# Test and Evaluation Challenges of Embodied Artificial Intelligence and Robotics

Technical Report UT-CS-08-628

Bruce J. MacLennan*

Department of Electrical Engineering & Computer Science
University of Tennessee, Knoxville
`www.cs.utk.edu/~mclennan/`

August 22, 2008

**Abstract**

Recent developments is cognitive science, artificial intelligence (AI), and robotics promise a new generation of intelligent agents exhibiting more of the capabilities of naturally intelligent agents. These new approaches are based on neuroscience research and improved understanding of the role of the body in efficient cognition. Although these approaches present many advantages and opportunities, they also raise issues in the testing and evaluation of future AI systems and robots. We discuss the problems and possible solutions.

* This report is an unedited draft of "Challenges of Embodied Artificial Intelligence and Robotics," an article invited for *The ITEA Journal of Test and Evaluation* of The International Test and Evaluation Association. It may be used for any non-profit purpose provided that the source is credited.

# 1 Introduction

Recent research into human and animal cognition has improved our understanding of natural intelligence, and has opened a path forward toward artificially intelligent agents, including robots, with much greater capabilities than those implemented to date. Improved understanding of the neural mechanisms underlying natural intelligence is providing a basis for implementing an efficient, robust, and adaptable artificial intelligence (AI). However, the nature of these mechanisms and the inherent characteristics of an AI based on them, raise significant issues in the test and evaluation of this new generation of artificially intelligent agents. This article briefly discusses the limitations of the "old AI," the means by which the "new AI" aims to transcend them to achieve an AI comparable to natural intelligence, the test and evaluation issues raised by the new AI, and possible means for dealing with these issues in order to deploy robust and reliable systems capable of achieving mission objectives.

# 2 The Nature of Expertise

It used to be supposed that human expertise consists of internalized rules representing both knowledge and inference. Knowledge was considered a collection of (general or specific) facts that could, in principle, be expressed as sentences in a natural language. Similarly, the process of thought was supposed to be represented by rules of inference, such as the laws of logic, also expressible in natural language. It was granted that natural language was vague, ambiguous, and imprecise, and so artificial languages, such as symbolic logic, were proposed as more adequate vehicles for knowledge representation in the brain. *Knowledge representation languages*, which were often used in AI, were effectively programming languages for operating on knowledge represented by language-like data structures.

A full critique of this model of knowledge and cognition is beyond the scope of this article, so I will just mention a few key points (for more, see, e.g., Dreyfus 1979; Dreyfus and Dreyfus 1986). One objection was that neuroscience provides no evidence that the brain is structured like a stored-program computer. In answer it was argued that the abstract idea of a general-purpose computer (i.e., of a universal Turing machine) could be implemented in many radically different ways, and so the brain could be such a machine even though it is very different from ordinary computers, and it was argued that in any case there was no reason for artificial intelligence to slavishly imitate the brain; we could use our technologically superior digital computers. Another objection was that, while we are sometimes conscious of following verbalizable rules, much of our intelligent behavior takes place without conscious rule following. In answer it was argued that well-learned behaviors were "compiled" into unconscious neural operations, much as programs written in high-level languages are compiled into machine code. A third objection was that, while it might be plausible that human knowledge and inference were represented in language-like rules, this was implausible as a model for nonhuman animal cognition, especially in simpler animals with no language-using ability. One answer was that nonhuman animals don't have conceptual knowledge, which is "true knowledge," as opposed to concrete memory and instinctive stimulus-response behaviors; only humans exhibit "true" cognition. An overarching defense of rule-based models of knowledge and inference was that they are "the only game in town," that is, that there were no defensible alternative models. Nevertheless, there are additional objections to rule-based approaches.

