

Understanding Associative and Cognitive Explanations in Comparative Psychology

Dr. Cameron Buckner
The University of Houston
cjbuckner@uh.edu

This is an Accepted Manuscript of a book chapter published by Routledge/CRC Press in *The Routledge Handbook of Philosophy of Animal Minds* on July 7, 2017, available online: <https://www.routledge.com/The-Routledge-Handbook-of-Philosophy-of-Animal-Minds-1st-Edition/Andrews-Beck/p/book/9781138822887>

I. Introduction

In the introduction to their influential anthology on comparative cognition research, Wasserman & Zentall (2006: 4-5) summarize what I have called that discipline's "Standard Practice":

[Cognition is] an animal's ability to remember the past, to choose in the present, and to plan for the future....Unequivocal distinctions between cognition and simpler Pavlovian and instrumental learning processes... are devilishly difficult to devise...[but] unless clear evidence is provided that a more complex process has been used, C. Lloyd Morgan's famous canon of parsimony obliges us to assume that it has not; we must then conclude that a simpler learning process can account for the learning.... The challenge then is to identify flexible behavior that cannot be accounted for by simpler learning mechanisms. Thus, a cognitive process is one that does not merely result from the repetition of a behavior or from the repeated pairing of a stimulus with reinforcement.

Several ideas can be unpacked from this short characterization of the field. First, there is a default concern for associative explanations of behavior; associative processes must be considered as a possible explanation for any experimental data. Second, there is a default preference for "simpler" associative explanations; producing a plausible associative account of some behavior is seen as a trump card which undermines a cognitive interpretation of the results. Third, these practices are only cogent if associative and cognitive explanations of behavior are mutually-exclusive alternatives.

Combined, these three ideas outline a clear research agenda for the discipline: to carefully devise experimental tasks that could be solved only by the use of a cognitive strategy,

and not by any plausible associative strategy. Though some form of Standard Practice has been with us at least since C. Lloyd Morgan formulated his famous Canon (Morgan 1903), this research program became dominant in the 1960s and 1970s due to the challenge fledgling cognitivists faced in justifying their approach to skeptical behaviorists. They defended their approach by arguing that animals were capable of certain feats which could not be explained in terms of the stock components of the behaviorist toolkit. Love it or hate it—and many influential theorists have recently expressed some ire—there is little doubt most comparative cognition research still fits this mold.

Though this methodology has produced a fine body of research, without a great deal of additional conceptual work it will soon lead the discipline to disaster. We must confront two problems, the first conceptual and the second empirical. The first problem is that the terms ‘associative’ and ‘cognitive’ are equivocal in contemporary practice. The second is that it recently appears that all cognitive processes will be fruitfully describable by associative models. We consider each in turn.

II. Defining ‘cognition’ and ‘association’

Over the millennia, something like a cognitive/associative distinction has manifested itself in a variety of forms, and as a result much discussion about the distinction today involves equivocation and talking-past. Vague dichotomies are notorious in their ability to absorb the hopes and fears of many incompatible perspectives, so a first step to reform is to recognize the terminological diversity in the literature and require theorists to clarify key terms, especially ‘cognition’ and ‘association’.

Let us begin with ‘cognition’. At one extreme, Shettleworth defines ‘cognition’ as any process “by which animals acquire, process, store, and act on information from the environment” (Shettleworth, 2010, 4). As a justification for this inclusive definition, it might be noted that the

term is commonly taken this way in cognitive science more broadly, where it is used to delimit the lower bounds of the subject matter studied by cognitive scientists. However, this definition would class even the most basic forms of classical and instrumental conditioning as cognitive, leaving Standard Practice obviously confused in at least two different ways. Firstly, a label that does not discriminate does no classificatory work, so it would be strange for comparative psychologists to expend so much energy trying to determine whether a process is cognitive. Secondly, such an inclusive definition rules out by fiat the possibility that cognition and association could be mutually exclusive, rendering the attempt to experimentally distinguish them clearly incoherent.

