that produced by S-R mechanisms) must be, beyond reasonable doubt, false (see also Iselin-Chaves et al. [2005] for learning under anaesthesia in humans). It warrants repeating that any dissociation between proposition knowledge and learning is fatal to the current account, and such dissociations do exist (albeit that unambiguous evidence is not widespread).

The insistence that associative accounts rely on nodes that represent whole stimuli in a symbolic manner is also a mischaracterisation. Foreshadowed by Estes's (1950) stimulus sampling theory, associative models explicitly acknowledge that any stimulus comprises multiple features that might each be shared with other stimuli (e.g., Blough 1975; Brandon et al. 2000). This undermines attempts in the target article to characterise the concept of generalisation as an unjustified assumption by which "freedom is gained to explain results" (sect. 6.1, para. 3). In fact, this is an integral and fully specified feature of almost all current associative learning models, and flows directly from the idea that whole stimuli should be considered as collections of potentially overlapping features. Within-compound associations are also treated as "get-out clauses" despite following naturally from, and being explicitly predicted by, standard associative principles. Furthermore, there is evidence for their existence (Rescorla & Durlach 1981) and influence upon cue-competition (e.g., Batsell et al. 2001; Durlach & Rescorla 1980). Although there are examples of particular associative-link models being modified in light of an inability to account for particular results, this does not undermine the fact that principles of generalisation and within-compound associations are instantiated within associative-link models as a class.

Finally, Mitchell et al. criticise associative theory for lacking parsimony because it must predicate two sources for human learning (associative-link and propositional mechanisms). However, associative-link theories are very parsimonious in other ways. Most notably, they can explain aspects of human learning (e.g., sensitization, habituation, perceptual learning) which lie beyond propositional mechanisms. Although associative models inherently require dual-process accounts of human learning, propositional accounts are inherently multiple-process with respect to other phenomena. Hence, proposition-only accounts of human learning are no more parsimonious than dual-process accounts when considered in a broader context.

Operating principles versus operating conditions in the distinction between associative and propositional processes

doi:10.1017/S0140525X09000958

Bertram Gawronski[a] and Galen V. Bodenhausen[b]

[a]*Department of Psychology, The University of Western Ontario, Social Science Centre, London, Ontario N6A 5C2, Canada;* [b]*Department of Psychology, Northwestern University, Evanston, IL 60208-2710.*

bgawrons@uwo.ca **http://publish.uwo.ca/~bgawrons/**

galen@northwestern.edu

http://www.psych.northwestern.edu/psych/people/faculty/bodenhausen/

Abstract: Drawing on our Associative-Propositional Evaluation (APE) Model, we argue for the usefulness of distinguishing between basic operating principles of learning processes (associative linking vs. propositional reasoning) and secondary features pertaining to the conditions of their operation (automatic vs. controlled). We review empirical evidence that supports the joint operation of associative and propositional processes in the formation of new associations.

In contrast to a common assumption of dual-process models, Mitchell et al. argue that the formation of new associations in human memory is an exclusive product of controlled, propositional inferences, and that there is no empirical evidence for automatic processes of associative linking. In response to Mitchell et al.'s conclusion, we argue that their analysis conflates the distinction between the basic operating principles of a given process (i.e., associative linking vs. propositional reasoning) and secondary features pertaining to the conditions of its operation (i.e., automatic vs. controlled). If the conceptual independence of these dimensions is taken into account, the reviewed evidence regarding features of automaticity will be diagnostic about the operation of a particular type of process only to the degree that there is perfect overlap between the two dimensions (automatic = associative; controlled = propositional) – which seems debatable on both conceptual and empirical grounds.