Even for humans, who have complex and expressive linguistic abilities, research shows that rules don't account well for expert behavior. As an example, I will use the book of Hubert L. and Stuart E. Dreyfus (1986), who summarize much of the research. Based on characteristic cognitive processes, they identify five levels of expertise: (1) novice, (2) advanced beginner, (3) competence, (4) proficiency, and (5) expertise (Dreyfus and Dreyfus 1986, 16–51). They apply this classification to "expert systems," which are rule-based AI systems incorporating a large knowledge-base, oriented toward some domain of knowledge, and appropriate inference rules. They argue that these systems operate at best at the "competence" level, which is characterized by goal-directed selection and application of rules (Dreyfus and Dreyfus 1986, 23–27, 101–121). However, expert systems cannot perform at the "proficient" level, which is characterized by unconscious, similarity-based apprehension of the situational context in which cognition should occur, rather than by conscious, rational "calculation" (rule-based determination) of the context (Dreyfus and Dreyfus 1986, 27–30). This apprehension of context is critical to proficient behavior, since it allows cognition to focus on stimuli that are relevant to the situation, without wasting time considering and rejecting those that aren't. Experts apply rules, if at all, in a flexible, nonrigid, context-sensitive way, which is why it is difficult to capture expertise in rules (Dreyfus and Dreyfus 1986, 30–36, 105–109). How, then, can we design artificially intelligent agents that exhibit true expertise?

# 3 Connectionism

The rule-based approach to knowledge representation and inference continued to dominate AI so long as there did not seem to be any viable alternative. However H.L. Dreyfus (1979) and others pointed the way to a different approach. First, since human and animal intelligence is realized in the physical brain, it seemed apparent that an artificial intelligence would be possible, although the AI system might have to be more like a brain than a conventional computer. Second, in the 1960s and '70s, Pribram, Dreyfus, and others had observed that human pattern recognition and memory seemed to have properties similar to optical holograms, as did simple models of neural networks (e.g., Anderson, Pellionisz and Rosenfeld 1990, ch. 7; Dreyfus 1979, 20, 25, 51; Dreyfus and Dreyfus 1986, 58–63, 90–92, 109; Haugeland 1978; Hinton and Anderson 1989; Pribram, Nuwer, and Baron 1974). These considerations helped to revitalize, in the early 1980s, the study of neural network computation, which had been languishing for about a decade (for more on the history of neural networks and connectionism, see MacLennan 2001; seminal papers are collected in Anderson and Rosenfeld 1988; Anderson, Pellionisz and Rosenfeld 1990; Haugeland 1997).

*Connectionism* is used to refer to approaches to knowledge representation and inference that are based on simple neural-network models. In rule-based approaches, knowledge is represented in language-like discrete structures, the smallest units of which are *features*: predicates for which many languages have words (e.g., "feathered," "winged," "two-legged," "egg-laying" are some features of birds). Connectionist representations, in contrast, are based on large, unstructured arrays of *microfeatures*. A microfeature is a property localized to one of a large number of parts of a sensory or memory image (e.g., a green pixel at a particular location in an image), which generally does not have any meaning in isolation. They are *not* the sorts of things for which natural languages have words, because they are not normally significant, or even perceptible, in isolation. In a typical neural-net representation, the activity level of a neuron (usually a continuous quantity) represents the relative degree of presence of a corresponding microfeature in the representation (e.g., the amount of green at that location in the image). As a consequence of the foregoing, connectionist representations are typically holistic in that individual elements have meaning only in the context of the whole representation.

Connectionism derives its name from the fact that knowledge is encoded in the connections between neurons. Because these are connections among (typically large) numbers of neurons representing microfeatures, connectionist knowledge representation is characteristically distributed and nonlocal. It is *distributed* in that the representation of what we normally think of as one fact or behavior is distributed across many connections, which affords connectionist knowledge representations a high degree of useful redundancy. It is *nonlocal* in that each connection participates in the representation of a large number of facts and behaviors. Therefore, a large number of connections in a neural network can represent a large number of facts and behaviors, but not in a one-to-one manner. Rather, the entirety of the connections represents the entirety of the facts and behaviors.

Biological neurons are notoriously slow compared to contemporary electronics; their maximum impulse rate is less than 1 KHz. And yet brains, even of comparatively simple animals, solve problems and coordinate activities that are beyond the capabilities of state-of-the-art computers, such as reliable face recognition and locomotion through rough and complex natural environments. How is this possible? Part of the answer is revealed by the "100-Step Rule" (Feldman and Ballard 1982). This is based on the simple observation that if we take the time for a simple cognitive action, such as recognizing a face ($\ll$ 1 sec.) and divide it by the time it takes a neuron to fire ($\gg$ 1 msec.), we find that there can be at most about 100 sequential processing steps between sensation and action. This reveals that brains process information very differently from contemporary computers. Information processing on traditional computers is *narrow-but-deep*, that is, it depends on the sequential execution of very large numbers of very rapid operations; even if execution is not completely sequential, the degree of parallelism is very small compared to a brain's. In contrast, information processing in brains is *shallow-but-wide*: there are relatively few sequential layers of information processing, as reflected in the 100-Step Rule, but each layer is massively parallel on a scale that is qualitatively different from contemporary parallel computers. For example, even in the retina approximately 100 million retinal cells preprocess visual data in order to be transmitted by approximately one million optic nerve fibers, which indicates the degree of parallel processing in visual information processing. Since neural density is at least 146 000/mm$^2$ throughout human cortex (Changeux 1985, 51), most neural modules operate with degrees of parallelism on the order of hundreds of thousands or millions.

