Advanced Review

Inductive reasoning

Brett K. Hayes,[1*] Evan Heit[2] and Haruka Swendsen[3]

Inductive reasoning entails using existing knowledge or observations to make predictions about novel cases. We review recent findings in research on category-based induction as well as theoretical models of these results, including similarity-based models, connectionist networks, an account based on relevance theory, Bayesian models, and other mathematical models. A number of touchstone empirical phenomena that involve taxonomic similarity are described. We also examine phenomena involving more complex background knowledge about premises and conclusions of inductive arguments and the properties referenced. Earlier models are shown to give a good account of similarity-based phenomena but not knowledge-based phenomena. Recent models that aim to account for both similarity-based and knowledge-based phenomena are reviewed and evaluated. Among the most important new directions in induction research are a focus on induction with uncertain premise categories, the modeling of the relationship between inductive and deductive reasoning, and examination of the neural substrates of induction. A common theme in both the well-established and emerging lines of induction research is the need to develop well-articulated and empirically testable formal models of induction. © 2010 John Wiley & Sons, Ltd. *WIREs Cogn Sci* 2010 1 278–292

Inductive reasoning involves making predictions about novel situations based on existing knowledge. These predictions are necessarily probabilistic. For example, if you are told that grizzly bears had a certain kind of enzyme you might be moderately confident, but by no means certain, that this property generalizes to other bears. Induction corresponds to much of the reasoning that people do in everyday life. Predicting whether it is likely to rain tomorrow, how your partner will react to the box of chocolates you purchased as a gift, or whether stock prices will rise in the next six months all involve some form of induction. More generally, induction is involved in a range of cognitive activities such as categorization, probability judgment, analogical reasoning, scientific inference, and decision making. The pervasive nature of induction is one of the reasons why it has become an important area of study for cognitive scientists. Another reason is that inductive reasoning seems to address one of the core questions of cognitive science, namely how knowledge is generalized from known to unknown cases.

Much of what we have learned about the cognitive processes involved in inductive reasoning has come from studies of category-based induction. This typically involves making an inference about the properties of the members of some conclusion category, based on knowledge of the properties of some premise category or set of categories. Hence, a property of the premise category of grizzly bears might be generalized to the conclusion category of 'all bears' because of your knowledge of the relations between these categories.

This review summarizes many of the key empirical phenomena identified in research on category-based induction and critically evaluates major theoretical models of inductive reasoning. This is not the first review of category-based induction.[1–3] The past few years, however, have seen an increase in the range of inductive phenomena studied, with a particular emphasis on the effects of causal and property knowledge. As the number of phenomena to be explained has expanded so too have the number and scope of theoretical models. A number of entirely new areas of investigation have also opened up. This review examines two such areas including studies of induction with objects whose category membership is uncertain and examination of the neural bases of reasoning.

*Correspondence to: B.Hayes@unsw.edu.au

[1]School of Psychology, University of New South Wales, Sydney, Australia

[2]School of Social Sciences, Humanities and Arts, University of California, Merced, Merced, CA, USA

[3]Mood and Anxiety Disorder Research Program, National Institutes of Health, Bethesda, MD 20892, USA

DOI: 10.1002/wcs.44

PHENOMENA THAT DEPEND ON TAXONOMIC SIMILARITY

Empirical studies of category-based induction with adults and children have identified the touchstone phenomena in Table 1. These are divided into two types. Those in the top panel involve the generalization of "blank" or unfamiliar properties (e.g., "has property P"). Because people have little prior knowledge of the distribution of such properties, generalization is assumed to be driven mainly by taxonomic relations between the premise and conclusion categories. Those in the lower panel involve more complex premise–conclusion relations, sometimes including prior knowledge about the property being generalized.

Premise–Conclusion Similarity Promotes Induction

A robust finding is that the probability of generalizing a novel property from a premise to a conclusion category is a positive function of their perceived similarity.[4] Hence, people are more likely to generalize a blank property from leopards to lions than from leopards to koalas. This is one of the first induction principles to emerge developmentally, appearing to be well understood by 13 months of age,[5] although there is an ongoing debate as to whether adults and children compute similarity in the same way when doing induction.[6–8]

Premise (and Possibly Conclusion) Typicality Promotes Induction

Another robust inductive principle honored by both adults and young children is premise typicality.[4,9–11] Premises that are more typical or representative of a general category are more likely to support property induction than are less typical premises. Hence, people are more likely to generalize a property from sparrows to geese than from penguins to geese. A related finding is premise–conclusion asymmetry[4,10,12] whereby the projection from sparrows to geese is seen as stronger than from geese to sparrows.

Evidence of an effect of the typicality of the conclusion has been less forthcoming. Early work[9,13] suggested that people are insensitive to conclusion typicality. However, when the relative similarity of premise and conclusion items is controlled, induction to more typical conclusions is stronger than induction to atypical conclusions.[14] For example, a novel property like 'needs vitamin K for liver function', is more likely to be generalized from koalas to tigers than from koalas to guinea pigs.

Conclusion Homogeneity Promotes Inductive Inference

All things being equal, adults[4,15] and children[16] prefer to generalize to less variable (homogenous) categories than to heterogeneous categories. For example, stronger inferences are made going from a premise category (e.g., bears) to a more homogenous conclusion category (e.g., mammals) than to a more general and heterogeneous category (e.g., animals) (but see Ref 17 for evidence on the limits of the effects of conclusion homogeneity).

Greater Diversity and Number of Premises Promotes Induction

Philosophers of science have suggested that the relative diversity of evidence should affect its generalization so that properties shared by more diverse premises are more likely to generalize than those of less diverse premises.[18] Adults are sensitive to this principle, generalizing properties more strongly from diverse sets of premise categories (e.g., lions and mice) than from nondiverse sets (e.g., lions and tigers).[19,20] Consistent with normative principles, adults are also sensitive to premise monotonicity. Increasing the number of premises that share a property increases the likelihood that the property will be generalized.[4,21,22]

Some Inductive Inferences Violate Rational Norms

Most of the inductive phenomena considered so far are broadly consistent with what many believe are normative principles.[23] Equally interesting are cases where inductive intuitions deviate from rational norms. One example is the *inclusion fallacy* in which generalization from a specific premise to a more general conclusion category (e.g., from crow to bird) is stronger than generalization to a conclusion (e.g., to ostrich) nested within the more general set.[4] Even more striking deviations from rationality involve perceived differences in the inductive strength of arguments that are deductively equivalent. Sloman, for example, demonstrated a *premise specificity* effect where generalizations from items like birds to sparrows were rated as stronger than generalizations from animals to sparrows.[24,25]

SIMILARITY-BASED INDUCTION MODELS

Similarity–Coverage Model

The first detailed model proposed to explain these induction phenomena focused on taxonomic

 © 2010 John Wiley & Sons, Ltd.

