Cultural Affordances in AI Perception

Zachariah A. Neemeh (zaneemeh@memphis.edu)
Department of Philosophy, University of Memphis, Memphis, TN 38152

Abstract

Affordances offer AI research an alternative from representations for linking perception to action in autonomous systems. Affordances are based in the informational structure of the environment and the somatic capacities of the agent and arise in their interaction. AI implementations of affordance perception typically utilize relatively basic, natural affordances such as the graspability of a handle. Culturally-scaffolded affordances, such as the letter-mailing capacity of a postbox, pose a more intractable problem for affordance-based robotics. This class of affordances requires acculturation and is highly culture-specific. AI implementations of affordance perception typically bypass this difficulty by making recourse to representations. I begin by reviewing affordance perception and the difference between natural and cultural affordances. I then critically discuss implementations of cultural affordance perception in autonomous agents. Finally, I argue that AI affordance perception does not require a robust representationalism in order to implement cultural affordances.

Keywords: affordances; AI perception; embodied cognition; philosophy of AI; representations

Introduction

The perpetually shifting nature of the environment poses a significant challenge to robotics research. Autonomous agents must negotiate and adapt to dynamically changing environments. Representational architectures limit autonomous agents' capacity to do so, however (Raubal & Moritz, 2008). Problems arise with the bandwidth, processing power, computational time, and programming time required to represent shifting environments (Rome et al., 2006). Nonrepresentational and affordance-based architectures have been proposed to overcome these difficulties (Braitenberg, 1984; Brooks, 1990, 1991; Horton, Chakraborty, & St. Amant, 2012). These architectures do not rely on separate layers for perception, action, and planning or reasoning. Instead, they offload part of the computational process onto the environment. As Rodney Brooks, the pioneer of embodied robotics and the inventor of the Roomba, said, "the world is its own best model" (1990, p. 4).

Many nonrepresentational architectures utilize affordances to replace otherwise separate perceptual and actional layers. Natural affordances, like the graspability of a handle, are embedded in the basic informational structure of the immediate environment (Gibson, 1979/2015). The agent picks up on information available in the environment, such as the light waves and pressure feedback of the handle. The affordance arises as the agent interacts with the object and is an opportunity for action that is highly constrained by the agent's form of embodiment.

Implementations of affordance perception are beset by a difficulty, however, once they encounter cultural affordances, or affordances that implicate background knowledge that is culturally mediated. Few AI implementations of affordance perception have attempted to incorporate such higher-order affordances (see Awaad, Kraetzschmar, and Hertzberg, 2015; Chu, Fitzgerald, and Thomaz, 2016; Raubal & Moritz, 2008). While the graspability of a handle can be modeled as an online, dynamical interaction unfolding between the agent's sensorimotor processes and the object's properties, the mail-ability of a letter implicates a vast background knowledge of letters, the postal service, postboxes, writing, and interpersonal communication. This background knowledge poses a significant problem. If all the rules and background knowledge pertinent to the mail-ability affordance must be represented, then affordance-based robotics offers little advantage over traditional architectures.

In this paper, I critically review extant AI implementations of cultural affordance perception and sketch a framework for perceiving cultural affordances with minimal recourse to representations. My aim is to show that a robust representationalism is not *conceptually* necessary for cultural affordance perception.

Affordance Perception

In classical computational models, the perception of the environment involves the creation of an inner, representational model (Fodor, 1985; Marr, 1982/2010).[1] The agent is decoupled from its environment and interacts with it through the medium of representations, which are processed computationally. The term 'representation' is used in a wide variety of senses, from a minimal sense of a covariation between an internal and external state, to a more robust internal mapping of an external state. To avoid the deep complexities involved in this term, I here use it to pick out any internal state that tracks an external state in the world, when that state is decoupled from perception and action. In autonomous agents, typically representations are instantiated in a planning or reasoning layer mediating between perceptual and actional layers. Processing can be

[1] While computationalism and representationalism do not necessarily entail one another (see Dennett, 1969), in practice they usually work in tandem.

through classical, serial processing architectures (as in Turing machines), or they can be massively parallel (as in connectionist, neural network, and similar architectures).

