Accepted for publication in *Cybernetics & Systems*, special issue on “Belief systems in science and engineering,” edited by Josué Antonio Nescolarde-Selva, José Luis Usó-Doménech and Hugh Gash.

Knowledge and Belief: Some Clarifications

Alexander Riegler, ariegler@vub.ac.be
Center for Logic and Philosophy of Science, Vrije Universiteit Brussel, Belgium

ABSTRACT

The notions of knowledge and belief play an important role in philosophy. Unfortunately, the literature is not very consistent about defining these notions. Is belief more fundamental than knowledge or the other way around? Many accounts rely on the widely accepted strategy of appealing to the intuition of the reader. Such an argumentative methodology is fundamentally flawed as it lets the problems of common sense reasoning in through the front door. Instead, I suggest that philosophical arguments should be based on formal-computational models to (a) reduce the ambiguities and uncertainties that come with intuitive arguments and reasoning, and (b) capture the dynamic nature of many philosophical concepts. I present a model of knowledge and belief that lends itself to being implemented on computers. Its purpose is to resolve terminological confusion in favor of a more transparent account. The position I defend is an anti-realist naturalized one: knowledge is best conceived as arising from experience, and is fundamental to belief.

INTRODUCTION

In orthodox epistemology, knowledge is defined as justified true belief (JTB). This view has been challenged for either being too broad and incomplete, or for being too restrictive because it does not cover different varieties of knowledge, or for referring to epistemically questionable notions. This heterogeneity of often diametrically opposed claims suggests that the discussion so far may have relied too much on appeals to intuition. Therefore, any alternative account should preferably draw on a different methodology in order to capture better the nature of knowledge and belief and how they relate to each other.

In this paper I distinguish between (consciously held) beliefs and (non-consciously or tacitly formed) concepts and schemata, which both make up a subject’s knowledge. Concepts consist of the aggregation of simple expectations while schemata are conditional structures composed of concepts and actions. I argue that beliefs are reflections of an observer who observes either others or herself. Much has been written about the former case, e.g., Daniel Dennett’s system of different stances an observer can use in order to attribute belief to what or whom she observes. In the latter case of self-observation, having a belief is the (conscious) act of reflecting on and comparing the constitutive elements in one’s own knowledge. For example, Alyssa’s knowledge of an apple consists of the expectations that its color is a certain type of green and that it has a certain size. While these expectations have been formed non-consciously in the course of Alyssa's ontogenesis, she can consciously relate the color expectation with "apple" and thus form the belief that "apples are green."

I will first present arguments that the JTB account of knowledge has little epistemic value for radical constructivism. This motivates a different definition, with the intention to transcend mere appeals to intuition in order to resolve terminological confusion. A promising candidate for bringing about progress in such an alternative conceptual analysis is the formal-computational approach. I will sketch how a computationally motivated account of knowledge and belief could look, and how it relates to the philosophical literature.

THE ORTHODOX PICTURE: KNOWLEDGE ENTAILS BELIEF

Deriving from Plato's dialogue Theaetetus and in particular Bertrand Russell's work on the theory of propositions and propositional attitudes ("Believing [is] the central problem in the analysis of mind," Russell 2005: 231), a belief is commonly referred to as an attitude towards a proposition: "Beliefs are a species of propositional attitude distinguished by their having the mind-to-world direction of fit" (Ludwig 2006: 532). In orthodox epistemology, beliefs are thought to be constitutive for knowledge, which is defined as justified true belief. This definition has been attacked from many angles. In particular, attention was drawn to the Gettier cases,[1] which seem to require additional conditions for the definition of knowledge. But the components "justification" and "truth" have also been criticized – as has "belief" itself, for many are hesitant to attribute belief to animals and young humans.

Truth

For many philosophers it seems intuitively clear that truth is uncontroversial and mind-independent. However, others have contended this view, for example pragmatists and social constructivists, for whom truth is a social regulative. Such non-factive conceptions of truth rely on what society or a certain professional group considers true or accepts as true.