Recognizing these difficulties, others have argued that Standard Practice operates instead with a more restrictive “supercognitive” (Heyes 2012) or “rational” (Dickinson 2012) notion of cognition that the simplest forms of classical and instrumental conditioning do not satisfy. Since the simplest forms of associative learning are ubiquitous in the Animal kingdom, the interesting empirical questions in Standard Practice concern which nonhuman animals have which supercognitive or rational processes and whether the category of supercognitive or rational processes is mutually exclusive with associative processing. To be clear, in the remainder of this article, when I use the word ‘cognition’, I use the term in this more restrictive sense. This interpretation still allows for the possibility that Standard Practice is confused, of course; but if so, it would be a substantive empirical discovery.

Thus, I have argued that Standard Practice holds that cognition requires the manipulation of declarative knowledge, higher-order processes, or symbolic, rule-based reasoning (Buckner, 2011; 2015). Here, learning that a process is cognitive tells us something interesting about the nature of its representational structure and consequently about the flexibility of the behavioral

capacities it enables. Specifically, it suggests forms of processing that are not rigidly bound to particular stimuli and perceptual similarity, enabling adaptive and flexible responding in perceptually novel circumstances. When an animal can arrive at the “rational” solution to a problem that is perceptually dissimilar from those which it has faced in the past—but similar, perhaps, in terms of its underlying logical or causal structure—it is said to display “reasoning” or “insight” that is cognitive in nature. This account leaves much to be desired in terms of empirical precision—significant leeway remains for researchers to disagree as to what counts as a empirical test for rational insight or stimulus independence, leeway we shall begin to constrain below.

Before proceeding further, though, the interpretation of ‘association’ must also be clarified. We might think association by comparison easy to define by indexing it to behaviorist theory circa 1950, perhaps as any learning that can result from the pairing of one stimulus with another or with a behavioral response (a ‘stimulus’ here being any event that can be registered by the sensory organs, such as a light, sound, or odor). The difficulty here is that associative learning theory has progressed in leaps and bounds since the advent of the cognitive revolution—with prominent associationists also now going to great pains to distinguish their approach from behaviorism (Rescorla 1988). As a result, associative learning theory now covers a dizzying and highly-technical array of higher-order stimulus relations, preprocessing of stimuli, cue competition, and even complex architectural ideas (see Table 1). An ecumenical way to delimit the scope of associative learning might be as any form of processing that can be accounted for with a fixed set of relations learned amongst representations of stimuli by observing spatiotemporal contingencies between cues and/or responses. Many authors also add the constraint that the nature of the links themselves—whether causal, temporal, or modal—not also

be represented by the system. As I use the term here, an explanation is associative if it shows how an animal could produce a behavior only by tracking a fixed set of relations amongst stimuli and/or responses presented in its learning history.

Learning Effect/Paradigm	Schematic
Stimulus Generalization	A+ perceptually similar variants of A?
Higher-order Conditioning	A+ AB B?
(Forward) Blocking	A+ AB+ B?
Backward Blocking	AB+ A+ B?
Higher-order Backward Blocking	AC+ CB+ B- A?
Overshadowing	AB+ B?
Sensory Preconditioning	AB- B+ A?
Latent Inhibition	A- A+ A?
Reversal Learning	A+, B- A-, B+ A?, B?
Context-Shifting	A+ in X A? in Y
Negative Patterning	A+, B+, AB- A?, B?, AB?
Value Transfer	A ₁₀₀ B ₀ , C ₅₀ D ₀ BD?

