (Mandelbrot, 1997), reaction time distributions in psychology (Van Orden, Holden, & Turvey, 2003), and distribution of earthquake magnitudes (Bak & Tang, 1989) all obey a $(1/f)^{\alpha}$ power law. Although the mechanism of this distribution is usually not known in a particular situation, this power law signature is ubiquitous. Bak (1996) has suggested that all complex systems tend to organize to a preferred state, which he called self-organized criticality, in which there is no characteristic spatial or temporal scale, that is, which obeys a $(1/f)^\alpha$ power law. In the case of EAAs, the dynamics of the controller of an agent whose behavior has no structure in its interaction with the environment would represent a totally random signal that would be characterized by a $(1/f)^\alpha$ distribution with $\alpha = 0$. The more predicatable the agent, the more constrained the response, and the larger the value of α. It appears that an $\alpha = 1$ represents the best balance between two tendencies of spontaneity (randomness) and constraint (predictability). If a measure of the THO of an EAA can be calculated in real time and if the agent has access to this measure, then the agent would have information as to whether its dynamics are balanced between spontaneity and constraint. Maintenance of a particular dynamic with $\alpha = 1$ can be used as an additional constraint in the evolution of the agents. Since the THO can also be viewed as the naturalized version of temporality, an agent that has access to it also will fulfill the model of a conscious experience defined by the sensible flux framed by its temporal horizons. The problem, then, reduces to the problem of finding a means by which an agent can get on online, real time, measure of its THO.

An optimized network must have long-range memory or power law correlation. A mathematical structure that implements this power law correlation is a fractal structure, exhibiting self-similar scale-invariant properties. These properties permit embedding of statistically similar dynamics at different time scales suggesting a mechanism by which information concerning the nature of the THO over an extended epoch can be determined through analysis of the THO over a shorter epoch, such as a short interval preceding and up to the present sensible experience. The details of such an implementation remain a goal for the real working model.

In an agent that is skillfully coping, that is, is interacting in a manner that is expected and without perturbations, its dynamics follow a trajectory in phase space and its THO, although available to the network, may not be specifically accessed. The type of understanding characterized by such a situation is prepredicative or gestalt understanding. With breakdown, the dynamics changes so that the THO is characterized by an increased number of short time scale components characteristic of an experience that is more immediate and determinate. Because of the self-similarity implicit in a power law relationship, details of future or past states can be accessed if needed for the resolution of breakdown, possibly through replay of the dynamics in a contracted time scale, a phenomenon that has been documented in animals (Diba & Buzsaki, 2007; Euston, Tatsuno & McNaughton, 2007; Lee & Wilson, 2002). Any access of the network to a particular time scale is, by definition, accessing an item in conscious experience. There is no part of the controller that accesses this information. Rather, the fitness function dictates the mechanisms by which this information is accessed and utilized by the network as a whole. Since the fundamental sensory-motor interaction is normative, what actually occurs after breakdown is the production of an activity that maintains optimal grip on the environment, where optimal grip includes maintaining a $(1/f)^1$ distribution. But after resolution, the (1/f) distribution is different than if the disturbance had not occurred because synaptic weights are activity dependent. In the EAA model developed, this activity-dependent change in synaptic weights was implemented through the use of Hebbian synapses as suggested by Urzelai and Floreano (2001). The fact that synaptic weights are activity-dependent is essential for this type of model. If the network has no permanent indication that a particular trajectory was followed in the resolution of the disturbance, this information will not be available in the fractal dynamics at a later time. Optimal grip, in this situation, thus entails the possibility of history-dependent adaptation.

Extrapolation to the Humanities

The sensory-motor experience of an embodied agent situated in a particular environment was used as the model of the hermeneutic circle. Similarly, with exposure to a text, "with the first vision, the first contact, the first pleasure, there is initiation, that is, not the positing of a content, but the establishment of a level in terms of which every other experience will henceforth be situated" (Merleau-Ponty, 1968, p. 151). This is the gestalt understanding that grounds more detailed analyses of the text. In addition, "each time we want to get at it immediately, or lay hands on it, or circumscribe it, or see it unveiled, we do in fact feel that the attempt is misconceived, that it retreats in the measure that we approach. The explication does not give us the idea itself; it is but a second version of it, a more manageable derivative" (Merleau-Ponty, 1968). Any more detailed access to the text, such as the meaning of a word or phrase, is associated with an altered dynamics that differs from the dynamics established with the first contact. Even though the text is composed of words and phrases, their understanding derives from and depends on the gestalt level of understanding. These ideas mirror the dynamical changes in sensory-motor interaction reflective of gestalt and conceptual understanding associated with skillful coping and breakdown in evolutionary autonomous agents. Such an interpretation obviously requires a dynamical view of language. Elman (1995) has proposed such a model and has suggested that "instead of symbolic rules and phrase structure trees, we have a dynamical system in which grammatical constructions are represented by trajectories in phase space." A reader's approach to the text is still subsumed under the overarching requirement to maximize grip. If the reader is a literary scholar, the types of understanding experienced during the reading of the text and the relationship between the different levels of understanding will differ from that of a person who is simply reading for pleasure. The normative nature of perception allows for a causal explanation not only for the types of understanding associated with reading a text but also the explanations for why particular types of understanding occur at particular times.