Radical constructivists[2] such as Ernst von Glasersfeld draw attention to the skeptical argument that, from the perspective of the cognitive subject, any attempt to verify the truth content of a proposition needs to be carried out by the very same sense organs that

---

[1] Edmund Gettier (1963) presented two counterexamples in which a person forms true and well-justified beliefs that (arguably for most people) intuitively do not count as knowledge.

[2] The notion "radical construction" was coined by von Glasersfeld (1974). It has been applied in various disciplines such as educational research (e.g., Dykstra 2005; Ulrich et al. 2014), communication science (e.g., Scholl 2010), sociology (e.g., Glasersfeld 2008), psychotherapy (e.g., Raskin 2011), ethics (e.g., Glasersfeld 2009; Quale 2014), mathematics (e.g., Glasersfeld 2006; Cariani 2012) and artificial intelligence (e.g., Riegler, Stewart & Ziemke 2013; Füllsack 2013).

participated in formulating the proposition in the first place.[3] In other words, there is no independent arbiter that could decide the truthfulness of any given statement.

The irrelevance of the notion of truth for epistemic purposes was also emphasized by logicians such as Peter Gärdenfors (1988) who consider "truth" irrelevant for the analysis of belief systems (see also further below).

Instead of "truth" von Glasersfeld suggests "viability" as a regulative, i.e., "actions, concepts, and conceptual operations are viable if they fit the purposive or descriptive contexts in which we use them" (Glasersfeld 1995: 14). This relieves radical constructivists from metaphysical commitments (ibid: 22). For them truth has a mere rhetorical rather than an epistemic function (Mitterer 1992; Riegler 2001): identifying one's own perspective as the truth makes it unassailable, it justifies research expenses and it seduces people into doing what they would not be doing otherwise.

Justification

Justifying a belief means either finding evidence supporting the belief or having formed the belief in reliable ways (Goldman 1979).

Many arguments have been made to the effect that justification should not be part of the JTB definition of knowledge because: (a) it does not cope with "being lucky" situations; (b) it does not account for practical knowing-how (how can we justify riding a bicycle or baking a fine cake?); and it (c) runs into the Münchhausen trilemma of infinite, circular or arbitrarily stopped regresses (Albert 1968). In particular, from the perspective of RC, the notions of viability and justification conflate: if certain knowledge is viable, its use is justified by its usefulness.[4]

Even in the case of the less problematic propositional knowing-that, it is questionable whether human subjects are, in fact, capable of justifying their beliefs. Frank Keil's (2003) studies revealed that people's explanatory understanding of everyday and complex phenomena is shallower and less coherent than they think. Such "illusion of explanatory depth" must also impact one's ability to justify one's belief: are we correct in assuming that we considered all necessary details and their mutual interdependencies? Keil's results leave us with doubts as to whether justification is a *useful* criterion, for it cannot be applied with certainty. This is not about demanding an infallibilist standard but rather asking for Occam's razor: why convolute the notion of knowledge by introducing a criterion that cannot save it from ambiguity? Justification simply does not seem to motivate the distinction between knowledge and belief.

Both justification and the truth conditions are of little or no value for the radical constructivist: truth is replaced by viability, and justification is just a synonym for being formed in reliable cognitive processes. So perhaps the philosophers' intuition about

---

[3] In particular, von Glasersfeld draws attention to the impossibility of executing Wittgenstein's instruction "In order to discover whether the picture is true or false we must compare it with reality" (Wittgenstein 1922, §2.223).

[4] This argument can be extended to cover innate knowledge, which reflects knowledge that has proven useful on a phylogenetic scale.

knowledge and belief is misplaced? In what follows I first discuss the notion of intuition before I suggest reversing the entailment relation between knowledge and belief.

SHOULD WE RELY ON INTUITION?