TABLE 1 | Touchstone Inductive Phenomena with Examples (Terms to the Left of the / are Premises; Terms to the Right are Conclusions. Predicates Are Shown in Italics)

Panel A. Phenomena That Depend on Taxonomic Similarity	Example of a Stronger Argument	Example of a Weaker Argument
1. Premise–conclusion similarity promotes induction	Leopards *have property X*/lions *have property X*	Leopards *have property X*/koalas *have property X*
2. Premise (and possibly conclusion) typicality promotes induction	1. Sparrows *have property X*/geese *have property X*	1. Penguins *have property X*/geese *have property X*
	2. Vultures *have property X*/sparrows *have property X*	2. Vultures *have property X*/quail *have property X*
3. Conclusion homogeneity promotes inductive inference	Bears *have property X*/mammals *have property X*	Bears *have property X*/animals *have property X*
4. Greater diversity and numerosity of premises promotes induction	1. Lions + mice *have property X*/mammals *have property X*	1. Lions + tigers *have property X*/mammals *have property X*
	2. Lions + mice + cows + Bears *have property X*/*mammals have property X*	2. Lions + mice *have property X*/mammals have *property X*
5. Some inductive inferences violate rational norms	1. Crows *have property X*/birds *have property X*	1. Crows have *property X*/ostriches *have property X*
	2. Birds *have property X*/sparrows *have property X*	2. Animals *have property X*/sparrows *have property X*
Panel B. Phenomena That Depend on Background Knowledge	**Example of a Stronger Argument**	**Example of a Weaker Argument**
1. Domain expertise can alter or reverse standard induction phenomena.	Coatimundi + raccoons *get disease X* [less diverse premises]/mammals *get disease X*	Shrews + otters *get disease X* [more diverse premises]/mammals *get disease X*
EG Nondiversity in Biological reasoning of Itza Maya		
2. Property knowledge alters the relations used in induction	1. Sparrows *have a ulnar artery*/hawks *have an ulnar artery*	1. Tigers *have an ulnar artery*/hawks *have an ulnar artery*
	2. Tigers *study their food before attacking*/hawks *study their food before attacking*	2. Sparrows *study their food before attacking*/hawks *study their food before attacking*
3. Salient relations between premises and conclusions can override similarity-based induction	1. Polar bears + antelopes *have property X*/animals have property X	1. Polar bears + penguins *have property X*/animals have property X
	2. Brown bears *have property X*/buffalo *have property X*	2. Brown bears + polar bears +grizzly bears *have property X*/buffalo *have property X*
4. Causal relations override similarity relations	1. Bananas *have property X*/monkeys *have property X*	1. Mice *have property X*/monkeys *have property X*
	2. Rabbits + zebras *have property X*/butterflies *have property X*	2. Rabbits + carrots *have property X*/butterflies *have property X*
	3. Carrots *have property X*/rabbits *have property X*	3. Rabbits *have property X*/carrots *have property X*

similarity between premise and conclusion categories. The similarity–coverage model[4] assumes that the probability that a blank property generalizes from a set of premises to a conclusion is determined by two factors. The similarity component involves computing the feature overlap between the premise and conclusion category. The second component, *coverage*, involves computing the average maximum similarity of the premise categories to the lowest-level category that includes both premises and conclusions.

The similarity component is sufficient to explain premise–conclusion similarity effects. The other

 © 2010 John Wiley & Sons, Ltd.

phenomena in the top panel of Table 1 are assumed to be driven by similarity and coverage. Premise typicality, for example, arises because more typical premises have higher mean similarity to the more inclusive conclusion category (i.e., better coverage) than less typical premises. Because more diverse premises provide better coverage, they will be seen as stronger basis for inductive projection than less diverse premises. More specific or homogenous conclusion categories also give higher coverage for premises included within the category than more general conclusions.

The model also suggests a straightforward explanation of developmental findings on category-based induction by assuming that computing coverage is a more complex operation than computing similarity. Hence, an understanding of arguments based only on similarity should be acquired first, followed by arguments involving coverage with a more general conclusion supplied (e.g., lions, mice/mammals), followed by arguments where the inclusive conclusion has to be generated (e.g., lions, mice/rabbits). This predicted progression is close to what has been observed in many developmental studies[11,26,27] (but see Ref 28 for a different view).

Feature-Based Induction Model (FBIM)

FBIM[24] also emphasizes the similarity of premise and conclusion categories as a basis for induction but does so without the retrieval of superordinate categories. FBIM is implemented as a connectionist network that learns associations between input nodes representing the features of the premise categories and an output node for the property to be considered. The conclusion activates the same output node in proportion to the features shared with the premise. Generalization increases as a function of similarity between the premise and the conclusion but will be reduced by the presence of 'rich' conclusion categories that contain many features.

FBIM explains many of the effects attributed to coverage in the Similarity–Coverage Model (SCM) without generating superordinate conclusions. It can also account for nonnormative fallacies like premise specificity. Unfortunately, FBIM fails to give a straightforward account of premise typicality. Although it predicts a typicality effect for arguments with general conclusions, it has difficulty explaining typicality for specific conclusions. It also has difficulty explaining cases where premise typicality has an independent effect on induction when the respective similarities of typical and atypical premises to a conclusion are equated.[9]

Summary

Both the similarity-coverage model and FBIM explain an impressive range of inductive phenomena, although neither model captures all of these results, at least not without making additional assumptions (Table 2). More seriously, over the past decade, a new set of inductive phenomena have been documented that represent a serious challenge to both models (see Section B of Table 1). These phenomena involve inductive generalization based on more complex relations between premises and conclusions (e.g., causal links) and/or specific knowledge about the property being generalized. The next sections summarize the main findings and then examine theoretical explanations of these more complex forms of induction.

INDUCTIVE PHENOMENA THAT DEPEND ON BACKGROUND KNOWLEDGE

Domain Expertise Can Alter or Reverse Standard Induction Phenomena

There are many cases where certain kinds of prior experience or expertise with the domain from which premise and conclusion categories are drawn undermines many of the effects summarized in the top panel of Table 1 or even reverses them. Many such studies have compared cultural groups with different levels of experience with the biological world. One series examined inductive reasoning about local mammals and birds in the Itza Maya of Central America and North American undergraduates. Unlike undergraduates, the Itza Maya showed either no diversity effects or a reversal of the effect.[29,30] Similar patterns have been found in comparisons of biological induction between urban and rural children and between urban American undergraduates and native Americans living in traditional communities.[31–33] Nondiversity and weak or absent typicality effects have also been found in those with expertise in specific areas of biology.[29,33]

The most likely explanation for these findings is that domain experts often generalize properties on the basis of relations that differ from those used by nonexperts. While nonexperts rely primarily on taxonomic similarity, experts invoke causal and ecological principles. For example, tree experts often reasoned about arguments involving novel tree diseases in terms of how widely planted different kinds of trees are and their susceptibility to disease.[34] Notably though, violations of the standard induction phenomena among experts are usually found only

 © 2010 John Wiley & Sons, Ltd.