Based on our own Associative-Propositional Evaluation (APE) Model (Gawronski & Bodenhausen 2006; 2007), we argue that the formation of a new association in memory should be understood as an *effect* that could be the result of two conceptually distinct *mechanisms*, associative linking and propositional reasoning. In our APE Model, we define *associative linking* as the creation of a new association between two concepts based on the mere co-occurrence of objects or events independent of the perceived validity of their relation. *Propositional learning* is defined as the creation of a new association as a result of syllogistic inferences about the validity of a given relation. The primary difference between the two processes is their dependency on subjective validity, in that only propositional learning, but not associative linking, takes the perceived validity of relations into account (see also Strack & Deutsch 2004). As such, the two mechanisms should lead to the same outcome when the co-occurrence of two objects or events is interpreted as reflecting a valid relation. However, the two mechanisms may lead to different outcomes when the co-occurrence between two objects or events is regarded as non-diagnostic or invalid. This conceptualization incorporates Mitchell et al.'s emphasis of truth values as a core feature of propositional reasoning. However, it differs from Mitchell et al.'s approach, in that assumptions about automatic features represent empirical claims about the boundary conditions of the operation of the two processes, rather than defining characteristics that could be conversely used to identify their operation in a particular case.

To empirically distinguish between the two processes, we suggest that the actual operation of associative and propositional processes should be identified by means of their interactive effects on associations and beliefs. In the APE Model, we define *associations* as mental links between concepts independent of their subjective truth or falsity; *beliefs* are defined as the endorsed relations that are implied by validated or invalidated associations. This distinction has proven its usefulness in the social-cognitive literature, showing that activated associations can produce behaviors that are congruent with these associations, even when the relations implied by these associations are regarded as invalid (for a review, see Strack & Deutsch 2004).

More importantly, there is suggestive evidence that such dissociations can sometimes be due to antagonistic effects of associative linking and propositional reasoning during the encoding of new information (e.g., Gawronski et al. 2008; Rydell et al. 2006), supporting the usefulness of the proposed distinction in the formation of new associations. The basic notion of these studies is that the mere co-occurrence between two objects can create a mental association between these objects, even though the validity of the implied relation is rejected at the propositional level. Empirically, these differences are often reflected in dissociations between implicit and explicit measures (Fazio & Olson 2003), such that implicit measures (e.g., sequential priming tasks) reflect the mere co-occurrence between the two objects, whereas explicit measures (i.e., self-reported judgments) reflect the perceived validity of the implied relation.

Other evidence that is consistent with the notion of associative linking comes from research on spontaneous trait transference (e.g., Skowronski et al.1998), in which communicators have

been shown to become associated with the traits they ascribe to others. In most cases, there is no logical basis to infer that a communicator has a particular trait (e.g., tidy) simply because he or she describes that trait in another person. Hence, it seems reasonable to assume that any such associations are the product of associative linking rather than propositional reasoning (Carlston & Skowronski 2005). To be sure, such associative linking processes may still depend on perceivers' attention, processing goals, or awareness of the co-occurrence. However, this by itself does not make the underlying learning process propositional, as defined in the proposed conceptualization.

Another important issue in this context is Mitchell et al.'s concern that proposing mutual interactions between associative and propositional processes would make the distinction between the two processes obsolete. Such interactions are a core assumption of our APE Model, which assumes that mutual interactions between the two processes are reflected in different mediation patterns of experimentally induced effects on activated associations and endorsed beliefs (Gawronski & Bodenhausen 2006). Specifically, we argue that associative linking will often produce parallel effects on associations and beliefs, such that newly created associations provide the basis for explicitly endorsed beliefs. Conversely, newly created associations may be the product of propositional inferences, such that new beliefs generated in the course of validating currently accessible information may be stored in associative memory. Drawing on the abovementioned distinction between implicit and explicit measures, the first case is assumed to produce parallel effects on both kinds of measures, with effects on the explicit measure being fully mediated by the implicit measure. In contrast, the second case should produce parallel effects on both kinds of measures, with effects on the implicit measure being fully mediated by the explicit measure.