Despite all the careful argumentation and sincerity philosophers use, one cannot help but notice an enormous variation of how concepts and their mutual relations are defined.[5] There is a mainstream view of defining beliefs as constituents of knowledge but there is also the reverse and widely accepted view that knowledge is more fundamental than belief – although, as Myers-Schulz & Schwitzgebel (2013) complain, these alternatives have not been clearly formulated. Some philosophers argue convincingly that animals cannot have beliefs while others, equally convincingly, arrive at the opposite conclusion. They all make appeals to common-sense intuition to convince their audience. How does this situation compare to the sciences and engineering? Are there, for example, various (mutually contradicting) views about static properties of bridges to be built? Are physicists divided over the question of whether fire releases phlogiston in wood? Evidently, I am referring to three different domains of inquiry:[6] the conceptual-philosophical domain (D2), which refers to the intuitions formed in the domain of common sense (D1) manifest in folk psychology and other forms of layman reasoning, and the formal domain of strict quantitative exactness and systematicity, which is in the service of science (D3).

For many, D1 and D3 are at opposite poles. Lewis Wolpert, in his *The Unnatural Nature of Science*, is rather blunt about it: "I would almost contend that if something fits in with common sense it almost certainly isn't science" (Wolpert 1992). Clearly, the world of common sense reasoning and having intuitions is afflicted with fundamental weaknesses, which become evident when being confronted with phenomena of exponential nature: long-term calculation of interests do not follow the layman's intuitions and neither do probabilistic estimations and logical inferences (e.g., Wason's 1966 empirical result that only a small minority of people is able to apply *modus tollens* inferences properly). The more complex the systems under inquiry, the more difficult it is to deal with these systems. Dietrich Dörner (1996) confronted test subjects with the complex effect network in a fictitious developing country. Despite their intentions to the contrary, most subjects ruined the country more rapidly than if they had not intervened at all. Also, in a computer simulation of a typical European town, subjects were incapable of appropriately analyzing the network structure. Rather, they relied on *intuitive interpretations* of the state, neglecting side-effects and future long-term impacts, and treated the complex net of interdependencies among variables as simple linear accumulation of facts. They focused on a single core variable that then became the starting point for a long chain of causal connections. This strategy reduced cognitive efforts and provided the subjects with the illusion that the system is controllable.

---

[5] Yet, these are not necessarily clear-cut contradictions: "Identifiable contradictions are vanishingly rare in philosophy, and are something of a badge of honour when they do occur" (Stove 1991).

[6] These three domains reflect different methodological tools but do not imply that physics and engineering supply the only epistemic standard that all forms of knowledge must meet (cf. Figure 1 below).

Without a single doubt, intuition must have been indispensible in our ancestors' struggle for survival but the insights gained from these experiments undermine the epistemic trust we have in intuition:[7] they are blinders that focus on one detail while ignoring the context. As such, intuition may be invaluable for scholars who want their audience to overlook obvious gaps in their argumentation and theories; but for philosophers interested in "finding connecting links" (Wittgenstein 1993), appeals to intuition should be a less favorable course of action.[8]

Using such appeals is tempting: Arguments are accepted because *intuitively* they *seem* plausible and reasonable.[9] However, these are the same intuitions that let people fail when interacting with exponential and nonlinear problems. But there is a further concern. For Ludwig Wittgenstein, philosophy is "a battle against the bewitchment of our intelligence by means of language" (Wittgenstein 1953, §109) – the same inexact language that we avail ourselves of in philosophical arguments and that is an obstacle to forming a perspicuous representation that "brings about the understanding which consists precisely in the fact that we 'see the connections'" (Wittgenstein 1993: 133). Meaning depends on the context and on the experiential background of the readers or listeners, in whom most diverse associations are triggered, and "uttering a word is like striking a note on the keyboard of the imagination", as Wittgenstein puts it. If argumentation is portrayed as *orienteering* between the premise and the goal, in which one argument has to be checked before the next argument in the chain can build on it,[10] it is hard to see how such semantic uncertainty could possibly constitute progress in any way: philosophical arguments presented in everyday language offer numerous connecting points whose link with the topic in question can be *arbitrarily* (and sometimes *deliberately*) stressed or neglected.[11] What is clearly missing is the