TABLE 2 | Comparison of Explanatory Scope of Theoretical Models of Inductive Reasoning

	Similarity–Covera	FBIM	Relevance Theory	SimProb	Heit's Bayesian Model	Structured Statistical Representations
Similarity Coverage-based phenomena						
1. Premise–conclusion similarity	✓	✓	✓	✓	✓	✓
2. Premise and conclusion typicality	Only premise typicality	X	✓	X	✓	✓
3. Conclusion homogeneity	✓	✓	✓	X	✓	✓
4. Premise diversity and numerosity	✓	✓	✓	X	✓	✓
5. Violations of rational norms	Some cases	✓	Some cases	X	X	✓
Knowledge-based phenomena						
1. Domain expertise	X	X	✓	X	✓	✓
2. Property knowledge	X	✓	✓	✓	✓	✓
3. Salient relations between premises and conclusions	X	X	✓	X	✓	✓
4. Causal relations	X	X	✓	X	X	X
General model characteristics						
Computationally implemented	✓	✓	X	✓	X	✓
Applied to phenomena in inductive development	Some cases	X	Some cases	Some cases	X	Some cases

for stimuli that lie within the domain of expertise. Fish experts, for example, used causal knowledge to generalize a novel disease property ('has a disease called sarca') but used taxonomic similarity to generalize a novel blank property ('has a property called sarca').[34] Another important finding is that experts use more complex reasoning only when they have sufficient time to consider an inductive problem. Under time pressure, experts often default to induction based on taxonomic similarity.[35] Hence, the best way of characterizing property induction by domain experts is that it involves consideration of a broad range of premise–conclusion relations that includes, but is not limited to, taxonomic similarity.

Property Knowledge Alters the Relations Used in Induction

A parallel effect of the selective generalization of properties has been found with nonexperts in studies that use more familiar (i.e., nonblank) kinds of properties.[36–38] Heit and Rubinstein, for example, found that anatomical properties (e.g., 'has an ulnar artery') are more likely to be generalized from sparrows to hawks than from tigers to hawks but that for behavioral properties (e.g., "studies its food before attacking") the pattern reverses. This suggests that property knowledge can alter the way that people compute the similarity between base and target stimuli (birds are likely to have similar biological features but predators are likely to have similar hunting behaviors). In a similar vein, 7-year olds (but not younger children) were found to be more likely to select a taxonomic category as the appropriate conclusion when generalizing biochemical properties but selected a script category when properties were situational.[33,39]

Salient Relations Between Premises and Conclusions Can Override Similarity-Based Induction

Another important finding is that many of the standard induction effects listed in the top panel of Table 1 can be altered or reversed when premises and conclusions are presented in certain contexts.[33,40] Medin et al. found that the diversity effect can be reversed when more diverse premise categories share some distinctive relation (but see Ref 19). For example, a novel property was judged less likely to generalize from polar bears and penguins (high diversity set) to other animals than from polar bears and antelopes (low diversity set) to other animals. This *nondiversity through property reinforcement* suggests that people look for shared distinctive relations

 © 2010 John Wiley & Sons, Ltd.

between multiple premises (arctic habitat seems a likely contender in the above example). If the relation is not shared by the conclusion category, then it is unlikely to be generalized. A related effect is *nonmonotonicity via property reinforcement* in which the addition of premises that highlight a distinctive relation reduces generalization to other categories.[40]

Priming of a distinctive relation among premise conclusion categories can also lead to nonnormative reasoning.[40] For example, people are more likely to generalize a property from pasta to both potatoes and rice than to just one of these categories, although the extent to which people are susceptible to this fallacy may depend on their general ability.[41]

Causal Relations Override Similarity Relations

When causal relations between premise and conclusion are present, these are often used as a basis for induction. Arguments with salient causal relations between dissimilar premise and conclusion categories are perceived to be as strong or stronger as arguments with more similar premise–conclusion pairs lacking any causal links.[40,42,43] Moreover, children as young as five are capable of appreciating the inductive potency of causal relations.[44,45]

Causal relations are also preferred as a basis for inductive projection over heuristics like premise typicality, diversity, and monotonicity.[40,43,46,47] In a series of studies with artificial categories, Rehder orthogonally manipulated typicality and diversity and the presence of causal relations linking a to-be-generalized feature to the category. Similarity-based effects were almost entirely eliminated when generalization could be based on causal relations.[43]

The presence of causal relations between premises and conclusions can also lead to asymmetries in inductive projection, with generalization flowing from causes to effects (e.g., a property of carrots passing to rabbits) viewed as stronger than generalization from effects to causes (e.g., a property of rabbits generalizing to carrots). They can also lead to conjunction fallacies such as perceiving that the properties of some causal agent (e.g., cows) are as likely to generalize to multiple conclusions (e.g., milk and ice cream) as to a single conclusion.[41,48]

THEORETICAL EXPLANATIONS OF KNOWLEDGE EFFECTS

Similarity-Based Theories

Because the similarity coverage model assumes that induction is driven primarily by the taxonomic similarity of premise and conclusion categories (and their superordinates), it has difficulty accounting for the effects of experience, premise context, property, and causal knowledge on induction. It may be that the model could be adapted to accommodate some of these effects by assuming that different similarity metrics are applied (e.g., similarity between premises and conclusions based on shared causal relations) but the details have yet to be specified, and little of the explanatory work would be done by the core SCM mechanisms.

In part because it is a more flexible model, FBIM fares slightly better in explaining context effects. The model allows features to be weighted by the number of categories in such a way that features shared by all premise categories would have the greatest weight. Such selective weighting could explain the effects of property reinforcement manipulations on diversity and monotonicity. It is not clear, however, how FBIM could accommodate the effects of knowledge about casual links between premises and conclusions or the effects of property knowledge.

Because of these difficulties, recent interest has shifted to broader theoretical frameworks for explaining inductive reasoning. Still, models like SCM and FBIM may be incomplete rather than incorrect. They still account for many results, and even when people have rich domain knowledge they often default to taxonomic similarity when reasoning about novel properties.

Relevance Theory

This approach has the specific aim of predicting and explaining context sensitivity and causal knowledge effects in induction.[40,41] The core assumption is that people evaluating inductive arguments actively compare premise and conclusion categories, and (in the multiple premise case) different premises, and that this process activates distinctive relations between these categories that serve as candidates for inductive projection. All things being equal, causal relations are assumed to be highly distinctive and more likely to influence induction than taxonomic or thematic relations. When nonblank properties are used, these further specify the relevant dimensions on which premises and conclusions should be compared, so that different kinds of nonblank properties (e.g., biological versus behavioral) can produce different patterns of inductive projection.

A positive feature of relevance theory is that it applies the same general principles to explain the touchstone effects involving taxonomic similarity (Table 1, panel A) and the effects of context and

 © 2010 John Wiley & Sons, Ltd.

knowledge (Table 1, panel B). For example, a property of lions and tigers is less likely to project to other animals than a property of polar bears and lions, because the less diverse set activates a distinctive property ('large cats') that is not shared by most instances of the conclusion. The same mechanism predicts that the even more diverse set of polar bears and penguins will be less potent for the same conclusion.

The effects of causal relations are explained using similar mechanisms. A projection from grass to cows is strong because the premise–conclusion order suggests a very distinctive causal process (ingestion). Note that this same principle may also explain the asymmetries in premise–conclusion inference and conclusion typicality effects that are problematic for many other models.

Although the breadth of relevance theory is appealing, a number of its assumptions need to be refined so that it can generate further testable assumptions. Most notably, the process by which the informativeness or distinctiveness of the features of premise and conclusions is computed and how this varies with different reasoning contexts, needs to be specified more clearly. Moreover, a number of the key process assumptions of the model (e.g., that for multiple premise arguments people begin induction by comparing premises) have yet to be tested.