Another difference between most contemporary computers and biological neural networks is that neurons are fundamentally *analog* computing devices. Continuous quantities are represented by the frequency, and in some cases the phase, of neural impulses propagating down a neuron's axon (output fiber). Knowledge is stored in "strength" of the connections between neurons, which depends on diffusion of chemical signals from a variable number of sources to a variable number of receptors, and is best treated as a real-valued "weight." The resulting electrical signals propagate continuously down the dendrites (input fibers) of a neuron, obeying electrical "cable equations" (Anderson 1995, 25–31), and are integrated in the cell body into a continuous membrane potential, which governs the frequency and phase of the neuron's spiking behavior (Gerstner and Kistler 2002).

It should be noted that analog signal processing in the brain is low-precision: generally continuous quantities are estimated to be represented with a precision of less one digit. Paradoxically, humans and other animals can perform perceptual discriminations and coordinate sensorimotor behaviors with great precision, but brains use statistical representations, such as "coarse coding" and other population codes (Rumelhart, McClelland, et al. 1986, 91–96; Sanger 1996), to achieve high-precision representations with low-precision components. These techniques, which exploit large numbers of neurons, have additional benefits in terms of reliability, robustness, and redundancy.

Similarly, artificial neural networks are usually based on analog computing elements (artificial neurons or *units*) interconnected by real-valued weights. Of course these continuous computational systems, like other continuous physical systems, can be simulated on ordinary digital computers, and that is the way many artificial neural networks are implemented. However, many advantages can be obtained by implementing artificial neural networks directly in massively-parallel, low-precision analog computing devices (Mead 1989), a topic outside the scope of this article (MacLennan in press).

The ability to adapt to changing circumstances and to learn from experience are hallmarks of intelligence. Further, learning and adaptation are critical to many important applications of robotics and artificial intelligence. Autonomous robots, by their very autonomy, may find themselves confronting situations for which they were not prepared, and they will be more effective if they can adapt appropriately to them. Autonomous robots should also be able to adapt as the parameters and circumstances of their missions evolve. It is also valuable if AI systems and robots can be trained in the field to perform new tasks and if they can generalize previous training to new, unanticipated situations. How can learning, training, and adaptation be accomplished?

An important capability of connectionist AI systems is that they can learn how to do things that we do not know how to do. This is the reason that connectionist systems are said to be *trained*, but not *programmed*. In order to program a process, you need to understand it so well that it can be reduced to explicit rules (an algorithm). Unfortunately, there are many important problems that are not sufficiently well understood to be programmed, and in these cases connectionist learning may offer an alternative solution. Many connectionist (or neural network) learning algorithms have been developed and studied over the last several decades. In *supervised learning*, a network is presented with desired input-output pairs (e.g., digital images and their correct classifications), and the learning algorithm adjusts the network's interconnection weights so it will produce the correct outputs. If the training is done properly the network will be able to generalize from the training inputs to novel inputs. In *reinforcement learning*, the network is told only whether it has performed correctly or not; it is not told the correct behavior. There is a very large literature on neural network learning, which is beyond the scope of this article (see, e.g., Haykin 1999).

One characteristic of connectionist learning is that, while connectionist systems can sometimes adapt very quickly, they can also adapt gradually, by subtle tuning of the interconnection weights. Rule-based systems can also adapt, but the fundamental process is the addition or deletion of a complete rule, a more brittle procedure. Thus connectionist systems are better able to modulate their behavior as they adapt and to avoid instability.

