

Induction and knowledge-what

Peter Gärdenfors¹ · Andreas Stephens¹

Received: 21 April 2017 / Accepted: 23 November 2017 / Published online: 13 December 2017

© The Author(s) 2017. This article is an open access publication

Abstract Within analytic philosophy, induction has been seen as a problem concerning inferences that have been analysed as relations between sentences. In this article, we argue that induction does not primarily concern relations between sentences, but between *properties* and *categories*. We outline a new approach to induction that is based on two theses. The first thesis is epistemological. We submit that there is not only knowledge-how and knowledge-that, but also *knowledge-what*. Knowledge-what concerns relations between properties and categories and we argue that it cannot be reduced to knowledge-that. We support the partition of knowledge by mapping it onto the long-term memory systems: procedural, semantic and episodic memory. The second thesis is that the role of inductive reasoning is to generate knowledge-what. We use conceptual spaces to model knowledge-what and the relations between properties and categories involved in induction.

Keywords Inductive inferences · Knowledge-what · Semantic memory · Conceptual knowledge · Naturalistic epistemology

1 Introduction

One of the most impressive features of human cognitive processing is our ability to perform *inductive inferences*. We generalise from a very limited number of observations, sometimes with overwhelming confidence. A central problem in philosophy of science concerns how the mechanism of inductive reasoning can be described and motivated.

✉ Peter Gärdenfors
peter.gardenfors@lucs.lu.se

Andreas Stephens
andreas.stephens@gmail.com

¹ Department of Philosophy, Lund University, Lund, Sweden

We do not perform inductive inferences in an arbitrary manner. Peirce notes that there are certain forms of constraints that delimit the vast class of possible inferences. As he puts it:

Nature is a far vaster and less clearly arranged repertory of facts than a census report; and if men had not come to it with special aptitudes for guessing right, it may well be doubted whether in the ten or twenty thousand years that they may have existed their greatest mind would have attained the amount of knowledge which is actually possessed by the lowest idiot. But, in point of fact, not man merely, but all animals derive by inheritance (presumably by natural selection) two classes of ideas which adapt them to their environment. In the first place, they all have from birth some notions, however crude and concrete, of force, matter, space, and time; and, in the next place, they have some notion of what sort of objects their fellow-beings are, and how they will act on given occasions. (Peirce 1955, pp. 214–5)

Here, Peirce hints at an *evolutionary* explanation of why “the human intellect is peculiarly adapted to the comprehension of the laws and facts of nature” (Peirce 1955, p. 213).

Within analytic philosophy, induction has been seen as a problem concerning inferences that have been analysed as relations between sentences. Inductive inferences were important for the logical positivists, being a cardinal component in their verificationist program (see, e.g., Carnap 1950; Hempel 1965; Rosenberg 2000; Ladyman 2002; Creath 2014; Vickers 2014). However, it soon became apparent that their logical approach resulted in paradoxes. The most well-known are Hempel’s (1965) ‘paradox of confirmation’ and Goodman’s (1983) ‘new riddle of induction’. If we use logical relations alone to determine which inductions are valid, the fact that all predicates are treated on a par induces *symmetries* which are not preserved by our intuitions concerning which inductive inferences are permissible: ‘Raven’ in Hempel’s paradox is treated on a par with ‘non-raven’, ‘green’ in Goodman’s with ‘grue’, etc. What is needed is a non-logical way of distinguishing the predicates that may be used in inductive inferences from those that may not.

In this article, our diagnosis of why the paradoxes have emerged in the traditional treatment is that induction does not primarily concern relations between sentences, but between *properties* and *categories*. We outline a new approach to induction that is based on two theses. The first one is epistemological. We argue that there is not only knowledge-how and knowledge-that, but also *knowledge-what*. Knowledge-what concerns relations between properties and categories and we argue that it cannot be reduced to knowledge-that. We motivate our approach by giving it a naturalistic grounding in cognitive neuroscience.

The second thesis is that the role of induction is to generate knowledge-what. This entails that we find much of the earlier discussion of induction misguided since it has focused on induction as generalisations generating knowledge-that. In this context, it should be noted that there are two meanings of ‘generalisation’ in the literature. One is logical, relating to the connection between sentences describing individual instances and universal sentences covering the individual sentences. The other, also called ‘stimulus generalisation’, is psychological and concerns the relations between reactions

to a particular stimulus and a class of *similar* stimuli. We argue that human inductive inference is more related to the psychological notion of generalisation.

A central question then is how knowledge-what can be modelled. Here we build on the theory of conceptual spaces proposed by Gärdenfors (1990, 2000, 2014). In this theory, knowledge is organised into domains modelled as spatial structures. Properties are analysed as (convex) regions within such domains and categories as complexes of regions from different domains. There are several dimensional theories of categorisation, but the unique property of the theory of conceptual spaces is its strong reliance on geometric structures.

Before we present our analysis of knowledge-what and its relation to induction, we give, in section 2, a brief account of why induction has been seen as generating knowledge-that, and then outline our own cognitivist and naturalistic stance. We present arguments for dividing knowledge into knowledge-how, knowledge-what and knowledge-that in section 3. In section 4, we map this tripartition of knowledge onto three kinds of long-term memory – procedural, semantic and episodic – thereby connecting our account of knowledge to cognitive neuroscience. Then in section 5 we introduce conceptual spaces as a tool for modelling knowledge-what in the form of relations between categories and properties. Finally, in section 6 we argue that induction concerns methods for generating knowledge-what.

2 Two approaches to induction

In this section, we sketch an account of why inductive inferences have been seen as relations between sentences, and then present our alternative naturalistic approach. We derive the approaches from the underlying views of what constitute knowledge.

2.1 Induction from the perspective of language

Historically, the empiricist turn of the seventeenth century raised an interest in inductive inferences, although it remained uncertain how induction should be justified since it lacked the logical rigor of deduction. Hume (1988) argued that it is impossible to justify inductive inferences, although he acknowledged *habit* as an inevitable part of human reasoning. Other issues included pinpointing which evidence, and what amount, was enough for valid inductive inferences as well as finding methods that could separate good inferences from bad ones (see, e.g., Mill 1843; Vickers 2014).

In the mainstream debate within analytic philosophy, a major distinction has been that between *knowledge-how* and *knowledge-that* (Ryle 1949). It has been a tacit assumption that induction does not concern knowledge-how. As part of the linguistic turn of analytic philosophy, there was a preference for analysing inferences, including induction, as relations between sentences. Hence, it was concluded that if induction is an epistemic process, it must deal with knowledge-that, since knowledge-that is propositional and can be expressed in sentences.

For the logical positivists, the basic objects of study were sentences in some more or less regimented language. Ideally, the language was a version of first-order logic where the atomic predicates represent observational properties. These observational predicates were taken as primitive, unanalysable notions. The main tool used when studying the

linguistic expressions was logical analysis. In its purest form, logical positivism allowed only this tool. A consequence of this methodology was that all observational predicates were treated in the same way since there were no logical reasons to differentiate between them. For example, Carnap (1950, sec. 18B) required that the primitive predicates of a language be logically independent of each other.

In this tradition, particular observations are used as evidence for inductive generalisations or predictions (Carnap 1950, 1971; Hempel 1965). When connections are found within the registered observations, inductive generalisations can be made, which then can be confirmed by additional observations. So if observations of objects $O_1, O_2, O_3 \dots$ all are C , the generalisation that *all* O are C can be made. The inference thus concerns a relation between individual and universal sentences. The evidence from the premises gives stronger or weaker support for the conclusion. Since this inductive process does not have the same logical rigor as a deductive process, the methodology of induction requires that supporting evidence preferably should come in large numbers, come from several different contexts and have no negative cases (Hempel 1965).