SimProb

At the other end of the spectrum, two models, GAP[38] and its more recent revision, SimProb,[49] have been implemented computationally but have more limited scope in explaining property knowledge effects. SimProb conceives of induction as a special case of conditional probability judgment. The conditional probability of a conclusion given a premise is assumed to be equal to the prior probability of the conclusion, weighted according to premise–conclusion similarity and the degree to which the premise and predicate are surprising. When premise–conclusion similarity is controlled, predicates that are more surprising in the light of background knowledge lead to a greater increase in inductive strength than less surprising premises. Hence, the model correctly predicts that arguments like 'Poodles can bite through wire therefore German Shepards can bite through wire' are perceived as having greater inductive strength than 'Dobermans can bite through wire therefore German Shepards can bite through wire'.

A strength of SimProb is that a single computational model can be applied to explain phenomena involving blank properties and nonblank properties. The model also provides a formal metric for measuring the extent to which premises are seen as surprising or informative. Empirically, the model has been shown to produce good fits to single and multiple premise arguments that vary in their similarity to a given conclusion and in their perceived prior probability.

One limitation is that premise–conclusion similarity is still conceived of as a fixed value. In that respect, SimProb fails to capture the notion that similarity is dynamic and may vary with the property in question. Although, in principal, the model could be applied to a wider range of phenomena including typicality, diversity, and causal knowledge effects, these extensions have yet to be articulated.

Bayesian Models

Heit[50] suggested that induction is best thought of as a form of Bayesian belief revision. His model assumes that when generalizing a novel property (e.g., disease X) between familiar categories (e.g., from sparrows to crows) people will access their prior knowledge about the distribution of familiar properties. They will know, for example, that certain properties are true of all birds including sparrows and crows, but that other properties are limited just to the premise or the conclusion. The inductive problem is to determine which of these distributions the novel property most closely resembles. To solve the problem, the Bayesian model treats the premises in an inductive argument as evidence, which is used to revise beliefs about the prior hypotheses according to Bayes's theorem. Once these beliefs have been revised, the plausibility of the conclusion is estimated.

This model successfully predicts most of the key results in the first panel of Table 1. It can be extended to arguments with nonblank properties by assuming that such properties prime the retrieval of specific types of priors. Hence, when asked whether a biological property of hawks is more likely to generalize to sparrows or tigers, people will retrieve prior knowledge about anatomical properties, whereas inductions about behavior will prime knowledge about familiar behavioral properties.

Like other early Bayesian models of induction,[51] Heit's model does not provide details about how prior probabilities are computed. It is also not clear whether the model gives priority to causal relations over other kinds of properties. Some of these issues have been addressed in more recent Bayesian accounts.[52,53] Like earlier models, the *structured statistical* approach[53] assumes that induction involves making an inference about the probability of a conclusion given the observed premises. Critically,

 © 2010 John Wiley & Sons, Ltd.

this approach assumes that the priors used as inputs into Bayesian calculus are based on intuitive theories. These theories are instantiated as structured representations of the relations between categories and prior beliefs about the distribution of features across categories. An important assumption is that different kinds of structured representations will be retrieved depending on the type of property being generalized. This idea is illustrated in Figure 1. When induction involves taxonomic properties, the default structure is a hierarchical tree. For spatial or quantitative properties, categories are represented according to their distance in dimensional space on the relevant property (e.g., size). Causal properties are represented in a directed causal graph. Each structure leads to the activation of a different set of priors about the distribution of known features. Hence, in the causal case knowing that a mouse has a disease may lead to retrieval of knowledge of relevant food-chain relations, activating many features of cats. In the case of blank properties, different conceptual domains are associated with different default representations. The tree structure, for example, is thought to be the default for generalizing blank properties in biological categories.[54] Once priors are derived from the appropriate structured representation, a single Bayesian inference mechanism is used to derive feature inferences.

This model (or set of models) was applied to a range of induction data sets, including many of the key findings from both panels of Table 1. In general, the fit of the structured statistical models was impressive compared with other quantitative models like similarity–coverage.[53] Critically, the relative fit of models based on different structured representations depended on the property being generalized. Bayesian predictions based on taxonomic 'tree' priors produced a good fit to induction involving genetic properties but a poor fit to induction with disease properties. Predictions based on a causal model showed the opposite trend.

The structured statistical approach is impressive in that it provides a clear mechanism for deriving prior probabilities from background knowledge, and allows for the flexible application of different kinds of knowledge. Not surprisingly, there are still some effects of knowledge on induction which the model has difficulty with. Although it can deal with causal relationships that can be represented in a chain or web, it is not clear how the model would deal with