---

[7] Proponents of experimental philosophy have criticized intuitions on other grounds, i.e., for being culturally biased: when philosophers appeal to them they refer to the intuitions of people affiliated with academic institutions and of European ancestry. Subjects outside these circles may have quite different intuitions about situations and events rendering the appeal to intuitions ambiguous (Weinberg, Nichols & Stich 2001).

[8] This is not to say that philosophy is mere "armchair reflection and informal dialogue." As Niki Pfeifer & Igor Douven (2014) point out, twentieth-century philosophy has seen several attempts to contain "the danger that philosophical speculation may go uncontrolled": First there was that of the logical empiricists, who insisted that "philosophers aim for the same rigor in their reasoning as mathematicians and logicians", then the demand of post-positivist philosophers in the 1960s to pay more attention to the sciences, the experimental philosophers who favor the return of the philosopher–scientist, and, finally, formal epistemologists, who use formal tools "to address what are for the most part traditional questions from mainstream epistemology."

[9] As Pfeifer & Douven (2014) point out, "many have come to often accept blanket appeals to intuition as providing evidence for philosophical theses." Certain philosophers defend their position by arguing that it is in line "with how the folk think…"

[10] In the context of epistemic dynamics, I will emphasis further below the importance of embedding checkpoints in the course of behavioral action and refer to it as "constructivist-anticipatory schema processing."

[11] A case in point are those philosophical thought experiments that ask the reader to transcend her horizon of experience, such as Sosa's New Evil Demon problem or Davidson's Swampman.

appropriate *quantitative* weighting of connecting arguments, and this makes it possible to fabricate any number of the most different and contradictory argumentative chains.[12]

Another way to view the differences between the domains D1 to D3 is to characterize them in terms of context-dependency, which is inversely correlated with formality. Since the members of a scientific community may have vastly different backgrounds, the intuitions triggered by verbal arguments will vary greatly as well. Words are intrinsically more prone to broader interpretation such that an argument in natural language exhibits many more degrees of freedom. To work with formal arguments offers a way to reduce the context-dependency: "human performance is affected by a number of factors, related to the content or context […], that are not, strictly speaking, relevant within a logical system" (Blanchette 2006: 1112). In other words, verbal structures invite many contexts that are detrimental to Wittgenstein's imperative of seeing the *proper* connection, while formal arguments are more precise and assist in avoiding unintended interpretations.[13] On this view, D1 to D3 define a spectrum of decreasing context-dependency and increasing formality (see Figure 1).

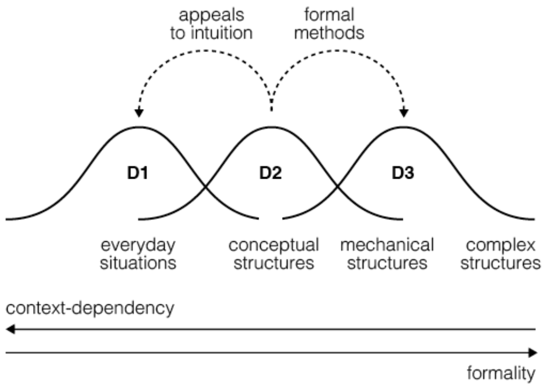

Figure 1.