One point that has been downplayed in the debate, however, is that not all universal sentences can function as conclusions in inductive inferences. Ever since Aristotle's classic "All men are mortal", inductive inferences only involve universal sentences that are *generics*, that is, express relations between categories and properties. A non-generic universal sentence such as "All persons in this room are Swedish" would not be acceptable as an inductive inference, even when perfectly supported by the given evidence. This is so since such 'accidental generalisations' do not support counterfactuals of the form "if a person came into the room he or she would be a Swede". So, even though the logical form of a law-like sentence is the same as that of an accidental universal sentence, we point to the connection between law-like sentences and generics. In the literature there has been attempts to distinguish 'law-like' (nomologic, nomothetic) generalisations from 'accidental' generalisations (see, e.g., Goodman 1983; Hempel 1965), and early steps to break the logical emphasis were taken by for example Dretske (1977), Tooley (1977) and Armstrong (1978, 1983) who focused on laws as relations of non-logical *necessitation* between universals (see also Carroll 2016).

2.2 Induction from a naturalistic perspective

Our alternative to the traditional propositional or sentential approach is *cognitivist* and *naturalistic*. We thus highlight that inductive inferences are possible since the world has moulded our cognitive faculties through evolution (Quine 1969b; Lorenz 1977; Gärdenfors 1990, 2000; Humphrey 1992; Kornblith 1993). We are cognitively imprinted to discover, recognise and categorise certain patterns in the world – otherwise our generalisations and predictions would be miraculous (Dennett 1991; see also Johansson 1998). So, in contrast to the discussion mentioned above on what constitute laws, our focus concerning inductive inferences is on the kind of *knowledge* that is involved.

We show in section 6 that psychological research on sensory and perceptual generalisations involved in learning concern properties and categories, rather than propositions. From this perspective, inductive inferences can be seen as natural processes in cognitive systems, rather than in language, that occur when an agent categorises its sensory input and then makes generalisations or predictions using its understanding of these categories.

Our approach to induction is naturalistic in the sense that we look to science for relevant input instead of relying on intuitions and language.¹ Methodologically, we endorse a form of ‘cooperative naturalism’, according to which relevant scientific findings always should be taken into consideration since they provide our best explanations (Rysiew 2016).²

In fact, there are a number of interconnected scientific practices inquiring into induction and knowledge, on many different levels of explanation and from different perspectives. Three research areas come fairly close to the traditional epistemological outlook to induction, namely, cognitive neuroscience, cognitive ethology, and cognitive psychology. In section 4 we single out and use cognitive neuroscience as a foundation for our partitioning of knowledge types and in section 6 we turn to cognitive psychology for experimental evidential input concerning inductive reasoning.

3 Knowledge-how, knowledge-what and knowledge-that

3.1 The contemporary debate

Ryle (1949) provides some influential arguments for upholding the distinction between ‘knowing-how’ and ‘knowing-that’. He argues that knowing-that is to *possess* knowledge whereas knowing-how is to be *intelligent*. Knowledge-that thus concerns relations between agents and true propositions, whereas knowledge-how instead concerns *abilities, dispositions* and *actions* of the agent.

However, not everyone agrees that there is a relevant distinction to be made (see, e.g., Stanley and Williamson 2001; Schaffer 2007; Stanley 2011). In particular, Stanley and Williamson (2001) and Stanley (2011) question the distinction and instead argue that knowledge-how is a form of knowledge-that. In the literature, this position is called *intellectualism*, in contrast to *anti-intellectualism* as exemplified by Ryle (1949), and for example Stanley (2011) claims that knowledge-how can be analysed as a state with propositional content.

In support of intellectualism, various examples like the following situation have been presented and discussed: “Suppose there is a certain complex ski manoeuvre, which only the most physically gifted of athletes can perform. A ski instructor might know how to do that manoeuvre, without being able to perform it herself.” (Stanley 2011, p. 128).³ The ski instructor is thought to know the relevant facts and propositions (knowledge-that) concerning the manoeuvre, which then can be used to ‘direct’ her or someone else’s actions. According to Stanley: “[... T]he acquisition of a skill is due to the learning of a fact [which] explains why certain acts constitute exercises of skill, rather than reflex. A particular action [...] is a skilled action, rather than a reflex, because it is guided by knowledge [...]” (Stanley 2011, p. 130).

¹ It should be pointed out that we consider philosophical questions important in their own right. Our point is that induction is not just a philosophical problem.

² Alternatives to our position can be found in for example ‘replacement naturalistic’ theories, where Quine (1969a) offers the most well-known account. Following a traditional understanding of Quine’s position, epistemology should “simply fall [...] into place as a chapter of psychology and hence of natural science.” (Quine 1969a, p. 82). Yet another alternative position is found in ‘substantial naturalism’, according to which epistemological questions ought to be re-formulated in more exact scientific terminology (Rysiew 2016).

³ Stanley attributes this example to p.c. with Jeff King.

The intellectualist position is, in our view, questionable since it fits only some aspects of highly technical skills and especially since it underestimates the importance and amount of non-conscious processes involved in intentional actions – even though there has been intellectualist attempts to better account for such aspects (see, e.g., Stanley and Krakauer 2013; Pavese 2015a, b). We agree that having propositional or theoretical knowledge (of true propositions), or receiving instructions (‘knowledge of the way’) that we should position and move our body in a particular manner might help us try to consciously improve our technique. Nevertheless, it is ultimately practical knowledge through repetitive training that eventually lets us *know how* to actually perform the action – it is only by going out on the slopes that we can learn how to ski. A myriad of non-declarative and non-conscious processes make up our motor-, perceptual- and cognitive abilities, which are required for us to know how to perform an action.⁴

3.2 Knowledge-what as knowledge of categories

Fantl proposes a more promising extension of the traditional dichotomy between knowledge-how and knowledge-that: “There’s the kind of knowledge you have when it is truly said of you that you know a person—say, your best friend.” (Fantl 2016). In our opinion, Fantl’s knowledge of ‘acquaintance’ is a special case of a third type of knowledge. We want to single out the ability to *categorise*, in particular to know *the relation between categories and properties* as a special form of knowledge, which we call *knowledge-what*.⁵

Not all relations between categories and properties are, however, relevant for induction. To make our use of the term knowledge-what more precise, three types of information about categories need to be separated: Defining properties, characteristic properties and accidental facts (Keil and Batterman 1984).⁶ Here we use these terms in the following way: *Defining* properties of a category refer to information that pertains to the meaning of the word for the category. *Characteristic* properties refer to general knowledge about the category, that is, properties that generally hold of the category (exceptions may be possible).⁷ In the case when characteristic properties are formulated in sentences, the distinction between defining and characteristic corresponds to the distinction between definitional and law-like sentences that has been made within

⁴ Another discussion of knowledge, which has received much less attention than knowledge-that and knowledge-how, is captured by the general formula *knowledge-wh*. This formula refers to the kind of knowledge involved when answering questions about who, when, where, why, whether, and what. If we consider knowledge-what, the examples that have been presented in the literature all concern singular facts rather than something general or categorical. Consequently, trying to identify the type of knowledge generated by induction by analysing answers to non-generic wh-questions does not seem to be a fruitful strategy. The intellectualist tradition claims that all forms of knowledge-wh, just as knowledge-how, reduces to declarative knowledge-that (Hintikka 1975; Lewis 1982; Boër and Lycan 1986; Higginbotham 1996; Stanley and Williamson 2001). We instead want to argue that knowledge-what is not directly connected to language but instead to properties and categories.