Autonomous agents in the real world encounter a wide variety of environments, no two of which will exactly be the same. Even the same environment often shifts in content through time. This creates a significant computational challenge, in addition to being resource- and energy-intensive. The existence of an inner representational layer places all the computational burden on the agent itself. Affordances, however, arise in the agent-environment interaction, offloading part of the processing burden onto the environment. Introducing affordance perception into autonomous agents enables them to continuously and dynamically adapt to shifting and changing environments.

Traditional autonomous agents separated sensing, planning/reasoning, and acting into different processes that would only link up at a later stage (Gat, 1998; Maes, 1991). The perceptual process sends information to the planning process, which in turn sends instructions for action (Horton, Chakraborty, & Amant, 2012). However, "[e]ven if an agent has perfect segmentation and feature recognition capabilities, this new form of information may be hard to translate into appropriate actions" (Nye & Silverman, 2012, p. 184). Affordance perception dispenses with the intermediary planning layer, instead generating affordances within a tight perception-action loop. What planning there may be is performed online through the perception-action loop, instead of offline between perception and action. This does not mean that the agent does no planning whatsoever; rather, it means there is often no representational layer mediating between perceptual and actional processes—at least, not at the level of basic perceptual processes (see Şahin et al., 2007). What planning there may be is performed online through the perception-action loop, instead of offline between perception and action. Furthermore, machine learning alone is insufficient; a robotic body is *required* for an affordance to be perceived. This is because affordances are not merely perceptual processes—they are perception-action processes and require dynamic engagement with the environment.

Natural and Cultural Affordances

Affordance perception in AI is complicated by the fact there are two very different types of affordances: natural and cultural. Natural affordances involve very basic cognitive processes. Cultural affordances are comparatively richer and involve culturally- and intersubjectively-mediated processes in order to be perceived and acted upon.[2] Cultural affordances, however, pose a particularly intractable problem. While natural affordances arise from the informational structure of the environment, cultural affordances require that the percipient be acculturated. There seems, *prima facie*, to be a level of decoupled, even representational, processing required to perceive a cultural affordance (Ramstead, Veissière, & Kirmayer, 2016).

Natural affordances are possibilities in the environment available for action (Dotov, Nie, & de Wit, 2012). Different agents can perceive different affordances based on their embodied capacities and species-typical behaviors. For example, a twig affords different actions to a cat, a finch, and a human. To the cat, the twig affords bite-ability and play-ability. To the finch, it affords graspability by the beak and build-ability for a nest. Finally, to the human, it affords manual manipulation. In each case, the embodied capacities and species-typical behaviors of the agent shape what kind of action the twig affords.

Affordances are based on the real information (light, pressure, scent) available in the environment. However, they do not themselves exist in the environment. They are generated in the agent-environment interaction. Affordance-perception occurs because the agent and environment form a complex, emergent system (Favela & Chemero, 2016; Thompson, Varela, & Rosch, 1991/2016; Gallagher, 2017; Thompson, 2007). That is, the agent is dynamically coupled with the environment. This coupling is modeled in ecological psychology and embodied cognition research using dynamical systems theory (Beer, 2014; Chemero, 2009; Turvey, 2019).

Several formalizations of affordances have been proposed (see Chemero, 2003; Stoffregen, 2003; Turvey, 1992). Stoffregen's (2003) formalization, which has been successfully utilized in AI affordance perception research (Nye & Silverman, 2012), is:

> "Let W_{pq} (e.g., a person-climbing-stairs system) = (X_p, Z_q) be composed of different things Z (e.g., person) and X (e.g., stairs). Let p be a property of X and q be a property of Z. The relation between p and q, p/q, defines a higher order property (i.e., a property of the animal−environment system), h. Then h is said to be an affordance of W_{pq} if and only if
>
> - W_{pq} = (X_p, Z_q) possesses h.
> - Neither Z nor X possesses h" (Stoffregen, 2003, p. 123).