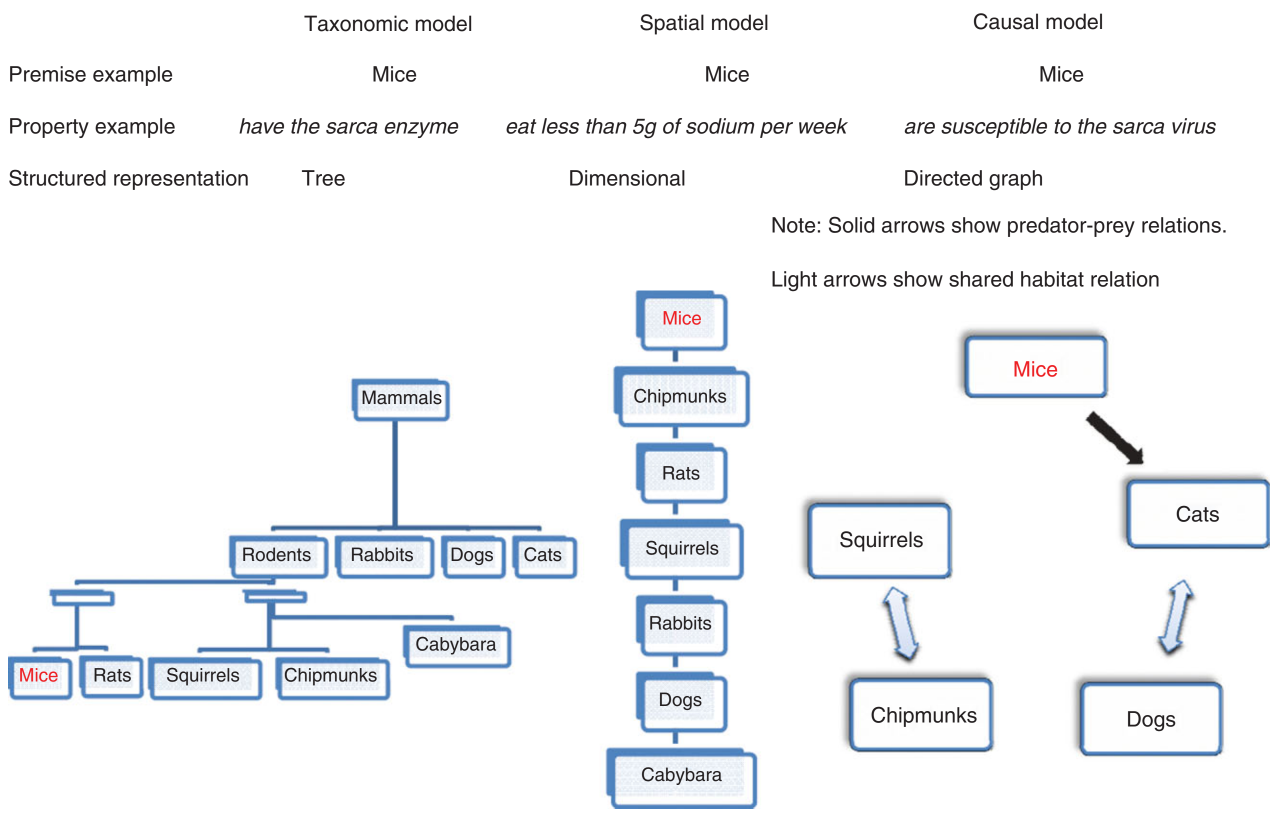

FIGURE 1 | Examples of tree, dimensional and causal structured representations. (Used with permission from Kemp and Tenenbaum.[53])

 © 2010 John Wiley & Sons, Ltd.

cases of induction where premises and conclusion categories are seen as similar because they fulfill similar causal *roles* (e.g., hawks and tigers are similar because they both take the role of predator). Perhaps, the main weakness of the structured statistical model though is a by-product of one of its principal strengths, namely its flexible application of different knowledge structures. The model says little about how people ensure that they retrieve the correct structured representation for a given problem. When told that mice have a certain disease, it seems likely that people would consider a range of possible routes for property generalization (taxonomic, predator–prey relations, ecological relations, etc.), each of which is associated with a different structured representation for generating priors. Exactly how a particular representation is selected (and others discarded) is not specified. This seems very similar to the problem of 'knowledge selection' in other cognitive domains such as category learning in knowledge-rich environments.[55] A first step in resolving this problem for inductive reasoning may be to adapt the knowledge selection approaches that have been successful in those domains.

Summary

This review shows that much of the competition between inductive theories involves how well they account for context- and knowledge-based effects (as summarized in Table 2). It is clear that this is a complex and difficult problem because it involves generating a theory of how knowledge is represented and then applied in specific contexts, which has proved to be an elusive goal in many areas of cognitive science.[56]

Nevertheless, it is clear that important progress has been made. A very positive step is that there are now a number of quantitative models that attempt to specify how background knowledge works together with similarity to determine inductive inference. The next step in theory development will involve a more detailed comparison of these quantitative models. At least three criteria need to be applied in adjudicating between models. First, the psychological plausibility of the model assumptions needs to be examined. Second, the tradeoff between model complexity and flexibility with quantitative fit needs to be evaluated. Finally, and perhaps most importantly, the models need to be compared in their ability to generate novel, testable predictions about how people do induction.

EMERGING DIRECTIONS IN INDUCTION RESEARCH

In this final section, we briefly review a number of emerging trends in induction research. The first of these areas involve forms of induction that are common in everyday life but have not been studied in the same depth as those reviewed so far. We then examine recent work on the relation between inductive and deductive reasoning, and how our understanding of this issue has been enriched through studies of the neural basis for induction.

Induction with Uncertain Categories

Although the phenomena summarized in Table 1 differ from one another in a number of important ways, they share one critical feature. All involve generalizing a property from a known or *certain* premise category. However, there may be many circumstances where inductive inferences are required about instances whose category membership is uncertain or ambiguous. Imagine you are hiking through a wilderness area and hear rustling in the scrub near your feet. Although you may know of many animals (e.g., snakes, small mammals, birds) that could produce the noise, you cannot be certain about the actual source. Nevertheless, you may wish to make a prompt prediction about some of the animal's features (e.g., how dangerous it is). Clearly, the various alternative categories have different conditional probabilities of being dangerous. The main psychological question is the extent to which people factor in these various alternatives into their feature predictions.

Normative (Bayesian) approaches[57] suggest that all of the alternatives should be considered. People should estimate the probabilities of the predicted property for each category alternative and then weight each according to the likelihood of the object being in that category. Early empirical work on this issue suggested that people do not reason this way. Experimental studies with artificial[58,59] and natural categories[60,61] found that people often ignored uncertainty about category membership, basing predictions solely on the category to which an object is most likely to belong (a heuristic termed 'single-category' reasoning).

Recent work has questioned the generality of this heuristic, demonstrating that whether or not people consider multiple category alternatives in inference depends on a range of factors.[62,63] These include experience in making predictions, the consequences of ignoring less likely categories and whether people are forced to make category judgments before making predictions.

An even more striking finding is that in some situations people may ignore category membership and base their predictions on a direct comparison between a test exemplar and other exemplars with

 © 2010 John Wiley & Sons, Ltd.

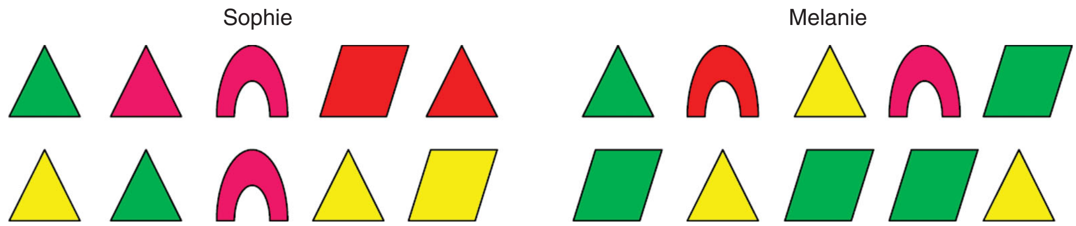

FIGURE 2 | Examples of stimuli used by Papadopoulos et al. to compare category-based reasoning and feature conjunction reasoning for a test item ('triangle') whose category membership is uncertain.[65]

similar features. This approach to reasoning, which has been termed 'feature conjunction', can be illustrated with reference to Figure 2. The stimuli are presented as categories of drawings done by different children. A novel test drawing with a single feature is then presented (e.g., a triangle). The participant has to classify this drawing ('who is most likely to have drawn it?') and make a prediction about its other feature ('what color is it most likely to be?'). Category membership is uncertain because although Sophie is the most likely artist, both children drew triangles. Category-based reasoning leads to the feature prediction of either 'green' (if both categories are used) or 'red' (if single-category reasoning is used). In contrast, feature conjunction ignores category bounds and compares objects that match the feature given in the test probe. Hence, you would examine the distribution of colors in all triangles, leading to the feature prediction 'yellow'. In several studies, an overwhelming majority preferred predictions based on feature conjunction over those based on either of the category-based approaches.[64,65] Hence, if object categorization is sufficiently uncertain, people may abandon using categories as a basis for feature induction, and instead use specific exemplars.