Table 1. An example of some learning paradigms considered part of associative learning theory. *A, B, C* indicate stimuli (such as lights or tones), *X* and *Y* indicate contexts (such as different rooms or times of day), + indicates reward, - indicates no reward, / indicates a break between trial blocks, subscripts indicate the percentage of time a stimulus is rewarded in training, and ? indicates the test situation where an effect is expected. To consider some examples, higher-order conditioning occurs when an animal is conditioned to respond to one stimulus, then the rewarded stimulus is repeatedly paired with a second, neutral stimulus, and the animal later responds to the previously neutral one in isolation (because it has been associated with the originally rewarded stimulus). Overshadowing is found when one stimulus is naturally more “salient” than another, and the overshadowing effect occurs when an animal is only conditioned to respond to one of two stimuli presented together with reward during training. Context shifting occurs when an animal is trained to respond to a stimulus in one context, but does not respond to that stimulus in a different context. Negative patterning occurs when an animal can be trained to respond to two stimuli in isolation, but not to their compound (which requires the animal to create a distinct third representation for the compound stimulus). Value transfer occurs when a more highly rewarded stimulus (such as A₁₀₀ or B₅₀) has some of its value “bleed” to other cues with which it co-occurs, (such as B₀), which can allow preferences to emerge between stimuli with equivalent elemental reward histories (such as B₀ and D₀) because the other stimuli with which each has co-occurred have been differentially rewarded.

III. The Return of Association

The second and even bigger problem with Standard Practice is that, under the ecumenical interpretations of ‘cognition’ and ‘association’ just described, the mutual exclusivity assumption that drives its experimental design appears to be empirically false. Sufficiently flexible associative processes can sometimes implement cognition; or in other words, the same process might be simultaneously, correctly described by both a cognitive and an associative model. Though this important possibility has been widely appreciated in other areas of cognitive science—especially in the debate between classicists and connectionists over cognitive architecture—it comes as a shock to some Standard Practitioners. Nevertheless, I argue it is an inevitable consequence of the other principles they already endorse, discussed above.

The source of this problem is associative learning theory’s surprising potential; its basic principles (discussed above) have not constrained its scope as much as cognitivists originally supposed. The number of processes that appear fruitfully describable in associative terms has dramatically expanded over the past few decades. Associative models are now live competitors as descriptions of many different cognitive capacities, including transitive inference, episodic memory, causal learning, metacognition, goal-directed behavior, imitation, early word learning, and many others. Though some theorists still insist that there is something crucial that associative models will never do, associative learning theory’s continued ability to exceed all predicted limits recommends some epistemic modesty. Considering our previous failures as inductive evidence, we should prepare for the possibility that associative models will eventually be able to fruitfully describe all psychological processes—lest we fall into the same kind of wishful thinking deployed by doomsday prophets continually pushing back the date of the expected apocalypse as it repeatedly fails to materialize.

In fact, this dramatic extension of associative learning theory has been a direct result of the empirical arms race between proponents and skeptics of animal cognition in Standard Practice. A typical pattern that emerges is that a clever cognitivist will devise an experimental test that cannot be passed using current principles of associative learning theory, and, after a high-profile publication, this test comes to be widely used as a benchmark for cognition across different species. A clever associationist will then devise a modest extension of prior associative learning theory that can allow associative models to pass the cognitivist's benchmark. The cognitivist in turn devises a yet more sophisticated behavioral test for cognition that controls for this revised associative mechanism, inspiring yet another modest innovation by the associationists. For many different faculties, this back-and-forth appears capable of continuing indefinitely.

If associative models can eventually accommodate any behavioral data, then theorists face a choice point. On the one hand, if we continue to endorse the assumption that associative models and cognitive models depict mutually exclusive kinds of psychological process, then we should all admit that the hard-nosed associationists will probably win the field—and that cognition does not exist. On the other hand, if (as I recommend) we abandon Standard Practice's mutual exclusivity assumption, then we need to provide specific guidance that allows researchers to know when associative processing has become sufficiently flexible to count as implementing cognition. In short, we would need to develop principled, empirically-plausible methods to distinguish (at least) two different kinds of associative processing, (at least) one of which serves as a deflationary alternative to cognition and the other of which implements cognition.