Cultural affordances require the agent to utilize "explicit or implicit expectations, norms, conventions, and cooperative social practices" (Ramstead, Veissière, & Kirmayer, 2016, p. 3). It is precisely these elements that seem, *prima facie*, to require a representational layer decoupled from perception-action processes. For example, Gibson (1979/2015) remarks that a buyer and a seller each afford one another opportunities for action (viz., buying and selling). However, he goes on to say,

> "The perceiving of these mutual affordances is enormously complex, but it is nonetheless lawful, and it is based on the pickup of the information in touch, sound, odor, taste, and ambient light" (p. 127).

The information, in this case, is directly out there in the environment, and the agent perceives it. The affordances for

[2] Although affordances may differ regarding their basicness or their cultural scaffolding, in practice it is difficult to disentangle these two (see Wagman, Caputo, & Stoffregen, 2016). Indeed, for human agents, even basic perception-action processes like picking up an apple are culturally mediated.

action, however, arise in the interaction of the agent with its environment. It is the information, not the affordance, that is objectively embedded in the immediate environment. However, how could a buyer be perceived *as such* merely based on light waves, sound pressure waves, and other ecological information?

Gibson also claims that "the real postbox…affords letter-mailing to a letter-writing human in a community with a postal system" (1979/2015, p. 130). In this example, we have a culturally-scaffolded process of perception and action that functions only within a highly-specific cultural framework. It is not clear, however, how these culturally-scaffolded processes could be "directly" perceived based on the immediate information available in the environment. Memory and background knowledge are required for the postbox to be perceived with a letter-mailing affordance. However, there is little in the postbox's shape, color, and size that informs the agent of the postal system, letter-writing culture, and letter-reading agents enabling it to have mail-ability. Either cultural affordances are representational (see Ramstead, Veissière, & Kirmayer, 2016), or they must somehow be generated in a cultural milieu and for an acculturated agent by utilizing nonrepresentational, memory-based processes (see Rietveld & Kiverstein, 2014).

AI Cultural Affordance Perception

Most AI implementations of affordance perception have focused on natural affordances. These are, no doubt, relatively easier to implement because they do not require background knowledge of culture or a process of enculturation in order to perceive and act upon them. They are based only on the informational structure of the immediate environment. The true challenge for AI affordance perception is to achieve the perception of *cultural* affordances. If, however, cultural affordance perception turns out to require a robust representationalism, it is not clear that it has any advantage over non-affordance-based AI.

Raubal and Moratz (2008) provide an AI implementation of cultural affordance perception whereby cultural affordances are scaffolded onto natural affordances by representations of cultural knowledge. Their target agent is the Bremen Autonomous Wheelchair *Rolland*, which interprets linguistic commands by its human occupant and navigates across the environment. The need for cultural affordance perception arises because the wheelchair does not blindly perform actions commanded by their users. For example, the user may request to visit a center outside of operational hours. In this case, the AI utilizes cultural affordances integrating knowledge of the institution and its operating hours when selecting for action outputs.

Cultural affordances arise in their system by a system of constraints upon natural affordances. A natural affordance is constrained within a given social and institutional context. For example, the mailbox affords a multiplicity of actions, including smashing, opening, inserting objects, and touching. In their model, it is the social and institutional context of the postal system, letter-readers, and letter-writing that constrain the possible natural affordances into a smaller subset of cultural affordances. The agent then performs internal actions on these cultural affordances—essentially, planning or reasoning processes—in order to act upon the more basic natural affordance of opening and inserting.

The cultural affordances utilized by Raubal and Moratz' (2008) agent are representational. A separate planning layer is retained by their AI wheelchair. Their conception of cultural affordances is simply a subset of natural affordances that are given social and institutional constraints. Knowledge such as closing and opening hours is certainly representational and linguistically-based. The problem with their implementation of cultural affordances is that there is little that distinguishes them from classical representations. The construct of 'cultural affordance' is not doing any work that the construct of 'representation' does not already do. Their agent is essentially a hybrid system incorporating affordance perception for low-level navigation and symbolic representations for higher-level constraints upon that navigation.