The idea of introducing formal-exact methods in philosophy is, of course, not new and has brought forth disciplines such as "formal epistemology" (Hendricks 2006) and "computational philosophy" (Thagard 1988; see also Riegler & Douven 2009). My goal here largely aligns with the ambitions of these disciplines. There are noteworthy differences, however. In contrast to logical accounts, I do not focus on normative aspects that come with the logical analysis of *static* ahistorical snapshots of the web of belief an adult human being might have. Rather it deals with the question of whether knowledge forms a *dynamic* network of hierarchical interdependencies whose components have become mutually dependent in the course of their ontogeny. Such relationships are difficult to capture in logic but may lend themselves to being modeled in computer algorithms, which offer formality in space and time. Furthermore, formal and computational methods in philosophy in general and epistemology in particular have been motivated by a realist understanding. What I propose here is not only a

[12] In his portrait of the quantification of the Western society, Alfred Crosby makes a similar point with regard to the medieval Schoolmen, who, despite their "precise definitions and meticulous reasoning, that is to say, *clarity*" (Crosby 1997: 65) eventually failed for they were "mathematicians without being quantifiers" (ibid: 67).

[13] The extreme form would be what Richard Feynman (1985) called "cargo cult science": understanding on the basis of superficially analogous structures.

different methodology but also a different metaphysical orientation, i.e., a computational model of the radical constructivist account of knowledge and belief.

A CONSTRUCTIVIST COMPUTATIONAL MODEL

One of the most influential formal models is the AGM model of belief revision. Named after its creators, Carlos Alchourrón, Peter Gärdenfors, and David Makinson, it models the change in a belief system in the face of new information. In AGM, the belief state is represented as a deductively closed set of beliefs, which makes AGM a coherentistic model that refrains from defining truth in terms of correspondence with a reality. In fact, it dispenses with the notion of truth altogether: "the concepts of truth and falsity are irrelevant for the analysis of belief systems" (Gärdenfors 1988: 20):

> "These concepts deal with the relation between belief systems and the external world, which I claim is not essential for an analysis of epistemic dynamics. […] My negligence of truth may strike traditional epistemologists as heretical. However, one of my aims is to show that many epistemological problems can be attacked without using the notions of truth and falsity." (ibid)

In a similar vein, radical constructivism aims at showing that any such correspondence with an alleged mind-independent reality is irrelevant for the analysis of concepts such as knowledge and belief. Rather, knowledge serves a particular purpose i.e., the viable organization of experience. "Viable" is understood in a pragmatist sense: a piece of knowledge is viable when it supports the survival of the subject: "The function of cognition is adaptive and serves the subject's organization of the experiential world, not the discovery of an objective ontological reality" (Glasersfeld 1988: 83). Several arguments have been called on in support of rejecting the correspondence-theory of truth:

(a) All that the brain deals with are indifferent electric impulses, as expressed in Heinz von Foerster's principle of undifferentiated encoding: "The response of a nerve cell does not encode the physical nature of the agents that caused its response. Encoded is only 'how much' at this point on my body, but not 'what'" (Foerster 1973: 214f).

(b) Even if the signals were to carry the semantics of whatever triggered them, we would be in no position to verify this information in ways independent of the sensors that carried that information in the first place (von Glasersfeld's skeptical argument). Therefore any perceptual activity must be understood as experiences, and mental activity consists of making sense of these experiences.

(c) The input–output sequence is the perspective of an observer, while from the perspective of the cognitive subject one takes action to control one's input:[14] "From the organism's point of view only actions which feed back to the organism's sensors can be

[14] Francisco Varela et al. sharply distinguish between living systems that are autonomous and systems that are defined in terms of input/output relations: "...the meaning of this or that interaction for a living system is not prescribed from outside but is the result of the organization and history of the system itself" (Varela, Thompson & Rosch 1991: 157).

observed [...] Any other action which simply disappears in the environment cannot be observed by the organism" (Porr & Wörgötter 2005).[15]

What exactly does "making sense of experiences" mean?[16] It has been argued that in order to be able to see a tiger you must see it *as* a tiger and to do that you must have the concept of a tiger. In orthodox philosophy, this has been referred to as holism, i.e., the relationship of belief to the web of beliefs a person holds in the sense of William Quine. John Searle called it "background," in psychology it has been termed "tacit knowledge" and in philosophy of science "theory-ladenness."[17] In other words, the already existing knowledge of a subject entitles her to have *expectations* about what to see. William James's example is a case in point: "The Fuegians, in Darwin's voyage, wondered at the small boats, but took the big ship as a 'matter of course'" (James 1890: 110f). There was simply no concept of a big ship in these indigenous South-American people that would allow them to make sense of that ship.