# 4 Embodied Cognition

An important recent development is the theory of *embodied cognition* and the related theories of *embodied AI* and *embodied robotics*. The theory of embodied cognition addresses the important — indeed essential — role that the body and its physical environment plays in efficient cognition. As Dreyfus (1979, 248–250, 253) observed long ago (1972), there are many things that humans know simply by virtue of having a body. That is, there is much knowledge that is implicit in the body's state, processes, and relation to its physical environment, and therefore this knowledge does not need to be represented explicitly in the brain. The theory of embodied intelligence has its roots in phenomenological philosophy (e.g., Dreyfus 1979, 235–255) and the pragmatism of William James and John Dewey (Johnson and Rohrer 2007).

For example, we swing our arms while we walk, which helps maintain balance for bipedal locomotion, but our brains do not have to calculate the detailed kinematics of our limbs. Rather, our limbs, joints, etc. have their characteristic frequencies etc., and all our brain must do is generate relatively low-dimensional signals to modulate these physical processes to maintain balance, as monitored by bodily sensors (inner ear, skin pressure, joint extension, etc.). The brain's goal is not to *simulate* the physical body in motion (a computationally intensive task), but to *control* the physical body in interaction with its physical environment in real time by means of neurally efficient computation. As opposed to a computer simulation of a robot, the brain's computations constitute a complete description of the body's motion only in the context of a specific physical body in its environment.

Because, in effect, an animal's brain can depend on the fact that it is controlling a body of a specific form, it can offload some information processing tasks to its physical body. For example, rather than calculating from first principles the muscle forces that will move its limb to a particular location, it can leave this "calculation" to the physical limb itself by learning correlations between effector signals and corresponding sensory responses (for which neural networks are ideally suited). Therefore also, if a weight (such as a cast) is put on a limb, or its motion is restricted by pain or an injury, an animal can adapt quickly to the change (an important goal for our robots too).

The power and efficiency of embodied cognition is exemplified by insects and other simple animals that behave very competently in their environments but that have very small brains. Understanding how they exploit embodiment for information processing — or, more precisely, how they obviate the need for information processing — will help us to design more competent autonomous robots, especially insect-size or smaller robots.

Studies of natural intelligence, and in particular of how the brain exploits the physical characteristics of the body and of its environment to control the body in its environment, has contributed to and will continue to contribute to the design of future robots (Brooks

1991; Pfeifer and Bongard 2007; Pfeifer, Lungarella, and Iida 2007). We are inclined to think of these problems in terms of equations and calculations (i.e., rule-like information representation and processing), but natural intelligence teaches how to use neural networks for efficient and competent behavior in real-world environments. This is a critical goal for future autonomous robots and indeed for artificial intelligence embedded in other physical systems.

# 5 Challenges

We have argued that connectionist artificial intelligence, based on neural network models and embodied cognition, provides a sounder, more effective basis for future AI and robotic systems than does rule-based knowledge representation and processing. Indeed there is widespread (though not universal) agreement on this, and many projects are pursuing these approaches. Therefore it is important to acknowledge that connectionism and embodiment present challenges for the test and evaluation of the systems in which they are used.

One problem is the *opacity* of neural networks. In a rule-based system the rules are expressed in an artificial language with some similarity to natural languages or to symbolic logic. The basic terms and predicates, in terms of which the rules are expressed, are generally those of the problem domain. Therefore the knowledge and rules of inference used by the system are *transparent*, that is, potentially intelligible to human beings. In a neural network, in contrast, the knowledge and inferential processes are implicit in real-valued connection weights among myriads of microfeatures. Further, representations are nonlocal and distributed. Therefore, individual microfeatures and connections do not usually have meanings that can be expressed in the terms of the problem domain.

Many people are troubled by the opacity of neural networks compared to the (potential) transparency of rule-based systems. With a rule-based system, they argue, you can look at the knowledge base and inferential rules, understand them, and see if they make sense. A human can, in effect, see if the system is making its decisions for the right reasons, or at least that it is not making them for the wrong reasons (e.g., on the basis of irrelevant factors). In contrast, a trained neural network might perform some task very well, but we will be unable to understand — in human-comprehensible terms — the bases on which it is doing its job. Perhaps it has found some totally irrelevant cues in the training and test data that allow it to perform well on them, but it will fail dismally when deployed.