Though we should not get bogged down in the details here, I have recommended a specific version of the latter approach (2015). The basic idea is to tie the distinction between

cognitive and associative psychological processes to the distinction between multiple memory systems in the brain, with the distinctively cognitive system centered on the hippocampus and other medial temporal lobe structures in mammals and its functional homologues in other classes. The theory of multiple memory systems has been richly elaborated in the field of cognitive neuroscience and is growing in popularity in comparative psychology itself. This body of work provides strong support for the conclusion that there are dissociable memory systems in the brain that, while all fruitfully describable by associative models, differ markedly in the degrees of behavioral flexibility they support—specifically in the forms that have been traditionally assessed by comparative psychology’s benchmarks for cognition. The methodology of Standard Practice can thus largely be salvaged if we reinterpret it as trying to determine which memory system controls some observed behavior.

This gross classification is only the initial stage of study, of course, but determining the memory system that controls a behavior can help guide its future investigation. I have suggested that the labels ‘cognitive’ and ‘non-cognitive’ should be seen as superordinate natural kind terms that organize a variety of more specific psychological kinds like transitive inference, cognitive mapping, theory of mind, and so on. To provide an analogy, they function in psychology like the similarly general labels ‘metal’/‘non-metal’ do in chemistry. Learning that a sample of some unknown element is a metal tells us only highly abstract information, but it does give us a general idea what kind of other properties we should expect the sample to possess (conducts electricity, ductile and solid at room temperature, etc.). In doing so, it tells us which future tests might produce useful results as we continue our investigation into that element’s distinctive characteristics.

Though many articles could be written linking these psychological and neural details, a few metaphors and examples may help explain the view and make it more accessible. Consider the contrasting pictures provided by Tolman (1948) in his classic “Cognitive Maps in Rats and Men”. In that work, Tolman distinguished two different approaches to the study of associative learning that were present in his day. The first, “stimulus response” school held that...

...the rat's central nervous system...may be likened to a complicated telephone switchboard...There are the incoming calls from sense-organs and there are the outgoing messages to muscles...*Learning*, according to this view, consists in the respective strengthening and weakening of various of these connections.

Behavior, according to this school, is generated by elemental stimulus-response links, akin to the telephone operator connecting stimulus inputs to motor outputs in a piecemeal fashion, following that linkage's individual history of reinforcement. The other school, Tolman's “field theorists,” held that...

...in the course of learning something like a field map of the environment gets established in the rat's brain...the intervening brain processes are more complicated, more patterned and often, pragmatically speaking, more autonomous...his nervous system is surprisingly selective as to which of these stimuli it will let in at any given time...the incoming impulses are usually worked over and elaborated in the central control room into a tentative, cognitive-like map of the environment. (Tolman 1948)

Several key points of contrast emerge: whether the animals' representation of its environment forms an integrated whole or a set of disorganized elemental links, whether the effect of any given stimulus is determined by that stimulus' informational value or each is treated indifferently, and whether behavior is determined in a centralized, coordinated manner or via independent stimulus-response links. Though Tolman intended to contrast two competing approaches to the study of associative learning, these metaphors work well if we hold that *both* approaches are right, but characterize different memory systems, with the map-like hippocampal system controlling cognitive processing. That the metaphors can be so easily repurposed may

not be so surprising, given that much of the foundational work on the hippocampal system was derived from O'Keefe & Nadel's classic work on the neural mechanisms behind cognitive mapping (1978).

IV. Looking Forward: Guiding Principles

I close by extracting several principles from an instructive and commonplace example of a clash between different memory systems: conditioned taste aversion, also known as the Garcia Effect. Conditioned taste aversion is a specialized, rapid, and long-lasting form of associative learning that can occur in a single trial between a taste stimulus and nausea, resulting in powerful and enduring aversion to that stimulus in the future. Anyone who has ever overindulged in tequila and later cringed away from a single harmless margarita is in the grips of conditioned taste aversion. No matter how many times one rehearses the fact that one drink poses no real threat, it is not possible to revise the taste-nausea association through explicit reflection alone. This insulation of one inflexible form of associative learning against revision by another, more flexible system provides a vivid example of the kind of dissociation between memory systems that I have been discussing. From this example, we can extract several important principles which can be used to guide future research in comparative psychology.