⁵ It should be noted that our use of ‘knowledge-what’ is not intended to cover all everyday uses of the term such as in “I know what time it is”.

⁶ This is sometimes referred to as the dictionary-encyclopaedia distinction, but this is a misnomer since dictionaries frequently use characteristic features in their definitions.

⁷ Among semanticists, it has been discussed how the borderline between defining and characteristic knowledge should be drawn, but this problem is not crucial for our arguments.

philosophy of science (Hempel 1965; Carroll 2016). *Accidental facts* contain information about particular instances of a category. We illustrate these three types of information with an example concerning the category ‘spiders’ web’:

- *Defining*: Spiders’ webs are made from a protein fibre extruded from the spider’s body.
- *Characteristic*: Spiders’ webs are used for catching insects that provide food for the spiders.
- *Accidental*: Spiders’ webs are abundant in my cellar.

Our take on knowledge-what is that it concerns defining and characteristic knowledge, while knowledge-that concerns facts – accidental facts as well as facts of the type $2 + 2 = 4$. Our central thesis (to be discussed in section 6) is that, as a special case of knowledge-what, inductive inferences result in knowledge about characteristic properties. We thus heed the anti-intellectualist distinction between knowledge-how and knowledge-that while adding knowledge-what as a third type of knowledge, which is central for processes of induction.

3.3 Knowledge-what is separate from knowledge-that

Even though knowledge-what is primarily non-linguistic, it can be expressed in language. We next present two arguments for why knowledge-what, even if formulated linguistically, should be separated from knowledge-that. The first one builds on the observation that it seems perfectly natural that the following two sentences can be accepted simultaneously:

- (1) Spiders have eight legs.
- (2) My neighbour’s spider has only seven legs.

However, if (1) is expressed – as is standard in writings on induction – in the form of a universal sentence:

- (1′) All spiders have eight legs.

Then (1′) contradicts (2). There are two reasonable ways out of the contradiction:

- (a) Deny that my neighbour’s creature is a spider.
- (b) Deny that (1) expresses the same knowledge as (1′).

A reason against option (a) can be found in that what characterises spiders can be thought of as a ‘pattern’ of properties. Despite having only seven legs it is still a spider, since it has other ‘essential’ properties of spiders (definitional properties).⁸ In favour of option (b), it is worth highlighting that (1) expresses definitional properties, while (2)

⁸ We need not subscribe to full-blown essentialism. It is sufficient that certain properties of spiders are considered cognitively more important than others. See Gärdenfors (2000, sec. 4.2.2) for a defence of such a ‘cognitive’ form of essentialism.

expresses an accidental fact. Interpreting (1) as (1') and putting it together with (2) conflates the two different types of knowledge. As an alternative way out of the contradiction one may propose the following formulation of (1):

- (1") Spiders *characteristically* have eight legs.

Barring the problem of explaining the meaning of 'characteristically' in a non-circular way, we consider that this formulation supports our position that the knowledge expressed in (1) is of the definitional or characteristic form, that is, knowledge-what.

3.4 Generic sentences express knowledge-what

A second argument for maintaining the distinction between knowledge-what and knowledge-that shows up in natural language, albeit in an indirect way, as the distinction between the meaning of *generic* universals versus the meaning of *factual* universals. For example, generic universals such as "Blue whales eat plankton" and "A wrench is a tool for fastening nuts" are used to express some of the characteristic properties of 'whale' and 'wrench'. In contrast, factual universals, such as "Blue whales can be seen around the Cape of Good Hope" and "Wrenches are expensive in this shop" express facts about the world that are not part of the characteristic properties about the concepts. And sentence (1) above is indeed a generic.

It is interesting to note that the two types of universals behave in different ways linguistically, as pointed out by Lawler (1973):

- (3a) Blue whales eat plankton.
- (3b) A blue whale eats plankton.
- (4a) Blue whales can be seen around the Cape of Good Hope.
- (4b) *A blue whale can be seen around the Cape of Good Hope.⁹

(3a) describes a characteristic property of *blue whales*. It can be exchanged for the indefinite singular version in (3b). It expresses a relation between the concept blue whale and the property of feeding on plankton. In contrast, (4a) is a factual universal that says something factual about blue whales. A test for this is that it cannot be exchanged for the indefinite singular version in (4b) (Carlson 2009; Krifka 2012). Lawler notes that generic universals (which he calls non-descriptive generics) "[...] seem most natural in definitional sentences, or ones used somehow to identify the nature of the thing specified by the generic by means of properties peculiar to it; they are less acceptable when an accidental quality is predicated on them." (Lawler 1973, p. 112).

The upshot is that although a generic universal is a sentence, it expresses a different kind of knowledge than factual universals. Philosophers who have analysed generics have noted that there is no linguistic operator associated with these sentences and that negations of generics cannot be handled in the traditional logical way (Leslie 2008). The fact that sentences (3a) and (3b) express the same content in spite of their very different logical form is a further indication that generics form a special class of sentences. This conclusion is also supported by the fact that generics are acquired

⁹ As is standard in linguistics, the asterix marks that the sentence is not acceptable.

earlier by children than explicit universal sentences (Gelman 2003), which indicates that the information contained in generics is of a more fundamental type (see also Hollander et al. 2002). Leslie (2008, p. 21) writes: “Thus the inclination to generalize, though aided by language, does not depend on language but is, rather, an early developing, presumably innate, cognitive disposition.”

Our position is that induction does not concern relations between sentences and hence it is not a logical problem involving relations between sentences. We submit that the focus should be on how relations between categories and properties are supported.

In this section we have argued that knowledge-how, knowledge-what and knowledge-that all fill important separate epistemic roles. We thus propose a tripartite division of knowledge. In the next section we present results from cognitive neuroscience that further support such a tripartition.

4 Memory and knowledge

Without memory there is no knowledge. In this section, we take a cognitive neuroscientific perspective and present a different kind of support for our thesis that knowledge-what is a separate form of knowledge by mapping our partitioning of the three types of knowledge onto different kinds of long-term memory. We build on Tulving’s (1985) categorisation of long-term memory into three kinds: procedural, semantic and episodic memories. Tulving’s position has been very influential and is still pertinent in recent analyses although it has been partially reinterpreted (see, e.g., Fletcher et al. 1999; Binder and Desai 2011; Yee et al. 2014; Kim 2016; see also Gazzaniga et al. 2002; Aizawa and Gillett 2009). In this paper, we follow the presentation in Yee et al. (2014).

4.1 Mapping forms of knowledge onto forms of memory

The non-declarative *procedural* memory, which is beyond our conscious reach, handles an agent’s skill in performing a task. This kind of memory can be described as generated by an automatic process, where an agent learns and remembers *how* to do something. Learning is achieved through repetition or practice, and procedural memory can easily be associated with operant conditioning since it is possible to describe in terms of stimulus and response. Procedural memory is something humans share with many other animals (Tulving 2002).

Semantic memory allows for agents to actively cognise about categories, concepts and objects. It is thus with the aid of semantic memory that agents think about categories and their relations (see, e.g., Herrnstein 1990; Martin et al. 1996; Martin and Chao 2001; Binder and Desai 2011; Yee et al. 2014). Semantic memory is general and does not depend on specific references to experiences. This kind of memory is needed for handling the environment as efficiently as possible. In particular, semantic memory is crucial for mapping categories to actions. Like procedural memory, some aspects of semantic memory are most likely hardwired through evolution, for example fear reactions to snakes.