Furthermore, some forms of social and institutional knowledge that Raubal and Moratz (2008) discuss, such as navigating across a city, are not necessarily fully representational processes. Unwritten norms such as walking on the right side of the sidewalk in many Western countries could be conceived of as representational rules. However, spontaneous pedestrian patterns can emerge without any specific intention (Moussaid et al., 2009).

Socialization and Supervised Learning

Awaad, Kraetzschmar, and Hertzberg (2015) provide an affordance-based model for AI agents that can "socialize" by learning expected uses of objects. The practical applications of this are in producing service robots that perform actions commanded by humans without being "robotic." When humans perform service tasks, an entire body of knowledge is brought to bear. Take the example of sweeping the floor. The human agent needs to know how to use a broom. However, the possibility space for utilizing a broom to sweep in *deviant* ways is quite large: one could sweep under the feet of others, sweep at the wrong times (e.g., while others are cooking), or sweep with furious movements and kick up dust. All these behaviors accomplish the task of sweeping but are social nuisances and perhaps even physically dangerous. There is an entire network of social expectations and etiquette surrounding the tool use in question. There is, in short, a "right way to do things." Furthermore, humans

> "effortlessly adapt our actions to unexpected situations, especially given the dynamic nature of our environment and the amount of uncertainty about it" (p. 422).

While moving a broom back and forth can be largely explained with natural affordances, these cultural constraints cannot. The broom affords more actions than are socially acceptable or considered appropriate to the task. Awaad, Kraetzschmar, and Hertzberg (2015) attempt to integrate them within an affordance-perception paradigm, however. They note that programming procedural knowledge is insufficient to cover these cases of "the right way to do things" because the agent will always encounter novel situations. They implement Hierarchical Task Network to decompose tasks into a set of individual tasks in order to accomplish a goal.

In order to reduce the possibility space for action to one for socially-appropriate action, Awaad, Kraetzschmar, and Hertzberg (2015) store information about the object, the commanding agent, and the intended uses of the object. These constraints are scaffolded on the natural affordance of the object. The broom, for example, is defined by its socially-intended purpose of cleaning. The commanding agent, the human who demands cleaning, would have a set of preferences and expectations as to how that task is accomplished. The authors implement this cultural scaffolding through coded representations. Like Raubal and Moratz (2008), they conceive of cultural affordances simply as subsets of natural affordances that arise through representational cultural constraints.

Although the authors use representations to implement the socially-scaffolded constraints on the object's affordances, their broader proposal shows how a nonrepresentational framework could be used to do the same work. While they programmed the constraint knowledge into their agents, they suggest that this would be better done by supervised learning, particularly learning by demonstration. In the following section, I argue such supervised learning by demonstration of affordances does not require a strong concept of representations for its implementation.

Joint Interaction and Cultural Affordances: Unsupervised and Supervised Learning

Chu, Fitzgerald, and Thomaz (2016) develop autonomous agents that learn to perceive and use affordances through a combination of unsupervised and supervised learning through interaction with a human. A human teacher physically guides the robot to certain affordances. For example, a robot is taught that drawers have an openable affordance by guiding its hand. The robot learns to mimic this movement and perceive the openability affordance of the drawer's handle.

While the openability of the drawer *prima facie* appears to be a natural affordance provided by the structure of the robot's hand and the drawer's handle, there is a large possibility space for socially-deviant drawer-opening behavior. Although Chu, Fitzgerald, and Thomaz (2016) do not note this, the human teacher is not merely teaching the autonomous agent how to perceive and act upon the openability affordance of the drawer. They are simultaneously teaching the AI agent the *acceptable* way to perform this action. The drawer is not to be forcefully opened or rapidly opened and closed in succession (as a small child may annoyingly do), for example. The process of supervised learning allows the AI agent to learn the socially-acceptable affordances. This makes the drawer's openability affordance not simply natural, but also culturally-scaffolded.