First Principle: Psychological kinds should be assessed by defeasible tests for property clusters, rather than definitive tests of necessary and sufficient conditions

One obvious difficulty posed by conditioned taste aversion is that it defies one of the most typical characterizations of associative learning: that it be slow and incremental. This complication demonstrates that we must move away from the idea that psychological kinds can be distinguished by neat sets of necessary and sufficient conditions, for accurate characterization of nearly any psychological category is complex and riddled with exceptions. Such exceptions

do not pose a fatal problem to the framework I proposed above, however, for conditioned-taste aversion is in nearly all other relevant ways highly inflexible.

Though it is good to insist that our cognitive and associative hypotheses generate clear predictions, we must give up on the idea of critical tests that can cleanly confirm or falsify such hypotheses in isolation. This simplistic philosophy of science should have died under the lash of the Quine-Duhem thesis, but it has persisted in corners of comparative psychology to this day. Some of the savviest comparative psychologists are now beginning to look instead for correlations amongst clusters of independent behavioral properties (Cheke & Clayton, 2015), which provides a better methodology for assessing the kinds of psychological categories I have been discussing here. In short, the task of assessing which memory system controls a psychological process through behavioral experiment is like trying to determine whether a car has a 4-cylinder or a 6-cylinder engine without opening the hood: both engines do many of the same things and in some conditions the 4-cylinder may outperform the 6-cylinder, but they will reliably differ in their full performance profiles.

Second Principle: Psychological models must be regarded as incomplete descriptions of real underlying phenomena, whose full nature is determined by the neural mechanisms those models target

A difficulty with the move just sketched, however, is that we want to be able to distinguish principled exceptions from unprincipled exceptions. In other words, why should we not count the admission that association may sometimes be more rapid than cognition as an unforgivably ad hoc attempt to salvage an empirically impugned hypothesis? The solution is to tie the criteria for various memory systems to underlying neural mechanisms, and decide whether an exception is principled by seeing whether the two different memory systems can still be successfully empirically distinguished by the other characteristic properties. The key (but often

neglected) idea here is that psychology is the study of the actual causes of behavior in humans and animals, so all models in comparative psychology must aim to describe, at some level of abstraction, real psychological processes operating in humans and animals. By contrast, they cannot—like models of ideally rational economic agents or perpetual motion machines—aim to describe some merely possible system under unrealistic assumptions.

This principle sounds obvious, but neglecting it can quickly lead to mischief. For example, consider the deflationary model of transitive inference proposed by De Lillo et al. (2001), a simple three-layer feed-forward connectionist network (Figure 1) that can demonstrate transitive-like choice when trained on the same sorts of stimuli as animals that have been said to demonstrate the cognitive solution to transitive inference problems. It does so by implementing the associative principle of “value transfer” (Table 1). Surely, the associationist might respond, such a simple model could not be thought to implement cognition, because it is incapable of any other forms of flexibility characteristic of cognition. Thus, they argue, this network shows that transitive-like choice in animals is not cognitive either.

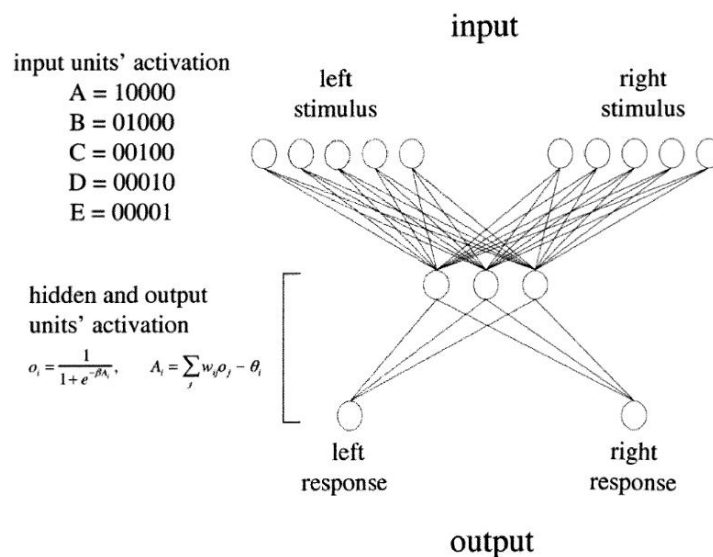


Figure 1. De Lillo et al.'s three-layer, feed-forward model of transitive inference. The five-digit strings representing stimuli are fed to the input units. The formula describes the sigmoidal

transfer function used to determine the activation of the hidden and output units in response to the input vectors.