Discovering how people make inductive inferences in uncertain situations is important for a number of reasons. It expands the range of phenomena that can be used to constrain and develop theories of induction. It strengthens the links between research on inductive reasoning and other domains of cognitive science like decision making.[66] Finally, induction with ambiguous categories is closely related to another important question, namely how people go about making inferences about cross-classified objects.[67]

Relations Between Induction and Deduction

More broadly, inductive reasoning is a kind of reasoning, but not the only kind. The traditional alternative to induction is deduction, which is linked to making valid inferences that are 100% certain given a set of premises, and do not depend on, or ideally even use, other background knowledge. Hence, at first glance, deduction seems very different from induction, which is probabilistic and knowledge-rich. But where do you draw the line between induction and deduction, and how are they related? Heit has distinguished between two different approaches, the problem view and the process view.[68] According to the problem view, induction and deduction refer to different kinds of reasoning problems that people solve. For example, deductive problems could be defined as arguments that are logically valid according to the rules of a well-specified logic, and other arguments could be referred to as inductive problems. Alternately, deductive problems could be defined as those with 100% (or perhaps 99.9%) likely conclusions, inductively strong problems could be those with very highly probable conclusions (perhaps 75–99.9%), and other problems could be considered weak.[69]

Unlike the problem view, the process view is concerned with cognitive processes. The question of interest is what processes underlie induction and deduction, and whether these are the same or different. Some researchers have suggested that induction and deduction depend on the same cognitive processes. This approach will be referred to as the one-process view. Several influential research programs embody the one-process view, by applying a common framework to both inductive and deductive problems. For example, the similarity–coverage model and FBIM would, without additional assumptions, account for some deductive reasoning phenomena. Similarly, Johnson-Laird[70] explained how mental model theory, typically applied to problems of deduction, can also be applied to problems of induction.

In contrast, according to two-process accounts,[21,71] both heuristic and analytic processes contribute to reasoning. Both induction and deduction could be influenced by these two processes, but in different proportions. Induction judgments would be particularly influenced by quick heuristic processes that tap into associative information about similarity and background knowledge but which do not determine the logical validity of an argument.

 © 2010 John Wiley & Sons, Ltd.

In contrast, deduction judgments would be more heavily influenced by slower analytic processes that encompass more deliberative, and typically more accurate, reasoning. Two-process accounts have provided an explanatory framework for many results (e.g., content effects, individual differences, effects of time pressure), but in general have not been implemented computationally.

A recent line of work[72–74] has directly pitted one- and two-process accounts against each other, by implementing the accounts as computational models and applying them to experimental data. The general method, adapted from Rips,[75] is to take a set of arguments and have one group judge whether the arguments are deductively valid and another group judge whether they are inductively strong. To the extent that different variables affect deduction versus induction judgments, there will be evidence of multiple cognitive processes affecting deduction and induction, rather than a single, shared process. Indeed, Heit and Rotello found that logical validity itself has a greater effect on deduction judgments, whereas more superficial variables such as similarity and the length of an argument (the similarity and numerosity phenomena in Table 1) have a greater effect on induction judgments. They implemented a two-dimensional model of reasoning (Figure 3) that successfully accounted for these results by assuming that deduction and induction rely on differing proportions of two kinds of underlying information, which could be the outputs of analytic and heuristic processes.

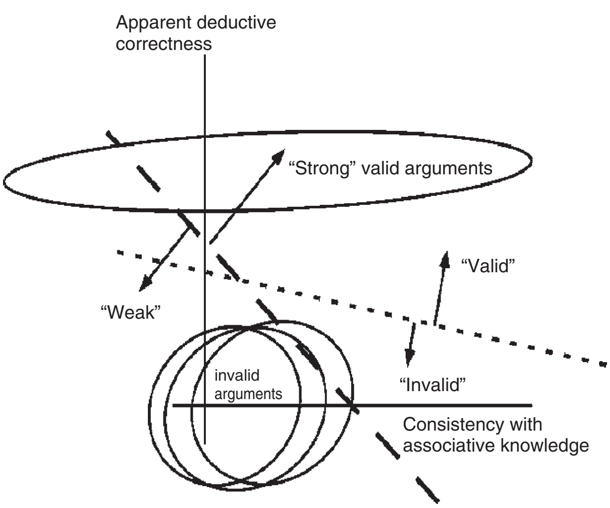

FIGURE 3 | Two-dimensional model of reasoning, showing arguments varying in apparent deductive correctness (y-axis) and consistency with associative knowledge (x-axis). The dotted line shows the decision boundary for judging whether an argument is deductively valid or invalid. The dashed line shows the decision boundary for judging whether an argument is inductively strong or weak. (Used with permission from Heit and Rotello[73]).

Neuroscience of Induction

Recent technological advances such as functional magnetic resonance imaging (fMRI) have spurred research on the neuroscience of reasoning, with the potential to address important theoretical questions in reasoning research. Researchers use fMRI to measure the hemodynamic response related to neural activity in the brain—sampling every few seconds from different positions and producing detailed, three-dimensional images of brain structure and activation.[76] Although it seems plausible that induction takes place throughout the brain, given the probabilistic and context-dependent nature of induction, some research has aimed to localize induction (and deduction) in regard to particular brain regions.

Several studies[77–80] have compared deductive and inductive reasoning tasks, in an effort to compare one- and two-process accounts of reasoning. To the extent that qualitatively different patterns of brain activity are observed for deduction versus induction, holding everything else equal between experimental conditions, dual-process accounts will be supported over single-process alternatives.[76] Indeed, these four studies found somewhat different patterns of brain activation for deduction versus induction. For example, three studies[77–79] found increased activation for induction, relative to deduction, in left frontal cortex, although in fairly different regions at a finer level. Overall, these results help make the case for two-process accounts over one-process accounts, although there are still unanswered questions such as why different brain regions are activated in different induction tasks.

Although cognitive scientists have typically used brain imaging techniques to compare theories[76] or identify the functions of specific brain regions,[81] another interesting approach is to use measures of brain activity as a kind of data to make behavioral predictions. Weber and Osherson[82] developed a simple mathematical model, related to SimProb,[49] for predicting inductive arguments from people's similarity ratings. Although the model performed well, they found slightly better performance when the model used patterns of brain activity rather than overt similarity ratings to predict inductive strength. In particular, subjects underwent fMRI scanning while identifying pictures of mammals such as dolphins and bears. Different mammals were associated with somewhat different patterns of activation in particular regions of the left occipital cortex and left ventral

© 2010 John Wiley & Sons, Ltd.

cortex. The degree of overlap in these patterns, between pairs of mammals, was used as a measure of similarity to successfully predict inductive judgments. This kind of finding suggests that neuroscientific approaches will play an increasingly important role in future studies of induction.

CONCLUSION

The review has highlighted two important trends in induction research. On the one hand, there is increasing attention being given to the development of more clearly articulated formal models that have sufficient scope to account for many, if not most, of the touchstone results in Table 1. On the other hand, induction researchers are increasingly turning their attention to a broader range of reasoning phenomena, examining the links between induction and other cognitive activities like decision making,[62] deductive reasoning,[73] and memory.[6,7,83] It seems likely that in the future these two trends will overlap. For example, some researchers have begun to examine the conceptual and empirical links between inductive judgments and recognition memory for a common stimulus set.[7,83] A positive outcome of this work may be the development of clearly articulated, computational models that explain how knowledge is generalized from known to unknown cases across a range of cognitive activities.