An illustrative demonstration of these diverging mediation patterns is a recent study by Whitfield and Jordan (submitted), who combined an implicit evaluative conditioning (EC) procedure (Olson & Fazio 2001) with a propositional impression formation task that used descriptive information about the conditioned stimulus. Their results showed that both the EC procedure and the impression formation task produced parallel effects on both explicit and implicit measures. However, in line with the predictions of the APE Model, EC effects on the explicit measure were fully mediated by the implicit measure, whereas impression formation effects on the implicit measure were fully mediated by the explicit measure (for related findings, see Gawronski & LeBel 2008; Gawronski & Strack 2004; Gawronski & Walther 2008).

Taken together, these results suggest that a conceptual distinction between associative and propositional processes in terms of their operating principles (rather than automatic vs. controlled features) has testable and empirically supported implications. More importantly, our analysis implies that the formation of new associations in memory can be the product of either associative or propositional processes, and that Mitchell et al.'s insightful review may speak only to the automatic versus controlled nature of these processes rather than to the general irrelevance of associative processes in human learning.

Rational constructivism: A new way to bridge rationalism and empiricism

doi:10.1017/S0140525X0900096X

Alison Gopnik

Department of Psychology, University of California at Berkeley, Berkeley, CA 94704.

Gopnik@berkeley.edu

Abstract: Recent work in rational probabilistic modeling suggests that a kind of propositional reasoning is ubiquitous in cognition and especially in cognitive development. However, there is no reason to believe that this type of computation is necessarily conscious or resource-intensive.

There is a paradox at the heart of cognitive science. Human beings (and some animals) seem to have abstract, hierarchical, structured representations of the world. These representations allow us to make a wide range of novel predictions and produce a wide range of novel behaviors. And these representations seem to be accurate – they capture the structure of the world, and they improve as we learn more about the world. But the information provided by our senses, our one direct source of evidence about the world, is very different from these representations. It is a noisy, probabilistic, and chaotic set of contingencies among specific concrete inputs, apparently far removed from the true structure of the world itself.

In the past 2000 years of western philosophy, and the past 50 years of cognitive science, there have been two very different approaches to resolving this paradox. One tradition (nativist, rationalist, propositional, "East Coast") argues that cognition does indeed involve abstract, hierarchical, structured representations. It only appears, however, that we infer these representations from the evidence of our senses. In fact, these representations must be there innately, and are only slightly modified by learning. Small details may be filled in by experience, or alternative parameters may be triggered by different experiences. But the fundamental structure of the representations is there from the start. The alternative tradition (empiricist, associationist, connectionist, "West Coast") argues that it only appears that we have abstract, hierarchical, structured representations. In fact, our novel predictions and behaviors are based on the complex contingency patterns among individual sensory inputs, patterns that we extract through associative mechanisms.

There have sometimes been arguments for a kind of dismissive co-existence between these two approaches. The rationalists say that most cognition is the result of innate abstract representations, but mere associationist processes may play a role in very automatic, low-level kinds of behavior. The empiricists say that associations are responsible for most cognition, but there may be explicit, conscious, and sophisticated propositional reasoning layered on top. These two-process views both suggest that there is some relationship between the sophistication, power, and likely domain of the representations and their computational character – associations are "low-level" and propositions are "high-level." They just disagree on whether most cognition falls on one side or the other.

The target article is in this general tradition, though it endorses the idea that propositional representations can account for even classical associationist phenomena, such as conditioning. But Mitchell et al. also argue that the propositional representations they endorse are resource-intensive, subject to conscious reflection, and can be understood as beliefs – they are "high-level."

In cognitive development, going back to Piaget, there has been a long tradition of trying to elude the rationalist/empiricist dichotomy with "constructivist" theories. A constructivist account should allow us to actually infer highly structured representations accurately from patterns of contingency in the data. The most recent constructivist project has been the "theory theory" – the idea that children develop intuitive theories from evidence in the way that scientists do. But the theory theory, like earlier constructivist theories, has suffered from a lack of computational precision and specific learning mechanisms.

However, in the last 10 years or so there has been increasing excitement about a new theoretical view that provides a computationally rigorous basis for the constructivist project. This approach might be called "rational probabilistic modeling." This view, unlike classical empiricist views, proposes structured, abstract, hierarchical representations. But unlike classical