In comparative psychology, however, it is of little consequence what a disembodied network can or cannot do in isolation. The real question is whether the brains of animals actually implement value transfer without also implementing the other forms of representational flexibility characteristic of cognition. And here, the many lesion and modeling studies on transitive inference suggest that the hippocampal system is responsible for value transfer in the mammalian brain. Thus, if the De Lillo et al. model is relevant at all in comparative psychology, it must be regarded as an incomplete depiction of the much more flexible hippocampal system—and so cannot stand as a general deflationary alternative to cognitive approaches to transitive inference. (For references and a longer discussion of this example, see Buckner 2015.)

To return to conditioned taste aversion, what is known about its neurobiology supports the claim that the exception in question is principled rather than ad hoc. It is for this reason that the exception does not threaten the cognitive/non-cognitive distinction any more than the fact that mercury is a liquid at room temperature threatens the metal/non-metal distinction in chemistry. In rats, at least, conditioned taste aversion appears to be controlled primarily by a specialized and evolutionarily older circuit located in the brain stem. Following lesion and electrophysiology studies, the taste-nausea associations are believed to form at the intersection of the midbrain and pons, in the parabrachial nucleus. Given this location's neurobiology and connectivity, conditioned taste aversion exhibits a number of other surprising features; for example, the lag between the taste stimulus and nausea onset can be extremely long—up to several hours—and can be formed without modulation by higher brain structures, during general anesthesia and deep hypothermia. These associations then trigger aversion reactions via a downstream connection to the amygdalae. Because the rapidity with which conditioned taste

aversion follows from distinctive neural architecture and connectivity that is inflexible in many other relevant ways, this exception does not impugn the strategy of tying the distinction to the theory of multiple memory systems.

A possibly painful corollary of the preceding discussion, however, is that comparative psychologists must give up on the idea that their discipline is independent and autonomous from neuroscience. The sorts of uncertainties of the previous paragraph will only become more common as comparative psychology continues to mature, diverse models proliferate, and the relationships between them—competition, complementation, or implementation?—become more difficult to determine. Not every researcher needs to be wholly multi-disciplinary, but it will become increasingly untenable to insist that every pressing question in psychology be answered by appeal to behavioral data alone.

Third Principle: Associative models do vital explanatory important work even when they are redundant with cognitive models

I close by attempting to forestall a mistaken conclusion that might be drawn from the preceding discussion: that, because associative models can depict implementations of cognition, they are somehow second-rate explanations or uninteresting “implementation stories” for cognitive processes. This unfortunate attitude has been endorsed by some in the older debate between classicists and connectionists about cognitive architecture (e.g. Fodor & Pylyshyn 1988), but it is based on bad philosophy of science. Instead, cognitive and associative models can be independently legitimate models that depict a process with different goals and at different levels of abstraction, with overlapping and complementary explanatory virtues.

A typical difference between cognitive and associative models of the same process is that associative models usually make predictions about fine-grained adjustments in response to the next stimuli observed, whereas cognitive models usually abstract away from this detail to predict

the learning outcomes that reliably emerge from diverse learning histories. Associative models would thus rank more highly on many criteria valued by philosophers of science, especially counterfactual explanatory power, the ability to answer “what if things had been different” questions. Associative models have more counterfactual power because they can make many more specific predictions about arbitrary interactions amongst low-level stimulus representations throughout the whole trajectory of learning. However, to make these predictions they require a daunting amount of background information—researchers must usually know the full associative learning history for that experimental subject regarding the relevant stimuli, information which is unavailable in many laboratory and field contexts and which tends to be highly idiosyncratic. Associative models thus excel at telling you where a particular subject is heading in the next step, whereas cognitive models excel at telling you where the average subject will tend to end up, given a typical learning history. Both explanatory goals are important, and neither reduces to a merely second-rate understudy of the other. (See Buckner 2014 for a case study in the predictive value of the latter kind of hypothesis in theory of mind research).