ACKNOWLEDGEMENTS

This work was supported by Australian Research Council Discovery Grant DP0663963 to the first two authors and collaborative National Science Foundation Grant BCS-0616979 to the second author. The authors would like to thank Helen Paton for her assistance in the preparation of the manuscript.

REFERENCES

1. Hayes BK. The development of inductive reasoning. In: Feeney A, Heit E, eds. *Inductive Reasoning: Experimental, Developmental and Computational Approaches*. Cambridge University Press; 2007, 25–54.
2. Heit E. Properties of inductive reasoning. *Psychon Bull Rev* 2000, 7:569–592.
3. Heit E. Models of inductive reasoning. In: Sun R, ed. *Cambridge Handbook of Computational Psychology*. New York, NY: Cambridge University Press; 2008, 322–338.
4. Osherson DN, Smith EE, Wilkie O, Lopez A. Category-based induction. *Psychol Rev* 1990, 97:185–200.
5. Keates J, Graham SA. Category markers or attributes: why do labels guide infants' inductive inferences?. *Psychol Sci* 2008, 19:1287–1293.
6. Sloutsky VM, Fisher AV. Induction and categorization in young children: a similarity-based model. *J Exp Psychol Gen* 2004, 133:166–188.
7. Heit E, Hayes BK. Relations between categorization, induction, recognition and similarity. *J Exp Psychol Gen* 2005, 134:596–605.
8. Hayes BK, McKinnon R, Sweller N. The development of category-based induction: re-thinking conclusions from the induction then recognition paradigm. *Dev Psychol* 2008, 44:1430–1444.
9. Rips LJ. Inductive judgements about natural categories. *J Verbal Learn Verbal Behav* 1975, 14:665–681.
10. Carey S. *Conceptual Change in Childhood*. Cambridge, MA Bradford Books; 1985.
11. Lopez A, Gelman SA, Gutheil G, Smith EE. The development of category-based induction. *Child Dev* 1992, 63:1070–1090.
12. Medin DL, Waxman S. Interpreting asymmetries of projection in children's inductive reasoning. In: Feeney A, Heit E, eds. *Inductive Reasoning: Experimental, Developmental and Computational Approaches*. New York, NY Cambridge University Press; 2007, 55–80.
13. Gelman SA, O'Reilly AW. Children's inductive inferences within superordinate categories: the role of language and category structure. *Child Dev* 1988, 59:876–887.
14. Hampton JA, Cannon I. Category-based induction: an effect of conclusion typicality. *Mem Cognit* 2004, 32:235–243.
15. Nisbett RE, Krantz DH, Jepson C, Kunda Z. The use of statistical heuristics in everyday inductive reasoning. *Psychol Rev* 1983, 90:339–363.
16. Gelman SA. The development of induction within natural kind and artefact categories. *Cogn Psychol* 1988, 20:65–95.
17. Coley JD, Hayes B, Lawson C, Moloney M. Knowledge, expectations and inductive reasoning within conceptual hierarchies. *Cognition* 2004, 90:217–253.

 © 2010 John Wiley & Sons, Ltd.

18. Nagel E. *Principles of the Theory of Probability*. Chicago, IL University of Chicago Press; 1939.
19. Heit E, Feeney A. Relations between premise similarity and inductive strength. *Psychon Bull Review* 2005, 12:340–344.
20. Lopez A. The diversity principle in the testing of arguments. *Mem Cognit* 1995, 23:374–382.
21. Feeney A. How many processes underlie category-based induction? Effects of conclusion specificity and cognitive ability. *Mem Cognit* 2007, 35:1830–1839.
22. Lo Y, Sides A, Rozelle J, Osherson D. Evidential diversity and premise probability in young children's inductive judgment. *Cogn Sci* 2002, 26:181–206.
23. Heit E, Hahn U, Feeney A. Defending diversity. In: Ahn W, Goldstone RL, Love BC, Markman AB, Wolfe P, eds. *Categorisation Inside and Outside the Laboratory: Essays in Honor of Douglas L. Medin*. Washington, DC American Psychological Association; 2005, 87–99.
24. Sloman SA. Feature-based induction. *Cogn Psychol* 1993, 25:231–280.
25. Sloman SA. Categorical inference is not a tree: the myth of inheritance hierarchies. *Cogn Psychol* 1998, 35:1–33.
26. Rhodes M, Brickman D, Gelman SA. Sample diversity and premise typicality in inductive reasoning: evidence for developmental change. *Cognition* 2008, 108:543–556.
27. Rhodes M, Gelman SA, Brickman D. Children's attention to sample composition in learning, teaching, and discovery. *Dev Sci* 2009, 1:1–9.
28. Heit E, Hahn U. Diversity-based reasoning in children. *Cogn Psychol* 2001, 43:243–273.
29. Bailenson JN, Shum MS, Atran S, Medin DL, Coley JD. A bird's eye view: biological categorization and reasoning within and across cultures. *Cognition* 2002, 84:1–53.
30. Lopez A, Atran S, Coley JD, Medin DL, Smith EE. The tree of life: universal and cultural features of folk-biological taxonomies and inductions. *Cogn Psychol* 1997, 32:251–295.
31. Atran S, Medin D. *The Native Mind and the Cultural Construction of Nature*. Cambridge MIT Press; 2008.
32. Bang M, Medin D, Atran S. Cultural mosaics and mental models of nature. *Proc Natl Acad Sci U S A* 2007, 104:13868–13874.
33. Shafto P, Coley JD, Vitkin A. Availability in category-based induction. In: Feeney A, Heit E, eds. *Inductive Reasoning: Experimental, Developmental, and Computational Approaches*. New York, NY Cambridge University Press; 2007, 114–136.
34. Proffitt JB, Coley JD, Medin DL. Expertise and category-based induction. *J Exp Psychol Learn Mem Cogn* 2000, 26:811–828.
35. Shafto P, Coley JD, Baldwin D. Effects of time pressure on context-sensitive property induction. *Psychon Bull Rev* 2007, 14:890–894.
36. Heit E, Rubinstein J. Similarity and property effects in inductive reasoning. *J Exp Psychol Learn Mem Cogn* 1994, 20:411–422.
37. Ross BH, Murphy GL. Food for thought: cross-classification and category organization in a complex real-world domain. *Cogn Psychol* 1999, 38:495–553.
38. Smith EE, Shafir E, Osherson D. Similarity, plausibility, and judgements of probability. *Cognition* 1993, 49:67–96.
39. Nguyen SP, Murphy GL. An apple is more than just a fruit: cross-classification in children's concepts. *Child Dev* 2003, 74:1783–1806.
40. Medin DL, Coley JD, Storms G, Hayes BK. A relevance theory of induction. *Psychon Bull Rev* 2003, 10:517–532.
41. Feeney A, Shafto P, Dunning D. Who is susceptible to conjunction fallacies in category-based induction? *Psychon Bull Rev* 2007, 14:884–889.
42. Lassaline ME. Structural alignment in induction and similarity. *J Exp Psychol Learn Mem Cogn* 1996, 22:754–770.
43. Rehder B. When similarity and causality compete in category-based property generalization. *Mem Cognit* 2006, 34:3–16.
44. Hayes BK, Thompson SP. Causal relation and feature similarity in children's inductive reasoning. *J Exp Psychol Gen* 2007, 136:470–484.
45. Sobel DM, Kirkham NZ. Blickets and babies: The development of causal reasoning in toddlers and infants. *Dev Psychol* 2006, 42:1103–1115.
46. Rehder B. Causal-based property generalization. *Cogn Sci* 2009, 33:301–343.
47. Rehder B, Burnett RC. Feature inference and the causal structure of categories. *Cogn Psychol* 2005, 50:364–314.
48. Crisp AK, Feeney A. Causal conjunction fallacies: the roles of causal strength and mental resources. *Q J Exp Psychol* 2009, 63:2320–2337.
49. Blok SV, Medin DL, Osherson DN. Induction as conditional probability judgement. *Mem Cognit* 2007, 35:1353–1364.
50. Heit E. A Bayesian analysis of some forms of inductive reasoning. In: Oaksford M, Chater N, eds. *Rational Models of Cognition*. Oxford, UK: Oxford University Press; 1998, 248–274.
51. Tenenbaum JB, Griffiths TL. Generalization, similarity and Bayesian inference. *Behav Brain Sci* 2001, 24:629–640.
52. Shafto P, Kemp C, Bonawitz EB, Coley JD. Inductive reasoning about causally transmitted properties. *Cognition* 2008, 109:175–192.