[C]ategorization is no saltation. It has turned up at every level of the animal kingdom where it has been competently sought. One reason for looking more

carefully at lower levels of categorization is that the continuity of cognitive processes linking humans and other animals is clear and undeniable here. And, as the evidence to be summarized suggests, it is probably at the upper end of this span that animal and human cognitive capacities diverge. (Herrnstein 1990, p. 138)

Humans share semantic memory with mammals and birds (Tulving 2002). Numerous findings support conceptual and categorical abilities in animals such as for example common squirrel monkeys (Thomas and Kerr 1976), rhesus monkeys (Spaet and Harlow 1943; Sands et al. 1982; Schrier and Brady 1987), chimpanzees (Nissen 1953), and pigeons (Vetter and Hearst 1968; Zeiler 1969; Cerella 1979).¹⁰

Episodic memory governs experienced knowledge that can be used in narratives. This kind of memory generates self-aware *remembrance* concerning single events such as they are experienced from a first person perspective (Tulving 1985, 2002).

Episodic memory makes it possible for humans to ‘time-travel’ in their minds. It allows us to remember individual events or episodes and the order in which they have occurred. Tulving (2002) claims that this form of memory is only found in humans. This position has, however, recently been challenged by researchers in animal cognition (Clayton and Dickinson 1998; Gärdenfors and Osvath 2010; Osvath 2015) who argue that episodic memory, albeit to a limited extent, can be found in animals such as great apes and corvids.

The three systems are viewed as separate systems, although they most likely work in parallel, something Tulving acknowledges (see, e.g., Tulving 2002, p. 6; see also Yee et al. 2014). We now propose a straightforward mapping between the three kinds of knowledge and the three long-term memory systems: Procedural memory handles knowledge-how, semantic memory handles knowledge-what, and episodic memory handles knowledge-that.¹¹ Since the characterisation of the knowledge handled by the three memory systems clearly maps onto our description of the three kinds of knowledge, this mapping supports that the three types we distinguish indeed have different functions in human cognition.

From the perspective of this article, it is interesting to note that Tulving claims that the order in which memory types are presented here corresponds to the order in which they have emerged in the evolution of the animal world. In Tulving’s words: “[...] Procedural memory entails semantic memory as a specialized subcategory, and [...] semantic memory, in turn, entails episodic memory as a specialized subcategory.” (Tulving 1985, pp. 2–3, italics removed). Both episodic and semantic memories therefore involve non-conscious aspects from procedural memory, which is prior. Since episodic memory is the memory-form most tightly connected with conscious experiences, thereby being connected to introspection and internalistic justification, it is no wonder that knowing-that is thought to be central for humans. However, for everyday problem solving and survival, the two other types are more essential. The fact that many animals have procedural and semantic memory while episodic memory is only well developed in humans indicates that, from an evolutionary point of view, knowing-how and knowing-what are more fundamental forms of knowledge than

¹⁰ For a more exhaustive overview see Thompson (1995).

¹¹ In different articles, Tulving vacillates in his characterisation of memories of facts. Here we link factual knowledge to episodic, rather than to semantic, memory (see also Yee et al. 2014).

knowing-that. Our argument therefore supports an anti-intellectualist position. Rather than investigating how the concept ‘knowledge’ figures in language, our mapping between knowledge and long-term memory focuses on how humans, and other animals, actually use their knowledge as shown by different cognitive tasks.

4.2 Semantic memory and neuroscience

Neuroscientific results provide ample support for Tulving’s distinction between the three memory systems, holding the procedural, semantic and episodic memory systems separate. In particular, the left and right prefrontal cortices are considered to play a key role for separating semantic and episodic memory. There is however an on-going debate concerning the details of this separation (Fletcher et al. 1999, p. 176; Goel and Dolan 2000; Kim 2016). Semantic memory is connected to conceptual knowledge, and fMRI studies show that in addition to the prefrontal cortex, the anterior cingulate, the inferior parietal cortex, the thalamus, and the hippocampus are also to various degrees involved in categorisation (Goel and Dolan 2000; Grossman et al. 2002b). Furthermore, the prefrontal cortex and hippocampus are directly linked to inductive inferences (Goel and Dolan 2000; Grossman et al. 2002a; Hayes et al. 2010; Yee et al. 2014; Fisher et al. 2015). Such findings offer a non-linguistic backing of our tripartitioning of knowledge as well as our linking of conceptual knowledge and induction to semantic memory. Specific brain-regions are correlated to categorical, conceptual and inductive inferential functions, all being ascribed to semantic memory.

The credibility of our tripartite account of knowledge, given the evidence from neuroscience, offers a counterargument against a reduction of knowledge-what (or knowledge-how) into knowledge-that. Conceptual knowledge-what should not be seen as reducible to propositional knowledge-that, since the memory system underlying knowledge-what is different from that underlying knowledge-that.

The upshot is that the mapping between types of knowledge and long-term memory systems thus provides us with a naturalistic argument for separating knowledge-what from knowledge-that. Indeed, memory science endorses the view that long-term memory can be partitioned into three types, which supports our corresponding distinction between three types of knowledge. In brief, knowledge-what is a special form of knowledge, just as semantic memory is a special form of memory.

5 Using conceptual spaces to model knowledge-what

We claim that knowledge-what is a different type than knowledge-that. What except for language can be used to model knowledge-what?¹² In this section we propose that *conceptual spaces* (Gärdenfors 1990, 2000, 2014) is an appropriate tool for this task. This notion can be seen as a development of the ‘quality spaces’ in Quine (1960), the ‘attribute spaces’ in Carnap (1971) and the ‘logical spaces’ in Stalnaker (1981). In section 6 we then argue that conceptual spaces help us understand induction as a way of achieving knowledge-what.

¹² We cannot use the word ‘describe’ instead of ‘model’ since that would presuppose a linguistic approach to knowledge.

There exist several other models of categories and their relations apart from conceptual spaces, for example models based on prototypes in the tradition of Rosch (1975) or exemplar-based models (Nosofsky 1988). However, the focus on geometrical structure, in particular the use of convexity in representing categories, make conceptual spaces particularly well suited for handling inductive processes (Gärdenfors 1990, 2000).

5.1 Dimensions and domains

A conceptual space consists of a number of quality dimensions. Examples of quality dimensions are temperature, weight, brightness, pitch, and force, as well as the three ordinary spatial dimensions of height, width, and depth. Some quality dimensions are of an abstract non-sensory character.

The quality dimensions are grouped into *domains*. For example, the space domain consists of the dimensions width, depth and height, and the colour domain of the dimensions hue, saturation and brightness. The domains are described with the aid of different topological or metric structures. For example the ordinary space domain forms a 3-dimensional Euclidean space, the colour domain forms a double spindle (Gärdenfors 2000), and the domain of tonal harmony forms a torus (Shepard 1982).

The primary function of the domains is to represent various qualities of objects. Distances in the domains are inversely correlated to the *similarities* between properties. For example the distance between orange and red in the colour domain is smaller than the distance between red and green. The domains of a conceptual space are related in various ways, since the properties of those objects modelled in the space co-vary. For example, in the fruit domain, the ripeness and colour dimensions co-vary and, of course, size and weight covariate strongly. Such covariations are central to inductive inferences.

The conceptual space framework presented here could be the answer to what for example Yee et al. (2014) are looking for in a domain-specific framework for semantic memory:

Many of the studies described in this chapter explored the organization of semantic memory by comparing the neural responses to traditionally defined categories (e.g., animals vs. tools). However, a more fruitful method of understanding conceptual representations may be to compare individual concepts to one another, and extract dimensions that describe the emergent similarity space. (Yee et al. 2014, p. 363)

5.2 Conceptual spaces as a tool for expressing properties, categories, and their relations

In first-order logic and other logical formalisms, *properties* are described with the aid of predicates. However, predicates are treated as atoms and not further analysed. In contrast, if conceptual spaces are used to define properties, more structure can be represented. The central role of similarity and the geometry of the spaces make it possible to represent features of concepts and their relations that are more or less impossible to express within a logical approach (that is, as part of knowledge-that).