In the case of music, the similarities with the sensory-motor model of the hermeneutic circle are even more clear cut. Voss and Clarke (1975) were the first to note that the loudness fluctuations and the pitch fluctuations in music exhibit a (1/f) power spectrum. By representing the auditory appreciation of a piece of music as a trajectory in phase space, analysis of the understanding associated with the piece is similar to that provided for sensory-motor interaction and suggested in the analysis of a text. The details of the relationship between the levels of understanding will again be subsumed under the fitness function of optimizing grip.

Discussion

In the model of conscious experience suggested, feedback of the temporal hierarchical organization imbedded in the dynamics of the agent into the network itself was the mechanism by which optimal grip was instantiated and by which the naturalized version of temporality instituted. The fractal structure of the dynamics suggested a mechanism by which the THO in a contracted time scale could surrogate the THO for the entire epoch. In addition, because the synaptic weights are activity dependent, details of the THO change with time and are dependent on the specific history of the agent. The computational consequence of this is that history dependent adaptation is possible. The phenomenological consequence is that the personal history and the meaning of experiences in this history play a role in the maintenance of optimal grip. The poles of understanding that define the hermeneutic circle occur in the stream of conscious experience as those dynamics best suited to maintain optimal grip in that particular agent.

With regard to the implications of these ideas for the humanities, Merleau-Ponty (1968) has summarized it well:

> Literature, music, the passions, but also the experience of the visible world are—no less than is the science of Lavoisier and Ampère—the exploration of an invisible and the disclosure of a universe of ideas. The difference is simply that this invisible, these ideas, unlike those of science, cannot be detached from the sensible appearance and be erected into a second positivity. The musical idea, the literary idea, the dialectic of love, and also the articulations of the light, the modes of exhibition of sound and of touch speak to us, have their logic, their coherence, their points of intersection, their concordances, and here also the appearances are the disguise of unknown "forces" and 'laws.' (p. 149)

It is the discovery of these "forces" and "laws" that bind the visible and the invisible that motivate the type of modeling presented and that frame the entire project of the naturalization of phenomenology. Dennett (2007) challenged that phenomenology leaves out "all the grubby details of implementation." Developing a neuroscientific explanation for the existence of the hermeneutic circle is indeed a formidable task, but the development of phenomenologically based autonomous agents (who not only "speak" but live Dennett's "robot-talk") has shown that such implementation is possible.

References

Baars, B. (1997). *In the theater of consciousness.* New York: Oxford University Press.

Bak, P. (1996). *How nature works: The science of self-organized criticality.* New York: Corpernicus.

Bak, P., & Tang, C. (1989). Earthquakes as a self-organized critical phenomenon. *Journal of Geophysical Research, 94,* 15635–15637.

Beer, R. (1997). The dynamics of adaptive behavior: A research program. *Robotics and Autonomous Agents, 20,* 257–289.

Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior, 11,* 209–243.

Blynel, J., & Floreano, D. (2002). Levels of dynamics and adaptive behavior in evolutionary neural controllers. In B. Hallam, D. Floreano, J. Hallam, G. Hayes, & J. Meyer (Eds.), *From animals to animats: Proceedings of the seventh international conference on simulation of adaptive behavior* (pp. 272–281). Cambridge MA: MIT Press.

Borrett, D., Khan, S., Lam, C., Li, D., Nguyen, H., & Kwan, H. (2006). Evolutionary autonomous agents and the naturalization of phenomenology. *Phenomenology and the Cognitive Sciences, 5,* 351–363.

Chella, A., & Manzotti, R. (2007). Towards robot conscious perception. In *Artificial Consciousness.* Exeter, United Kingdom: Imprint Academic.

Chemero, A. (2000). Anti-representationalism and the dynamical stance. *Philosophy of Science, 67,* 625–647.

Clark, A. (1997). *Being there: Putting brain, body and world together again.* Cambridge, MA: MIT Press.

Clark, A. (1998). Time and mind. *Journal of Philosophy XCV, 7,* 354–376.

Dennett, D. (2007). Heterophenomenology reconsidered. *Phenomenology and the Cognitive Sciences, 6,* 247–270.

Der, R., Hesse, F., & Martius, G. (2006). Rocking stamper and jumping snakes from a dynamical system approach to artificial life. *Adaptive Behavior, 14,* 105–115.

Diba, K., & Buzsaki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience, 10,* 1241–1242.

di Paolo, E. (2002). Organismically-inspired robotics: Homeostatic adaptation and teleology beyond the closed sensory-motor loop. In K. Murase & T. Asakura (Eds.), *Dynamical systems approach to embodiment and sociality* (pp. 19–42). Adelaide, Australia: Advanced Knowledge International.