Having such expectations are the foundation of any cognitive process: "Our talk of external things, our very notion of things, is just a conceptual apparatus that helps us to foresee and control the triggering of our sensory receptors in the light of previous triggering of our sensory receptors" Quine (1981: 1).[18] In order to see an apple we use a conjunction of expectations: x is green, x is round and x has a certain diameter. Together these expectations form the concept of "apple."

Evidently, this is a rather simplified account. Apples may differ in color and they may have different diameters, yet a subject will still call them apples. Furthermore, similar to Wittgenstein's prototype theory, such details may not be obligatory. There are red apples after all. In order to account for such variety it is necessary to define expectations as probabilities or, more precisely, as Gaussian distributions with a certain significance and specificity. To recognize an instance of a concept, the sum of the expectations subsumed under a concept needs to reach a certain threshold. The significance of an

---

[15] According to Peter Godfrey-Smith, this was also Quine's (1981) position: "Beliefs give rise to actions, those actions affect our environment, and this, in turn, affects what is later experienced" such that new beliefs are functions of preexisting beliefs and present experiences. This means that "future sensory stimuli are partly under the control of present actions, which themselves depend on present beliefs." Quine posited objects as providing the feedback that connects current actions with later experience: "External objects are […] nodes in a structure whereby we partially control, by means of action, the flow of sensory input" (Godfrey-Smith 2014: 64f).

[16] In his last paper, Ernst von Glasersfeld drew attention to the fact that experience can be understood in two ways: the "person's experience as lived through or witnessed *in the here and now* and the practical wisdom gained over time" (Glasersfeld & Ackermann 2011: 194, italic in the original). This distinction is reflected in the German translation into "Erlebnis" and "Erfahrung", respectively. In the context of the terminology used here, Erfahrung corresponds to "knowing-how" for which Erlebnis furnishes the raw material, i.e., that which is made sense of.

[17] The probably most entrenched tacit knowledge is that of space and time, which for Immanual Kant was an a priori category but which Hans Vaihinger (1913) referred to as useful fiction and which for Jean Piaget were gradually built-up categories that serve as "eminently useful tools for the categorization of experience" (Glasersfeld & Ackermann 2011: 197).

[18] Cf. also Quine (1951: 41): "the conceptual scheme of science as a tool, ultimately, for predicting future experience in the light of past experience." And Quine and Ullian (1978: 108): "The immediate utility of a good hypothesis is as an aid to prediction. For it is by predicting the effect of our actions or of other observed events that we are enabled to turn our environment to best advantage."

expectation for the encompassing concept correlates with the height of its Bell curve, while its width inversely correlates with its specificity. This agrees with both Wittgenstein's concept of "family resemblance" and Eleanor Rosch's (1978) experimental results about categorization.

Knowledge for knowledge's sake would not have contributed to our ancestors' survival. Rather, knowledge has a dispositional function: it enables a subject to act (see also next section). This is accomplished by combining concepts with actions such that when Alyssa recognizes an apple, the associated action sequence would enable her to walk towards the apple, grasp it and put it into her mouth. Conditionally joining concepts with action sequences creates a schema.

Again, this is a rather simplified account, as walking towards an object and grasping it is a sophisticated kinematic problem – as robot science tells us. Part of the problem is that in the course of the execution of the action, the subject needs regular feedback about it, e.g.: Am I still heading towards the apple? Has the apple been sufficiently approached for starting the grasp action?. This requires expectations to be embedded in the action sequence, expectations such as: Has the apparent diameter of the apple become big enough such that grasping for it will succeed? Has my hand reached the apple such that it can be closed and the apple be pulled off its stem?