The following criterion was proposed in Gärdenfors (1990, 2000), where the geometrical characteristics of the quality dimensions are used to introduce a spatial structure to properties:

- *Criterion P*: A natural property is a convex region in some domain.

That a region is convex means that, if some objects located at x and y in relation to some domain are both examples of a property, then any object that is located *between* x and y with respect to the same domain will also be an example of the property. As an application of criterion *P*, Jäger (2010) has provided strong support for the convexity of colour terms in 109 languages. Properties as defined in this criterion are natural in the sense that they emerge as results of learning in children, adults and many animal species. We will discuss the relations to learning further in section 6.

The notion of a natural property can also be extended to some discrete dimensions. For example, in a graph structure with nodes and arcs, we have a notion of betweenness, and thus we can identify the convex sub-sets of the graph (compare Johnson's (1921, pp. 181–3) notion of 'adjectival betweenness'). This means that in a biological classification, which can be represented by a tree structure, a property is 'natural' if it applies to all and only those parts of the classificatory tree that lie below one particular node in the tree. For example, the properties 'marsupial' and 'vertebrate' will be natural properties in the phylogenetic classification, while 'featherless' and 'biped' will not.

Properties, as defined by criterion *P*, should be distinguished from *categories*. Gärdenfors (2000, 2014) defines this distinction by saying that a property is based on a single domain, while a category is based on one *or more* domains. This distinction has been obliterated in the philosophical literature since both properties and categories are represented by predicates in first-order logic. A rule of thumb is that adjectives in a language typically express properties, while nouns express categories. This point is developed in Gärdenfors (2014).

When representing a category, one of the first problems one encounters is to decide which the relevant domains are. A typical example of a category that is represented in several domains is 'apple' (compare Smith et al. 1988). When we encounter apples as children, the first domains we learn about are those of colour, shape, texture, and taste (see, e.g., Son et al. 2008; Gärdenfors 2017). Later, we learn about apples as fruits (biology), and as things with nutritional value, etc.

Categories are not just bundles of properties. They also involve relations and *covariations* between regions from different domains that are associated with the category. The 'apple' category has a strong positive covariation between sweetness in the taste domain and sugar content in the nutrition domain, and a weaker positive covariation between redness and sweetness. Such considerations motivate the following definition for category¹³:

- *Criterion C*: A natural category is represented as a set of convex regions in a number of domains, together with information about how the regions in different domains are related.

¹³ For a more precise definition, see Gärdenfors (2000, ch. 4).

Within the philosophical tradition, a forerunner is Johnson's (1921, ch. XI) distinction between 'complex determinables' (corresponding to our domains) and 'determinates' (corresponding to points or regions of domains).

The theory of conceptual spaces has clear connections to prototype theory, according to which members in each category in a domain are more or less typical. The member that is most typical can be dubbed a prototype, although it can be pointed out that properties often do not have clear-cut lines but instead graded boundaries (Rosch 1975; Smith et al. 1988; Decock et al. 2013). This is easily translated into the terminology of conceptual spaces where a prototype can be described as lying at the centre of the region(s) representing a property or category.

One aspect that deserves to be highlighted is that conceptual spaces offer the ability to add domains to the representation of a concept (Gärdenfors 2000, 2104). To use our previous example, when we learn the meaning of 'apple' as children, the shape, colour and taste domains are the central ones. Later we learn that apples also have nutritional values, which can be represented by adding a new domain to the 'apple' category. Adding new domains is a form of *learning* about categories.

The connections to prototype theory and the possibility to learn about a category by adding new domains entails that our model of properties and categories is in conflict with the classical approach where concepts are defined in terms of necessary and sufficient conditions. The classical approach presumes a language-based description of concepts, something that is not presupposed when concepts are represented in terms of conceptual spaces.

Our interpretation of conceptual spaces is instrumentalistic. Nevertheless, our evolutionarily moulded cognitive faculties provide some natural quality dimensions for humans. Our quality dimensions are what they are because they have been selected to fit the surrounding world (Gärdenfors 2000, p. 82). In Quine's words: "To trust induction as a way of access to the truths of nature [...] is to suppose, more nearly, that our quality space matches that of the cosmos." (Quine 1969b, p. 125). His notion of 'quality space' is close to that of a conceptual space.

5.3 Knowledge-what as relations between categories and properties

We might not have unbiased contact with the world, but it is still the real world that provides the sensory input we get, and "[i]t is precisely because the world has the causal structure required for the existence of natural kinds that inductive knowledge is even possible." (Kornblith 1993, p. 35). There are only certain clusters of properties that are organised in a stable enough way as to stick together in natural categories enabling us to make inductive inferences.

As a way of capturing clusters of properties, criterion *C* introduces relations between domains as a factor of an object category. Our proposal is that knowledge-what consists of such relations. For example knowing what aspartame is, involves knowledge about the relation between the chemical domains that characterise aspartame and the sweetness region of the taste domain.¹⁴

There are, however, different kinds of relations. The strongest one is when all examples of a category fall within one region of a domain, as in for example "all

¹⁴ See Gärdenfors (2000) for a discussion of this domain.

ravens are black”. Another form of relation is *covariation*, for example, “metals expand when heated” which describes a covariation between the temperature and size domain, or the covariation between the colour and sweetness of fruits.

Even though there are many possible relations between categories and properties, we most often are able to discern which relations are relevant. We have a built-in understanding of the world’s structure, where some properties and patterns are intuitively grasped (Kornblith 1993, pp. 100–1; Johansson 1998). As Kornblith points out “[...] we are accomplished detectors of multiple, clustered patterns of covariation.” (Kornblith 1993, p. 104). It is primarily in contrived situations that our inductive inferences tend to go wrong; in natural settings we are quite apt at recognising essential ‘deep similarities’:

It is thus safe to say that we have a sensitivity to the features of objects which reside in homeostatic clusters. Indeed, the way in which we detect covariation is precisely tailored to the structure of natural kinds. [...] We conceptualize kinds in such a way in order to separate the properties of the members of a kind which are projectable from those which are not. We are aided in this task by our ability to detect clustered covariation. (Kornblith 1993, pp. 105–6)

An argument for focusing on covariations between domains comes from work by Billman (1983) and Billman and Knutson (1996) that indicates that humans are quite good at detecting covariations that cluster several domains (Hayes et al. 2010). A plausible explanation of this phenomenon is that our perceptions of natural objects show covariations along multiple domains, and, as a result of natural selection, we have developed a competence to detect such clustered relations. In line with this, the *basic level* categories of prototype theory (Rosch 1975) can be characterised by distinctive clusters of covarying properties (Holland et al. 1986, pp. 183–4).

6 Induction as generating knowledge-what

As we mentioned earlier, the propositional approach to induction led to unintuitive conclusions visible in numerous paradoxes. Quine’s (1969b) negative conclusions concerning the possibilities of defining ‘natural kind’ or the corresponding notion ‘similarity’ can be interpreted as indicating that we have to go beyond language to find a solution. What is needed is a way of tapping our sources of knowledge so that we become able to distinguish the properties that may be used in inductive inferences from those that may not.