 © 2010 John Wiley & Sons, Ltd.

53. Kemp C, Tenenbaum JB. Structured statistical models of inductive reasoning. *Psychol Rev* 2009, 116:20–58.
54. Shafto P, Coley JD. Development of categorization and reasoning in the natural world: novices to experts, naive similarity to ecological knowledge. *J Exp Psychol Learn Mem Cogn* 2003, 29:641–649.
55. Heit E, Bott L. Knowledge selection in category learning. In: Medin D, ed. *Psychology of Learning and Motivation*. Vol. 39. San Diego, CA: Academic Press; 2000, 163–199.
56. Rogers TT, McClelland JL. *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge MIT Press; 2004.
57. Anderson JR. The adaptive nature of human categorization. *Psychol Rev* 1991, 98:409–429.
58. Murphy GL, Ross BH. Prediction from uncertain categorization. *Cogn Psychol* 1994, 27:148–193.
59. Murphy GL, Ross BH. The two faces of typicality in category-based induction. *Cognition* 2005, 95:175–200.
60. Malt BC, Ross BH, Murphy GL. Predicting features of members of natural categories when categorization is uncertain. *J Exp Psychol Learn Mem Cogn* 1995, 21:646–661.
61. Ross BH, Murphy GL. Category-based predictions: influence of uncertainty and feature associations. *J Exp Psychol Learn Mem Cogn* 1996, 22:736–753.
62. Hayes BK, Newell BR. Induction with uncertain categories: when do people consider the category alternatives? *Mem Cognit* 2009, 37:730–743.
63. Hayes BK, Chen TJ. Clinical expertise and reasoning with uncertain categories. *Psychon Bull Rev* 2008, 15:1002–1007.
64. Hayes BK, Ruthven C, Newell BR. Inferring properties when categorization is uncertain: a feature-conjunction account. In: McNamara D, Tafton G, eds. *Proceedings of the 29th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum; 2007, 309–214.
65. Papadopoulos C, Hayes BK, Newell B. Induction with uncertain categories. In: Taatgen N, van Rijn H, eds. *Proceedings of the 30th Meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum; 2009, 83–88.
66. Medin DL, Bazerman MH. Broadening behavioral decision research: multiple levels of cognitive processing. *Psychon Bull Rev* 1999, 6:533–546.
67. Murphy GL, Ross BH. Induction with cross-classified categories. *Mem Cognit* 1999, 27:1024–1041.
68. Heit E. What is induction and why study it? In: Feeney A, Heit E, eds. *Inductive Reasoning*. Cambridge Cambridge University Press; 2007, 1–24.
69. Skyrms B. *Choice and Chance: An Introduction to Inductive Logic*. 4th ed. Belmont, CA Wadsworth; 2000.
70. Johnson-Laird PN. Mental models and probabilistic thinking. *Cognition* 1994, 50:189–209.
71. Evans JSBT. Dual-processing accounts of reasoning, judgment, and social cognition. *Annu Rev Psychol* 2008, 59:255–278.
72. Heit E, Rotello CM. Are there two kinds of reasoning? In: Bara B, Barsalou L, Eucciarelli M, eds. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum; 2005.
73. Heit E, Rotello CM. Modeling two kinds of reasoning. In: Sloutsky V, Love B, McCrae K, eds. *Proceedings of the Thirtieth Annual Meeting of the Cognitive Science Society*. Austin, TX Cognitive Science Society; 2008, 1831–1836.
74. Rotello CM, Heit E. Modeling the effects of argument length and validity on inductive and deductive reasoning. *J Exp Psychol Learn Mem Cogn* 2009, 35:1317–1330.
75. Rips LJ. Two kinds of reasoning. *Psychol Sci* 2001, 12:129–134.
76. Henson RN. What can functional imaging tell the experimental psychologist? *Q J Exp Psychol A* 2005, 58:193–233.
77. Goel V, Dolan RJ. Differential involvement of left prefrontal cortex in inductive and deductive reasoning. *Cognition* 2004, 93:B109–B121.
78. Goel V, Gold B, Kapur S, Houle S. The seats of reason: a localization study of deductive and inductive reasoning using PET (O15) blood flow technique. *Neuroreport* 1997, 8:1305–1310.
79. Osherson D, Perani D, Cappa S, Schnur T, Grassi F, et al. Distinct brain loci in deductive versus probabilistic reasoning. *Neuropsychologia* 1998, 36:369–376.
80. Parsons LM, Osherson D. New evidence for distinct right and left brain systems for deductive versus probabilistic reasoning. *Cereb Cortex* 2001, 11:954–965.
81. Poldrack RA. Can cognitive processes be inferred from neuroimaging data? *Trends Cogn Sci* 2006, 10:59–63.
82. Weber M, Osherson D. Similarity and induction. *Eur Rev Philos Psychol*. In press.
83. Heit E, Hayes B. Predicting reasoning from visual memory. In: Sloutsky V, Love B, McCrae K, eds. *Proceedings of the 29th Meeting of the Cognitive Science Society*. Hillsdale, NJ Erlbaum; 2008, 83–88.

FURTHER READING

Feeney A, Heit E, eds. *Inductive Reasoning: Experimental, Developmental, and Computational Approaches* New York, NY: Cambridge University Press; 2007.

 © 2010 John Wiley & Sons, Ltd.

Goodman N. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press; 1955.
Hempel CG. *Philosophy of Natural Science*. Englewood Cliffs, NJ: Prentice Hall; 1966.
Holland JH, Holyoak KJ, Nisbett RE, Thagard PR. *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press; 1986.
Mill JS. *A System of Logic*. New York, NY: Harper; 1874.
Sloman SA, Lagnado DA. The problem of induction. In: Holyoak K, Morrison R, eds. *Cambridge Handbook of Thinking and Reasoning*. New York, NY: Cambridge University Press; 2005.

 © 2010 John Wiley & Sons, Ltd.