Conclusion

The Standard Practice of comparative psychology presumes that cognitive and associative causes of behavior are mutually exclusive alternatives, and attempts to distinguish them by means of cleverly controlled experiments. I have provided reasons above for thinking that this methodology is due for a serious revision, but not the wholesale rejection recommended by many recent commentators. If we reinterpret the methodology as trying to determine the memory system under which a behavior is controlled—accepting that all memory systems, even the distinctively “cognitive” ones, can fruitfully be described by associative models—then this methodology can be largely salvaged, and indeed emerge with a strengthened self-understanding.

This revision requires numerous changes of perspective and especially a willingness to cooperate with neuroscience; but if we are up to the task, comparative psychology may continue to enjoy a bright future for many years to come.

References

- Buckner, C. (2011) "Two Approaches to the Distinction between Cognition and 'Mere Association,'" *International Journal of Comparative Psychology*, 24(4).
- (2014) "The Semantic Problem(s) with Research on Animal Mind-Reading," *Mind & Language*, 29(5), 566–589.
- (2015) "A Property Cluster Theory of Cognition," *Philosophical Psychology*, 28(3), 307–336.
- Cheke, L. G., & Clayton, N. S. (2015) "The Six Blind Men and the Elephant: Are Episodic Memory Tasks Tests of Different Things or Different Tests of the Same Thing?" *Journal of Experimental Child Psychology*, 137, 164–171.
- De Lillo, C., Floreano, D., & Antinucci, F. (2001) "Transitive Choices by a Simple, Fully Connected, Backpropagation Neural Network: Implications for the Comparative Study of Transitive Inference," *Animal Cognition*, 4(1), 61–68.
- Dickinson, A. (2012) "Associative Learning and Animal Cognition," *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367(1603), 2733–2742.
- Fodor, J. A., & Pylyshyn, Z. W. (1988) "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition*, 28(1), 3–71.
- Heyes, C. (2012) "Simple Minds: A Qualified Defence of Associative Learning," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1603), 2695–2703.
- Morgan, C. L. (1903). *An Introduction to Comparative Psychology*, London: Walter Scott, limited.

O'Keefe, J., & Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Vol. 3), Oxford: Clarendon Press.

Rescorla, R. A. (1988) "Pavlovian Conditioning: It's Not What You Think It Is," *American Psychologist*, 43(3), 151.

Shettleworth, S. J. (2010) *Cognition, evolution, and behavior*, 2nd Edition. London: Oxford University Press.

Tolman, E. C. (1948). "Cognitive Maps in Rats and Men," *Psychological Review*, 55(4), 189.

Wasserman, E. A., & Zentall, T. R. (2006) *Comparative Cognition: Experimental Explorations of Animal Intelligence*. Oxford: Oxford University Press.

Further Reading

J. Pearce *Animal Learning and Cognition 3rd edition* (New York: Psychology Press, 2013) presents an up-to-date and accessible review of recent advances in associative learning theory.

M. Gluck & C. Meyers *Gateway to Memory: An Introduction to Neural Network Modeling of the Hippocampus and Learning* (Cambridge: MIT Press, 2001) and H. Eichenbaum & N. Cohen *From Conditioning to Conscious Recollection: Memory Systems of the Brain* (Oxford: OUP, 2004) relate this learning theory to computational and anatomical neuroscience. Finally, an excellent collection of different perspectives on these methodological issues can be found in S. Hurley and M. Nudds *Rational Animals?* (Oxford: OUP 2006).

Author Bio:

Cameron Buckner began studying artificial intelligence at Texas Tech University and completed his Ph.D. in Philosophy at Indiana University, Bloomington. After an Alexander von Humboldt Fellowship at Ruhr-University, Bochum, he became an Assistant Professor at the University of Houston.