For Goodman (1983), the question of what makes certain generalisations law-like becomes the problem of which predicates are ‘projectable’, that is, which predicates can be used in inductive inferences.¹⁵ The solution we propose here is that only natural properties and categories, as defined in criteria *P* and *C*, are ‘projectable’, that is, allowed in inductive inferences (Gärdenfors 1990, 2000). Consequently, the feature of a conceptual space that is the most essential for a theory of induction is its topological

¹⁵ Goodman spells the term ‘projectable’ in his well-known discussion of induction and entrenched predicates (Goodman 1983).

and metric properties, while the logical structure of the language that ‘lives on’ in the conceptual space is secondary.

As an example, let us take a brief look at the categories that occur in Hempel’s (1965) paradox of confirmation. The paradox describes how all observations of black ravens confirm the generalisation that all ravens are black. However, all observed non-black non-ravens logically confirm the same generalisation, which might be considered counterintuitive or paradoxical. Observing a white shoe, for example, would confirm that all ravens are black. It seems odd that any such observation should support an inductive inference that all ravens are black. According to the theory of Gärdenfors (2000, 2014), object categories are represented in product spaces, where each subspace represents a property of the category. If the properties of the category all correspond to convex regions, then the product of the regions will be convex too.¹⁶ In contrast, the category ‘non-raven’ would be difficult to count as a natural category. The class of all objects that are non-ravens belong to many unrelated domains. The associated regions, let alone their product, cannot be specified as a convex region of some domain. Consequently, ‘non-raven’ does not qualify as a natural category. A similar analysis can be provided for Goodman’s (1983) example of ‘grue’ (Gärdenfors 1990). The properties used in these problems do not correspond to convex regions in domains and are hence not projectable. A more detailed discussion of this topic is presented in Gärdenfors (2000).

Conceptual spaces help us understand induction as a way of achieving knowledge-what. Given the characterisation of projectable predicates as natural properties and categories, our analysis of what is achieved in induction is knowledge about relations between such properties and categories. Induction thus lets us achieve new knowledge-what about categories. And to know what properties a specific category is related to is to have knowledge about characteristic properties of that category. Most scientific empirical discoveries are of this type. For example, when it was discovered that penicillin is an antibiotic, such a relation was established (Aldridge et al. 1999; Lax 2004). Or when it was discovered that a certain alloy of niobium and titanium was a superconductor (Berlincourt and Hake 1963) new knowledge about characteristic properties of the type knowledge-what was acquired.

Induction consists in generalising from a limited number of observations. In logical approaches to induction, ‘generalisation’ means forming some form of universal sentence. When conceptual spaces are used as a basis, however, the situation is different. The similarity structure of the domains allow generalisation in the form of extending the given observations to *similar* instances, in particular by applying the convexity criteria P and C . This form of generalisation therefore comes closer to what is called ‘stimulus generalisation’ in psychology. Unfortunately, this form of generalisation has not been discussed in the philosophical literature on induction.

In contrast to the sentential approach to induction in philosophy, there exist in psychology an active research programme dealing with ‘category-based induction’ (Osherson et al. 1990; Hayes et al. 2010; Fisher et al. 2015). Within this programme, the stimuli almost exclusively consist of generic sentences that, according to our classification, express knowledge-what. The inferences that are studied are typically

¹⁶ The model presented in Gärdenfors (2000, 2014) is slightly more complicated, involving also correlations between properties.

of two kinds: general, where the conclusion concerns a class that is superordinate to those of the premises, and specific, where the class of the conclusion is on the same categorical level as the premises. An example of a general argument is the following:

- Grizzly bears love onions
- Polar bears love onions

Hence: All bears love onions.

And an example of a specific argument is:

- Robins use serotonin as a neurotransmitter
- Bluejays use serotonin as a neurotransmitter

Hence: Geese use serotonin as a neurotransmitter.

Experimental subjects are asked to judge the validity of different inductive relations.¹⁷ A central question that is investigated is how the perceived similarities between the categories affect the judgments. Thus category-based induction is closely related to stimulus generalisation. In accordance with our analysis, the focus of this research programme is knowledge-what. As far as we are aware the distinction between knowledge-what and knowledge-that has not been discussed within this psychological tradition.

Further support for the thesis that inductive generalisations build on relations between categories comes from studies of how children reason (Sutherland and Cimpian 2017). It has been argued that the drive to learn about categories is an innate feature of human cognition (Csibra and Gergely 2009). Information about categories is privileged in memory since children are better able to recall new information about categories than to recall information about non-category sets (Cimpian and Erickson 2012). Furthermore, children find it easier to reason with categories (dogs) than with set-expressions (all dogs) (Hollander et al. 2002). Findings of this type indicate that knowledge-what is primary to knowledge-that, just as semantic memory is primary to episodic.

It should be noted that the change in how induction is perceived – from relations between sentences to relations between properties and categories – does not lead to any radical changes in the *methodology* used to establish inductive knowledge. Well-known requirements of repeated experiments, precision, variation and generalisability in experiments are still valid (Hempel 1965; Seltman 2015). Furthermore, these requirements turn out to be even more natural from the perspective of establishing knowledge-what. As noted above, generalisability achieves a different meaning and different methods for determining relations that take distances in domains into consideration should therefore be put in focus.

7 Concluding remarks

We have argued that the traditional problems for the logical positivists' analyses of induction have arisen because they confined themselves to narrowly to sentential

¹⁷ When we write about validity, we intend not just logical validity, but use the term in a broader sense including other forms of inference.

representations of information and to logical tools in their analyses. Instead we have shown the fruitfulness of using conceptual spaces as a way to represent knowledge-what and to investigate inductive inferences.

We have defended two theses. Firstly, there is not only knowledge-how and knowledge-that, but also knowledge-what. Secondly, induction concerns knowledge-what, that is, knowledge concerning the relation between categories and properties. We have presented support for these theses by connecting our tripartition of knowledge to the procedural, semantic and episodic long-term memory systems. We have specifically stressed the correlations in brain activity found between semantic memory, conceptual knowledge and induction. Knowledge-what should thus be included as a fundamental component of an account of human knowledge.

It is time to give up the focus on propositional knowledge in analytical philosophy. In our opinion, the many riddles of induction are a consequence of this focus and they will not appear if knowledge-what is accepted as a type of knowledge and induction is recognised as involving knowledge-what. By introducing the tripartition of knowledge, we hope to reboot epistemology in a naturalistic direction. Since philosophical (but not psychological) research on inductive processes during the last century has focused on symbolic representations of knowledge-that, we propound that the representations of categories and properties – as a way of modelling knowledge-what – should be given much more attention in the future.

Acknowledgements We want to thank Ingvar Johansson, Jan Faye, and our anonymous reviewers for valuable comments on earlier version of the text. Peter Gärdenfors gratefully acknowledges support from the Swedish Research Council to the Linneaus environment Cognition, Communication and Learning. He also thanks Wissenschaftskolleg zu Berlin for support during the writing of this paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Aizawa, K., & Gillett, C. (2009). The (multiple) realization of psychological and other properties in the sciences. *Mind & Language*, 24(2), 181–208. <https://doi.org/10.1111/j.1468-0017.2008.01359.x>.
- Aldridge, S., Parascandola, J., & Sturchio, J. L. (1999). *Discovery and development of penicillin*. American Chemical Society International Historic Chemical Landmarks. <http://www.acs.org/content/acs/en/education/whatischemistry/landmarks/flemingpenicillin.html/>.
- Armstrong, D. M. (1978). *A theory of universals*. Cambridge: Cambridge University Press.
- Armstrong, D. M. (1983). *What is a law of nature*. Cambridge: Cambridge University Press.
- Berlincourt, T. G., & Hake, R. R. (1963). Superconductivity at high magnetic fields. *Physical Review*, 131(1), 140–157. <https://doi.org/10.1103/PhysRev.131.140>.
- Billman, D. (1983). *Procedures for learning syntactic structure: A model and test with artificial grammars*. Doctoral dissertation: University of Michigan.
- Billman, D., & Knutson, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(2), 458–475.
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, 15(11), 527–536. <https://doi.org/10.1016/j.tics.2011.10.001>.
- Boër, S., & Lycan, W. (1986). *Knowing who*. Cambridge: Cambridge University Press.