Ramstead, Veissière, and Kirmayer (2016) invoke Gricean norms to understand these contexts. Grice (1975) articulated a set of rules governing conversation. These rules are ancillary to the communicative and phatic functions of language and facilitate nondeviant interactions. For example, one ought to convey as much detail as the topic requires without divulging too much detail. If one fails to do the former, one is perceived as terse, reticent, or uncommunicative. If one fails to do the latter, one is perceived as a windbag. In either case, deviation from the unwritten norm has the effect of interrupting the communicative act itself. Likewise, mundane actions have ancillary but unwritten norms guiding how they ought to be performed. These norms can only be learned by actual practice and observation of these actions in a social context. They are not symbolic rules because there may be no explicit representation of their content. They are merely habitual patterns of behavior used to accomplish certain tasks—e.g., opening a drawer slowly rather than forcefully.

While the robot may not develop shared intentions with the human teacher, in this case, it is significant that the robot only learns to perceive and act upon affordances through a process of interaction with a human (who is a "native" affordance perceiver-actor). In this case, representations are not necessary to explain how the AI agent learns to perceive and act upon the drawer's openability affordance in a socially-nondeviant way. While Awaad, Kraetzschmar, and Hertzberg (2015) programmed in cultural knowledge through representations, this supervised learning process does not specifically require the agent to store representations of cultural constraints, expectations, and other social rules. Rather than storing social rules and using them to constrain the agent's affordance perception and action, supervised learning allows the agent to learn to perceive and act upon the affordance in certain typical ways. Instead of inducing a rule based on the multiple supervised learning instances of opening the drawer—e.g., if drawer, then constraint x, y, z—the agent can simply follow the typical range of paths that have been learned.

One objection is that human agents are conscious of not deviating from socially-accepted norms of tool usage. These norms may be at a higher level than "not kicking up dust." One may be aware that one ought not to bother or annoy anyone. Nonetheless, even that does not require a specific rule. Even if the human agent has such a rule in mind, it is generally not the cause of their socially-nondeviant behavior. We do not walk around constantly thinking "I ought not to annoy x." If we can formulate such a rule, and even implement it in some cases, it is the exception (perhaps applying to a highly novel situation) rather than the norm. There is nothing here that cannot also be explained through processes of social learning, acculturation, and operant conditioning. These parallel the supervised learning trials in

AI affordance perception (Awaad, Kraetzschmar, & Hertzberg, 2015; Chu, Fitzgerald, & Thomaz, 2016).

AI Perception of Cultural Affordances without Representations: Learning and Habit

Representations such as rules can be used to constrain behavior in highly novel situations. Indeed, these kinds of rules may be part of the learning process itself. However, programming a database of representational cultural constraints for autonomous agents is a task just as formidable as that of traditional, non-affordance-based AI and computer vision. It is not clear that utilizing affordances in AI perception gains us anything. The problem, however, is that implementations of cultural affordance perception have generally been representational. The supervised learning in Chu, Fitzgerald, and Thomaz (2016) provides a way of thinking about what partially-nonrepresentational cultural affordances may look like in autonomous agents. Their autonomous agent learned how to open cabinet drawers in nondeviant and socially-acceptable ways. The drawer's handle information could afford multiple possibilities for action that are deviant, such as forcefully opening or rapidly opening and closing. During its supervised learning trials, the autonomous agent only learns the socially-acceptable way of opening the drawer. The agent does not perceive a natural affordance of open-ability. It perceives a cultural affordance of gentle-open-ability, one that is only salient within a given social structure and context.