Dreyfus, H. (2005). Merleau-Ponty and recent cognitive science. In *Cambridge Companion to Merleau-Ponty.* Cambridge, United Kingdom: Cambridge University Press.

Elman, J. (1995). Language as a dynamical system. In R. Port & T. van Gelder (Eds.), *Mind as motion*. Cambridge, MA; MIT Press.

Euston, D., Tatsuno, M., & McNaughton (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science, 318,* 1147–1149.

Franklin, S. (1995). *Artificial minds.* Boston: MIT Press

Freeman, W. (1997). Three centuries of category errors in studies of the neural basis of consciousness and intentionality. *Neural Networks, 10,* 1175–1183.

Freeman, W. (1999). *How brains make up their minds.* London: Weidenfeld and Nicolson.

Gadamer, H. (2004). *Truth and method.* New York: Continuum.

Gallagher, S. (1998). *The Inordinance of time.* Evanston, IL: Northwestern University Press.

Haikonen, P. (2003). *The cognitive approach to conscious machines.* Exeter, United Kingdom: Imprint Academic.

Heidegger, M. (1962). *Being and time.* New York: Harper and Row.

Kelly, S. (2005). Seeing things in Merleau-Ponty. In *Cambridge Companion to Merleau-Ponty* (pp. 222–222). Cambridge, United Kingdom: Cambridge University Press.

Kelso, S. (1995). *Dynamic patterns: The self-organization of brain and behavior.* Cambridge, MA: MIT Press.

Lee, A., & Wilson, M. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron, 36,* 1183–1194.

Mandelbrot, B. (1997). *Fractals and scaling in finance.* New York: Freeman.

Merleau-Ponty, M. (1962). *Phenomenology of perception.* London: Routledge and Kegan Paul.

Merleau-Ponty, M. (1968). *The visible and the invisible.* Evanston, IL: Northwestern University Press.

Metzinger, T. (2000). *Neural correlates of consciousness.* Cambridge, MA: MIT Press.

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence and technology of self-organizing machines.* Cambridge, MA: MIT Press.

Nolfi, S., & Marocco, D. (2001). Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences, 120,* 287–310.

Paine, R., & Tani, J. (2005). How hierarchical control self-organizes in artificial adaptive systems. *Adaptive Behavior, 13,* 211–225.

Port, R., & van Gelder, T. (1995). *Mind as motion: Explorations in the dynamics of cognition.* Cambridge, MA: MIT Press.

Quine, W. V. O. (2000). Epistemology naturalized. In S. Bernecker and F. Dreske (Eds), *Knowledge:*

Reading in contemporary epistemology (pp. 266–278). New York: Oxford University Press.

Ruppin, E. (2002). Evolutionary autonomous agents: A neuroscience perspective. *Nature Reviews: Neuroscience, 3,* 132–141.

Sun, R. (2001). Computation, reduction and teleology. *Cognitive Systems Research 1,* 241–249.

Taylor, C. (1993). Engaged agency and background. In C. Guignon (Ed.), *The Cambridge companion to Heidegger*. Cambridge, United Kingdom: Cambridge University Press.

Taylor, C. (2005). Merleau-Ponty and the epistemological picture. In *The Cambridge companion to Merleau-Ponty*. Cambridge, United Kingdom: Cambridge University Press.

Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of perception and action.* Cambridge, MA: MIT Press.

Thompson, E. (2007). *Mind in life: Biology, phenomenology and the sciences of mind.* Cambridge, MA: The Belknap Press of Harvard University Press.

Todes, S. (2001). *Body and world.* Cambridge, MA: MIT Press.

Urzelai, J., & Floreano, D. (2001). Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evolutionary Computation, 9,* 495–524.

van Dujin, M., Keijzer, F., & Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensory-motor coordination. *Adaptive Behavior, 14,* 157–170.

Van Orden, G., Holden, T., & Turvey, M. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General, 132,* 331–350.

Varela, F. (1999). The specious present. In J. Petitot, F. Varela, B. Pachoud, & J.-M. Roy (Eds.), *Naturalizing phenomenology: Issues in contemporary phenomenology and cognitive science* (p. 305). Stanford, CA: Stanford University Press.

Varela, F., & Shear, J. (1999). First person methodologies: What, why, how? *Journal of Consciousness Studies, 6,* 1–14.

Varela, F., Thompson, E., & Rosch, E. (1995). *The embodied mind: Cognitive science and human experience.* Cambridge, MA: MIT Press.

Von Foerster, H. (2002). *Understanding understanding: Essays on cybernetics and cognition.* New York: Springer-Verlag.

Voss, R., & Clarke, J. (1975). 1/f noise in music and speech. *Nature, 258,* 317–318.

Received February 1, 2007
Accepted February 1, 2008 ■