- Carlson, G. N. (2009). Generics and concepts. In F. J. Pelletier (Ed.), *Kinds, things and stuff: Mass terms and generics* (pp. 16–36). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195382891.003.0002>.
- Carnap, R. (1950). *Logical foundations of probability*. Chicago, IL: Chicago University Press.
- Carnap, R. (1971). A basic system of inductive logic, part 1. In R. Carnap & R. C. Jeffrey (Eds.), *Studies in inductive logics and probability* (Vol. 1, pp. 35–165). Berkeley: University of California Press.
- Carroll, J. W. (2016). Laws of nature. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition). URL = <http://plato.stanford.edu/archives/fall2016/entries/laws-of-nature/>
- Cerella, J. (1979). Visual classes and natural categories in the pigeon. *Journal of Experimental Psychology: Human Perception and Performance*, 5(1), 68–77.
- Cimpian, A., & Erickson, L. C. (2012). Remembering kinds: New evidence that categories are privileged in children's thinking. *Cognitive Psychology*, 64(3), 161–185. <https://doi.org/10.1016/j.cogpsych.2011.11.002>.
- Clayton, N. S., & Dickinson, A. D. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, 395(6699), 272–274. <https://doi.org/10.1038/26216>.
- Creath, R. (2014). Logical empiricism. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition). URL = <http://plato.stanford.edu/archives/spr2014/entries/logical-empiricism/>
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153. <https://doi.org/10.1016/j.tics.2009.01.005>.
- Decock, L., Dietz, R., & Douven, I. (2013). Modelling comparative concepts in conceptual spaces. In Y. Motomura, A. Butler, & D. Bekki (Eds.), *New Frontiers in artificial intelligence* (pp. 69–86). Heidelberg: Springer. https://doi.org/10.1007/978-3-642-39931-2_6.
- Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27–51. <https://doi.org/10.2307/2027085>.
- Dretske, F. (1977). Laws of nature. *Philosophy of Science*, 44(2), 248–268. <https://doi.org/10.1086/288741>.
- Fantl, J. (2016). Knowledge how. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition). URL = <https://plato.stanford.edu/archives/spr2016/entries/knowledge-how/>
- Fisher, A. V., Godwin, K. E., Matlen, B. J., & Unger, L. (2015). Development of category-based induction and semantic knowledge. *Child Development*, 86(1), 48–62. <https://doi.org/10.1111/cdev.12277>.
- Fletcher, P. C., Büchel, C., Josephs, O., Friston, K., & Dolan, R. J. (1999). Learning-related neuronal responses in prefrontal cortex studied with functional neuroimaging. *Cerebral Cortex*, 9(2), 168–178. <https://doi.org/10.1093/cercor/9.2.168>.
- Gärdenfors, P. (1990). Induction, conceptual spaces and AI. *Philosophy of Science*, 57(1), 78–95. <https://doi.org/10.1086/289532>.
- Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. Cambridge: Bradford Books, MIT Press.
- Gärdenfors, P. (2014). *The geometry of meaning: Semantics based on conceptual spaces*. Cambridge: MIT Press.
- Gärdenfors, P. (2017). Semantic knowledge, domains of meaning and conceptual spaces. In P. Meusburger, B. Werlen, & L. Suarsana (Eds.), *Knowledge and action. Knowledge and space* (Vol. 9, pp. 203–219). Cham: Springer. https://doi.org/10.1007/978-3-319-44588-5_12.
- Gärdenfors, P., & Osvath, M. (2010). Prospection as a cognitive precursor to symbolic communication. In R. K. Larson, V. Déprez, & H. Yamakido (Eds.), *Evolution of language: Bilingual approaches* (pp. 103–114). Cambridge: Cambridge University Press.
- Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (2002). *Cognitive neuroscience: The biology of the mind*. New York: W. W. Norton & Company.
- Gelman, S. A. (2003). *The essential child*. New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195154061.001.0001>.
- Goel, V., & Dolan, R. J. (2000). Anatomical segregation of component processes in an inductive interference task. *Journal of Cognitive Neuroscience*, 12(1), 110–119. <https://doi.org/10.1162/08989290051137639>.
- Goodman, N. (1983, [1955]). *Facts, fiction and forecast* (4th ed.). Cambridge: Harvard University Press.
- Grossman, M., Koenig, P., DeVita, C., Glosser, G., Alsop, D., Detre, J., & Gee, J. (2002a). The neural basis for category-specific knowledge: An fMRI study. *NeuroImage*, 15(4), 936–948. <https://doi.org/10.1006/nimg.2001.1028>.
- Grossman, M., Smith, E. E., Koenig, P., Glosser, G., DeVita, C., Moore, P., & McMillan, C. (2002b). The neural basis for categorization in semantic memory. *NeuroImage*, 17(3), 1549–1561. <https://doi.org/10.1006/nimg.2002.1273>.
- Hayes, B. K., Heit, E., & Swendsen, H. (2010). Inductive reasoning. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(2), 278–292. <https://doi.org/10.1002/wcs.44>.

- Hempel, C. G. (1965). *Aspects of scientific explanation, and other essays in the philosophy of science*. New York: Free Press.
- Herrnstein, R. J. (1990). Levels of stimulus control: A functional approach. *Cognition*, 37(1), 133–166. [https://doi.org/10.1016/0010-0277\(90\)90021-B](https://doi.org/10.1016/0010-0277(90)90021-B).
- Higginbotham, J. (1996). The semantics of questions. In S. Lappin (Ed.), *The handbook of contemporary semantic theory* (pp. 361–383). Oxford: Blackwell Publishers Ltd.
- Hintikka, J. (1975). Different constructions in terms of the basic epistemological verbs: a survey of some problems and proposals. In *The intensions of intentionality and other new models for modalities* (pp. 1–25). Dordrecht: D. Reidel.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge: MIT Press.
- Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology*, 38(6), 883–894. <https://doi.org/10.1037/0012-1649.38.6.883>.
- Hume, D. (1988 [f.p. 1748]). *An enquiry concerning human understanding*. Illinois: Open Court.
- Humphrey, N. (1992). *A history of the mind: Evolution and the birth of consciousness*. New York: Copernicus, Springer-Verlag. <https://doi.org/10.1007/978-1-4419-8544-6>.
- Jäger, G. (2010). Natural color categories are convex sets. *Amsterdam Colloquium 2009, LNAI 6042*, 11–20.
- Johansson, I. (1998). Pattern as an ontological category. In N. Guarino (Ed.), *Formal ontology in information systems* (pp. 86–94). Amsterdam: IOS Press.
- Johnson, W. E. (1921). *Logic, part I*. Cambridge: Cambridge University Press.
- Keil, F. C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior*, 23(2), 221–236. [https://doi.org/10.1016/S0022-5371\(84\)90148-8](https://doi.org/10.1016/S0022-5371(84)90148-8).
- Kim, H. (2016). Default network activation during episodic and semantic memory retrieval: A selective meta-analytic comparison. *Neuropsychologia*, 80, 35–46. <https://doi.org/10.1016/j.neuropsychologia.2015.11.006>.
- Kornblith, H. (1993). *Inductive inference and its natural ground: An essay in naturalistic epistemology*. Cambridge: MIT Press.
- Krifka, M. (2012). Definitional generics. In A. Mari, C. Beyssade, & F. Prete (Eds.), *Genericity* (pp. 372–389). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199691807.003.0015>.
- Ladyman, J. (2002). *Understanding philosophy of science*. London: Routledge. <https://doi.org/10.4324/9780203463680>.
- Lawler, J. M. (1973). *Studies in English generics*. (Doctoral dissertation, University of Michigan). University of Michigan Papers in Linguistics.
- Lax, E. (2004). *The mold in Dr. Florey's coat: The story of the penicillin miracle*. New York: Holt Paperbacks.
- Leslie, S. J. (2008). Generics: Cognition and acquisition. *Philosophical Review*, 117(1), 1–47. <https://doi.org/10.1215/00318108-2007-023>.
- Lewis, D. (1982). Whether report. In T. Pauli (Ed.), *Philosophical essays dedicated to Lennart Åqvist on his fiftieth birthday* (pp. 194–206). Uppsala: Filosofiska studier.
- Lorenz, K. (1977 [f.p. 1973]). *Behind the mirror: A search for a natural history of human knowledge*. London: Methuen.
- Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion in Neurobiology*, 11(2), 194–201. [https://doi.org/10.1016/S0959-4388\(00\)00196-3](https://doi.org/10.1016/S0959-4388(00)00196-3).
- Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature*, 379(6566), 649–652. <https://doi.org/10.1038/379649a0>.
- Mill, J. S. (1843). *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence and the methods of scientific investigation* (Vol. 1). London: John W. Parker.
- Nissen, H. W. (1953). Sensory patterning versus central organization. *The Journal of Psychology*, 36(2), 271–287. <https://doi.org/10.1080/00223980.1953.9712893>.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(4), 700–708.
- Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97(2), 185–200. <https://doi.org/10.1037/0033-295X.97.2.185>.
- Osvath, M. (2015). Putting flexible animal prospection into context: Escaping the theoretical box. *WIREs Cognitive Science 2015*. <https://doi.org/10.1002/wcs.1372>, 1, 5, 18.
- Pavese, C. (2015a). Knowing a rule. *Philosophical Issues: A Supplement to Nous*, 25(1), 165–188. <https://doi.org/10.1111/phis.12045>.
- Pavese, C. (2015b). Practical Senses. *Philosopher's Imprint*, 15(29), 1–25.