Their autonomous agent does not have to learn or be preprogrammed with a *rule* about acceptable ways to open drawers. It is through multiple supervised learning trials that the cultural affordance begins to emerge—it is, in short, a *habit*. By habit, I mean a pattern of behavior that develops through supervised and unsupervised learning. The agent's habitual patterns of behavior are not representational in the sense that they are not primarily guided by symbolic rules, although the latter may constrain habits in actual behavior. Habit emerges from a set of previous behaviors and continues to guide future ones without necessarily having any explicit formulation. Surely some affordances must be constrained by symbolic representations. The closing time of a building or institution is something that could be learned by habit. The agent could develop a sense of when it closes by a long process of trial and error. However, that would be far less efficient than simply having a rule representing its closing time. In many cases, though, the work being done by representations can just as well be done by supervised and unsupervised learning or habit.

Returning to the example of sweeping, when the AI agent learns how to sweep from a human teacher, the latter will only teach the socially-accepted ways to sweep. The teacher will not teach how to sweep under people's feet, around them while walking, vigorously so as to kick up dust, or any other socially-deviant manner. The agent would learn these patterns of use of the object. Inducting a specific representational rule to cover these cases is supernumerary and fails to add explanatory value. The agent does not need a representational rule ("do not kick up dust"), because they have been taught to use the broom in a set of patterns that do not include kicking up dust. Following Ockham's razor, if the explanation can be had without representations, then we ought to dispense with them as an *explanans* in those cases.

One of the challenges for robustly-representationalist implementations of cultural affordance perception is that they require just as thorough programming with rules as traditional representational architectures. Habit, however, could greatly reduce the set of background knowledge that needs to be programmed. This is also a primary way that human agents develop habits during development. Children do not learn about their culture's interpersonal distance—the typical distance people stand from one another during communication—by learning a rule about how many centimeters away from another person to stand. They merely develop a habit of standing a certain distance away from another person. This habit is reinforced by observation of others and by violations of the norm (e.g., standing too close to someone can be perceived as aggressive). They may not even be aware that there is such a social norm guiding their behavior. If a representational rule happens to be extracted by a reflecting agent, it is still habit and not that rule that continues to guide its behavior. A humanoid autonomous agent could likewise learn to communicate using nondeviant interpersonal distance without any representational rules dictating how many centimeters away to stand by recourse to supervised learning (observation and mimicry) and unsupervised learning (violations).

Conclusion

Affordance perception offers a new paradigm for perception and action in autonomous agents. While traditional three-level systems dissociate perception, planning or reason, and action into separate layers, nonrepresentational affordances involve a dynamic and bidirectional perception-action loop with online planning. Many implementations of affordance perception in AI research have retained the representationalist paradigm even as they seek to integrate affordances. While this is certainly feasible from a technical standpoint, the construct of 'affordance' loses much of its power. An affordance-based robotics that remains largely representationalist has no clear advantage over traditional architectures.

Examining several implementations of cultural affordance perception in AI research, I argue that representations are not *necessary* for cultural affordances. I sketched a possible way for autonomous agents to implement cultural affordance perception by habit gained through supervised and unsupervised learning. AI implementations of affordance perception do not conceptually require a robust representationalism. If affordance-based robotics is to have any advantage over traditional architectures, it may need to reconsider the role of representations in cultural affordance perception.

Acknowledgements

The author would like to thank the Institute for Intelligent Systems for a travel grant to present this paper at the 41st Annual Meeting of the Cognitive Science Society.

References

Awaad, I., Kraetzschmar, G. K., & Hertzberg, J. (2015). The role of functional affordances in socializing robots. *International Journal of Social Robotics*, *7*, 421-438.

Beer, R. D. (2014). Dynamical systems and embedded cognition. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge handbook of artificial intelligence*. Cambridge University Press.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, *6*, 3-15.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, *47*, 139-159.

Chemero, A. (2003). An outline of a theory of affordances. *Ecological Psychology*, *15*(2), 181-195.

Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.

Chu, V., Fitzgerald, T., & Thomaz, A. L. (2016). Learning object affordances by leveraging the combination of human-guidance and self-exploration. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 221-228.

Dennett, D. (1969). *Content and consciousness*. New York, NY: Humanities Press.

Dotov, D. G., Nie, L., & de Wit, M. M. (2012). Understanding affordances: History and contemporary development of Gibson's central concept. *AVANT*, *3*(2), 28-39.