These are legitimate concerns, but unavoidable. As we have seen, rule-following is characteristic of merely "competent" behavior, and therefore behavior that *can* be expressed in human-comprehensible rules will not surpass the merely competent level. Conversely, expert behavior — which is our goal for AI and autonomous robotics — will entail subtle discriminations, integrative perceptions, and context sensitivities that cannot be expressed in human-comprehensible terms. How then can we come to trust a connectionist AI system? In the same way we come to trust a human expert: by observing their reliably expert behavior in wide variety of contexts and situations. The situation is similar to that with the use of unsupervised trained animals to perform some mission. We cannot look into *their* heads either, but we can test their behavior in a variety of mission-relevant situations until we have sufficient confidence.

Much of the inflexibility and brittleness of rule-based systems — and indeed of many digital computer programs — is a consequence of their behaving the same in all contexts, whereas natural intelligence is sensitive to context and can modulate its behavior appropriately. Due to the ability of artificial neural networks to integrate a very large number of microfeatures, which may be individually insignificant, they can exhibit valuable context sensitivity. However, this presents a test and evaluation challenge for connectionist systems, since we cannot test such a system in a single or simple context (e.g., in a laboratory) and assume that it will work in all contexts. Rather, it is important to identify the contexts in which the system may find itself and ensure that it operates acceptably in all of them.

Context sensitivity and embodied cognition both necessitate use of the *implemented* robotic or AI system in almost all phases of test and evaluation. As previously mentioned, one of the advantages of connectionist AI is that it can be sensitive to the context of its behavior, but this implies an early transition of the test and evaluation activity into realistic physical contexts (i.e., field testing). Since we want and expect the system to make subtle contextual discriminations, it cannot be adequately tested or evaluated in artificially constructed situations that do not demand this subtlety. The same applies to the system's (hopefully robust) response to novelty. Further, embodied intelligence depends crucially on the physical characteristics of the system in which it is embedded and on its physical relationships to its environment. While preliminary testing and evaluation can make use of simulations of the physical system and its environment, such simulations are always incomplete, and are more computationally expensive the more complete they are. Whereas to some extent conventional AI systems can be tested and evaluated offline, embodied AI systems cannot. Therefore physical prototypes must be integrated earlier into the development cycle.

In effect, test and evaluation of embodied connectionist AI and robotic systems is no different from that of vehicles, weapons systems, and other physical devices and equipment. The difference is in our expectations, for we are used to being able to test and evaluate software systems offline, except in the later stages in the case of embedded software.

Finally, as discussed above, embodied connectionist systems are to some degree opaque, that is, their *cognitive* processes are not fully transparent (intelligible) to humans. Of course, neural networks and their embodiments obey the laws of physics, and are intelligible in *physical* terms, but that level of explanation is of limited use in understanding the intelligent behavior of a system. This seems like a distinct disadvantage compared to abstract, rule-based systems but, as we have argued, it is a necessary consequence of expert behavior. In this regard, the test and evaluation of embodied connectionist systems is not much different from that of other physical systems, for which abstract models and simulations are insufficient in the absence of field testing.

Further, the deployment of embodied connectionist systems is not qualitatively different from the deployment of trained animals or humans. Being able to recite memorized rules of procedure or to perform well in laboratory environments does not substitute for performance testing and evaluation in real, or at least realistic, situations.

# 6 Conclusions

We have argued that embodied connectionist AI and robotics promises a new generation of intelligent agents able to behave with fluent expertise in natural operational environments. Such systems will be able to modulate their perception and behavior according to context and to respond flexibly and appropriately to novelty, unpredictability, and uncertainty in their environments. These capabilities will be achieved by understanding natural intelligence, its realization in neural networks, and its exploitation of embodiment, and by applying this understanding to the design of autonomous robots and other intelligent agents.

However, a more natural intelligence is also an intelligence that responds more subtly to its environment, and guides its body in a fluent dance with its physical environment. As a consequence, such systems cannot be adequately tested or evaluated independently of their physical embodiment and the physical environment in which they act. Naturally intelligent systems typically lack both transparency of behavior and independence of information processing from physical realization, which we have come to expect in artificial intelligence.

Nevertheless, such systems may be tested and evaluated by similar approaches to those applied to other inherently physical systems; it is really only a shift of emphasis from abstract rules and programs to concrete physical interaction with the operational environment.