- Peirce, C. S. (1955 [f.p. 1883]). The general theory of probable inference. In J. Buchler (Ed.), *Philosophical writings of Peirce* (pp. 190–217). New York: Dover Publications, Inc.
- Quine, W. V. O. (1960). *Word and object*. Cambridge: MIT Press.
- Quine, W. V. O. (1969a). Epistemology naturalized. In W. V. O. Quine, *Ontological relativity and other essays* (pp. 69–90). The John Dewey Essays in Philosophy, No. 1. New York: Columbia University Press.
- Quine, W. V. O. (1969b). Natural kinds. In W. V. O. Quine, *Ontological relativity and other essays* (pp. 114–138). The John Dewey Essays in Philosophy, No. 1. New York: Columbia University Press.
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104(3), 192–233. <https://doi.org/10.1037/0096-3445.104.3.192>.
- Rosenberg, A. (2000). *Philosophy of science: A contemporary introduction* (2nd ed.). New York: Routledge.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson.
- Rysiew, P. (2016). Naturalism in epistemology. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2016 Edition)*. URL = <http://plato.stanford.edu/archives/sum2016/entries/epistemology-naturalized/>
- Sands, S. F., Lincoln, C. E., & Wright, A. A. (1982). Pictorial similarity judgments and the organization of visual memory in the rhesus monkey. *Journal of Experimental Psychology: General*, 111(4), 369–389. <https://doi.org/10.1037/0096-3445.111.4.369>.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2), 383–403. <https://doi.org/10.1111/j.1933-1592.2007.00081.x>.
- Schrier, A. M., & Brady, P. M. (1987). Categorization of natural stimuli by monkeys (*Macaca Mulatta*): Effects of stimulus set size and modification of exemplars. *Journal of Experimental Psychology: Animal Behavior Processes*, 13(2), 136–142.
- Seltman, H. J. (2015). *Experimental design and analysis*. Pittsburgh: Carnegie Mellon University.
- Shepard, R. N. (1982). Geometrical approximations to the structure of musical pitch. *Psychological Review*, 89(4), 305–333. <https://doi.org/10.1037/0033-295X.89.4.305>.
- Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, 12(4), 485–527. https://doi.org/10.1207/s15516709cog1204_1.
- Son, J. Y., Smith, L. B., & Goldstone, R. L. (2008). Simplicity and generalization: Short-cutting abstraction in children's object categorizations. *Cognition*, 108(3), 626–638. <https://doi.org/10.1016/j.cognition.2008.05.002>.
- Spaet, T., & Harlow, H. F. (1943). Solution by rhesus monkeys of multiple sign problems utilizing the oddity technique. *Journal of Comparative Psychology*, 35(2), 119–132. <https://doi.org/10.1037/h0059354>.
- Stalnaker, R. (1981). Antiessentialism. *Midwest Studies of Philosophy*, 4, 343–355.
- Stanley, J. (2011). *Know how*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199695362.001.0001>.
- Stanley, J., & Krakauer, J. W. (2013). Motor skill depends on knowledge of facts. *Frontiers in Human Neuroscience*, 7(503). <https://doi.org/10.3389/fnhum.2013.00503>.
- Stanley, J., & Williamson, T. (2001). Knowing how. *The Journal of Philosophy*, 98(8), 411–444. <https://doi.org/10.2307/2678403>.
- Sutherland, S. L., & Cimpian, A. (2017). Inductive generalization relies on category representations. *Psychonomic Bulletin and Review*, 24(2), 632–636. <https://doi.org/10.3758/s13423-015-0951-z>.
- Thomas, R. K., & Kerr, R. S. (1976). Conceptual conditional discrimination in *Saimiri Sciureus*. *Animal Learning & Behavior*, 4(3), 333–336. <https://doi.org/10.3758/BF03214060>.
- Thompson, R. K. (1995). Natural and relational concepts in animals. In H. L. Roitblad & J. A. Meyer (Eds.), *Comparative approaches to cognitive science* (pp. 175–224). Cambridge: MIT Press.
- Tooley, M. (1977). The nature of laws. *Canadian Journal of Philosophy*, 7(4), 667–698. <https://doi.org/10.1080/00455091.1977.10716190>.
- Tulving, E. (1985). Memory and consciousness. *Canadian Psychology*, 26(1), 1–12. <https://doi.org/10.1037/h0080017>.
- Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, 53(1), 1–25. <https://doi.org/10.1146/annurev.psych.53.100901.135114>.
- Vetter, G., & Hearst, E. (1968). Generalization and discrimination of shape orientation in pigeons. *Journal of the Experimental Analysis of Behavior*, 11(6), 753–765. <https://doi.org/10.1901/jeab.1968.11-753>.
- Vickers, J. (2014). The problem of induction. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2014 Edition)*. URL = <http://plato.stanford.edu/archives/fall2014/entries/induction-problem/>
- Yee, E., Chrysikou, E. G., & Thompson-Schill, S. L. (2014). Semantic memory. In K. Ochsner & S. Kosslyn (Eds.), *The Oxford handbook of cognitive neuroscience: Volume 1, core topics* (pp. 353–374). Oxford: Oxford University Press.
- Zeiler, M. (1969). Repeated measurements of reinforcement schedule effects on gradients of stimulus control. *Journal of the Experimental Analysis of Behavior*, 12(3), 451–461. <https://doi.org/10.1901/jeab.1969.12-451>.