Favela, L. H., & Chemero, A. (2016). The animal-environment system. In Y. Coello & M. H. Fischer (Eds.), *Foundations of embodied cognition: Volume 1: Perceptual and emotional embodiment*. New York, NY: Psychology Press.

Fodor, J. (1985). Fodor's guide to mental representation: The intelligent auntie's vade-mecum. *Mind*, *94*(373), 76-100.

Gallagher, S. (2017). *Enactivist interventions: Rethinking the mind*. Oxford University Press.

Gat, E. (1998). On three-layer architectures. In D. Kortenkamp, R. P. Bonasso, & R. Murphy (Eds.), *Artificial intelligence and mobile robots: Case studies of successful robot systems*. Menlo Park, CA: American Association for Artificial Intelligence.

Gibson, J. J. (1979/2015). *The ecological approach to visual perception* (Classic Ed.). New York, NY: Psychology Press.

Grice, H. P. (1975). Logic and conversation. In J. L. Morgan & P. Cole (Eds.), *Syntax and semantics 3: Speech acts*. London, UK: Academic Press.

Horton, T. E., Chakraborty, A., & St. Amant, R. (2012). Affordances for robots: A brief survey. *AVANT*, *3*(2), 70-84.

Maes, P. (Ed.). (1991). *Designing autonomous agents: Theory and practice from biology to engineering and back*. Cambridge, MA: MIT Press.

Marr, D. (1982/2010). *Vision: A computational investigation into the human representation and processing of visual information*. Cambridge, MA: MIT Press.

Moussaid, M., Garnier, S., Theraulaz, G., & Helbing, D. (2009). Collective information processing and pattern formation in swarms, flocks, and crowds. *Topics in Cognitive Science*, *1*, 469-497.

Nye, B. D., & Silverman, B. G. (2012). Affordances in AI. In N. M. Seel (Ed.), *Encyclopedia of the sciences of learning*. New York, NY: Springer.

Ramstead, M. J. D., Veissière, S. P. L., & Kirmayer, L. J. (2016). Cultural affordances: Scaffolding local worlds through shared intentionality and regimes of attention. *Frontiers in Psychology*, *7*, 1-21.

Raubal, M., & Moratz, R. (2008). A functional model for affordance-based agents. In E. Rome, J. Hertzberg, & G. Dorffner (Eds.), *Towards affordance-based robot control: International seminar, Dagstuhl Castle, Germany, June 5-9, 2006. Revised papers*. Berlin, Germany: Springer.

Rietveld, E., & Kiverstein, J. (2014). A rich landscape of affordances. *Ecological Psychology*, *26*(4), 325-352.

Rome, E., Hertzberg, J., Dorffner, G., Doherty, P. (2006). Towards affordance-based robot control. In *Dagstuhl Seminar 06231 "Affordance-Based Robot Control," June 5-9, 2006*, Schloss Dagstuhl, Germany: Internationales Begegnungs- und Forschungszentrum für Informatik.

Şahin, E., Çakmak, M., Doğar, M. R., Uğur, E., & Üçoluk, G. (2007). To afford or not to afford: A new formalization of affordances towards affordance-based robot control. *Adaptive Behavior*, *15*(4), 447-472.

Stoffregen, T. A. (2003). Affordances as properties of the animal environment system. *Ecological Psychology*, *15*(2), 115-134.

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

Turvey, M. T. (1992). Affordances and prospective control: An outline of the ontology. *Ecological Psychology*, *4*(3), 173-187.

Turvey, M. T. (2019). *Lectures on perception: An ecological perspective*. New York: Routledge.

Varela, F. J., Thompson, E., & Rosch, E. (1991/2016). *The embodied mind: Cognitive science and human experience* (Rev. Ed.). Cambridge, MA: MIT Press.

Wagman, J. B., Caputo, S. E., & Stoffregen, T. A. (2016). Hierarchical nesting of affordances in a tool use task. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(10), 1627-1642.