□

*Bruce MacLennan received a Bachelor of Science in mathematics from Florida State University in 1972, a Master of Science in computer science from Purdue University in 1974, and a Ph.D. in computer science from Purdue in 1975. From 1976 to 1979 he was a Senior Software Engineer at Intel and contributed to the architecture of the 8086 and iAPX-432 microprocessors. In 1979 he joined the computer science faculty of the Naval Postgraduate School in Monterey, CA, where he was Assistant Professor, Associate Professor, and Acting Chairmen. Since 1987 he has been a member of the computer science and electrical engineering faculty at the University of Tennessee, Knoxville, where he investigates novel computing technologies, nanotechnology, and biologically-inspired computation.*

# 7 References


Anderson, J.A. 1995. *An introduction to neural networks*. Cambridge: MIT Press.

Anderson, J.A., Pellionisz, A., and Rosenfeld, E. (Eds.). 1990. *Neurocomputing 2: Directions for research*. Cambridge: MIT Press.

Anderson, J.A., and Rosenfeld, E. (Eds.). 1988. *Neurocomputing: Foundations of research*. Cambridge: MIT Press.

Anderson, M.L. 2003. "Embodied cognition: A field guide." *Artificial Intelligence* 149: 91–130.

Brooks, R. 1991. "Intelligence without representation." *Artificial Intelligence* 47: 139–159.

Changeux, J.-P. 1985. *Neuronal man: The biology of mind,* tr. by L. Garey. Oxford: Oxford University Press.

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

Dreyfus, H.L. 1979. *What computer's can't do: The limits of artificial intelligence*, rev. ed. New York: Harper & Row.

Dreyfus, H.L., and Dreyfus, S.E. 1986. *Mind over machine: The power of human intuition and expertise in the era of the computer*. New York: Free Press.

Feldman, J.A., and Ballard, D.H. 1982. "Connectionist models and their properties." *Cognitive Science* 6(3): 205–254.

Gerstner, W., and Kistler, W.M. 2002. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge: Cambridge University Press.

Haugeland, J. 1978. "The nature and plausibility of cognitivism." *Behavioral and Brain Sciences* 1: 215–226.

Haugeland, J. (Ed.) 1997. *Mind design II: Philosophy, psychology, artificial intelligence*, rev. & enlarged ed. Cambridge: MIT Press.

Haykin, S. 1999. *Neural networks: A comprehensive foundation*, 2nd ed. Upper Saddle River: Prentice-Hall.

Hinton, G.E., and Anderson, J.A. (Eds.). 1989. *Parallel models of associative memory*, updated ed. Hillsdale: Lawrence Erlbaum.

Johnson, M., and Rohrer, T. 2007. "We are live creatures: Embodiment, American pragmatism, and the cognitive organism." In J. Zlatev, T. Ziemke, R. Frank,and R. Dirven (Eds.), *Body, language, and mind*, vol. 1, pp. 17–54. Berlin: Mouton de Gruyter.

MacLennan, B.J. 2001. “Connectionist approaches.” In N.J. Smelser and P.B. Baltes (Eds.), *International encyclopedia of the social and behavioral sciences*, pp. 2568–2573. Oxford: Elsevier.

MacLennan, B.J. in press. “Analog computation.” In R.A. Meyers et al. (Eds.), *Encyclopedia of complexity and system science*. New York: Springer.

Mead, C. 1989. *Analog VLSI and neural systems*. Reading: Addison-Wesley.

Pfeifer, R., and Bongard, J.C. 2007. *How the body shapes the way we think — A new view of intelligence*. Cambridge: MIT.

Pfeifer, R., Lungarella, M., and Iida, F. 2007. “Self-organization, embodiment, and biologically inspired robotics.” *Science* 318: 1088–1093.

Pribram, K.H., Nuwer, M., and Baron, R.J. (1974). “The holographic hypothesis of memory structure in brain function and perception.” In D.H. Krantz, R.C. Atkinson, R.D. Luce, and P. Suppes (Eds), *Contemporary developments in mathematical psychology*, vol. 2, pp. 416–457. New York: Freeman.

Rumelhart, D.E., McClelland, J.L., and the PDP Research Group. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 1: Foundations*. Cambridge: MIT Press.

Sanger, T.D. 1996. “Probability density estimation for the interpretation of neural population codes.” *Journal of Neurophysiology* 76: 2790–2793.