In the animal kingdom, there are various examples that show the efficacy of such schemata, in primitive and more sophisticated animals alike. Hungry young birds open their beaks when one of their parents shows up above them with food. This schema is a simple schema, and the birds can be easily tricked by presenting any sufficiently sized dummy above their heads which indicates that their concept of "parent bird" is underdetermined by the evidence. However, a more complex concept, say the concept of a mermaid, can also be underdetermined if its constituting expectations do not cover a crucial feature:

> "if I am not told that the fish's tail replaces the woman's legs, I may construct a notion that is more like a fish-tailed biped than like the intended traditional mermaid. My deviant notion could then be corrected only by further interaction, i.e., by getting into situations where my conception of a creature with legs as well as a fish's tail comes into explicit conflict with a picture or with what speakers of the language say about mermaids." (Glasersfeld 1983: 212)

So as long as there is no recalcitrant experience to the contrary, the incomplete but nevertheless coherent concept of mermaid allows the subject who holds the concept to understand all stories she might hear about mermaids.

Naturally, a subject will recognize and eat more than just one apple, which requires the schema to be readily accessible as a whole rather than by its constituting expectations and actions. In these cases it is convenient to – borrowing Bruno Latour's (1987) notion – "black box" experiences, concepts, action sequences and schemata and to attach a unique identification tag (e.g., a number) for later access. In this way they can easily be re-used for (similar) situations by making them available to other action sequences and schemata. For example, the "grasping an apple" sequence could be reused in a more general "food-foraging" schema.

Here is a more formal notation for the constructivist computational model:

An expectation e is defined as a Gaussian function characterized by its significance h and specificity w. An expectation takes an input value (which for an observer could be a visual receptor detecting the green-ness of an object, but see below) and maps it onto an output value (see Figure 2).

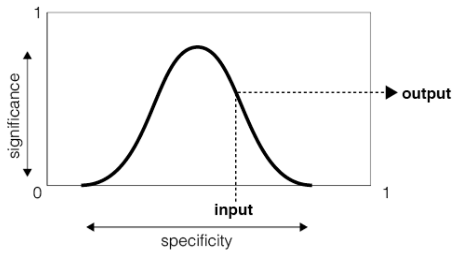

Figure 2.

A concept C is the sum of n expectations: $C = \sum_{i=1}^{n} e_i$. The concept is "satisfied" if the sum of the outputs of its constituting expectations reaches a certain threshold.

All expectations and concepts are stored in a library of concepts, L_e and L_C respectively. There are also libraries for action, action sequences and schemata: L_a , L_A and L_S.

An action sequence A is the concatenation of single actions but can also embed any other elements such as expectations, concepts, action sequences and even entire schemata: $A = \{x_n \mid x \in L_e \cap L_C \cap L_a \cap L_A \cap L_S\}$.

A schema S is a conditional of the form S: $C_i \rightarrow A_j$. Cognition is in perpetual motion. First it seeks to find the schema whose conditional part matches the current context (see below for how the context is defined). If the expectations can be met, the action sequence is executed, which in turn may call on other schemata and other elements. The execution of a schema is terminated when the last constituting element has been processed or when one of the embedded expectations or concepts are not satisfied. Consequently, a new schema is selected on the basis of the (meanwhile changed) context.[19]

At this point a philosophically inclined person may wonder what, in fact, distinguishes this model from other computational models that are defined over symbols with arbitrary meaning? And if the concept of "apple" requires expectations about its green-ness and size, are green-ness and size not concepts too?

To forestall misunderstanding, my suggestion that the input for expectations would be receptors measuring some property of some external objects in a mind-independent reality was pro tem only: for the sake of illustrating the main idea I used anthropomorphic notions to keep things simple. However, strictly speaking, for a subject's cognitive processes there are no such distinctions. Referring to von Foerster's principle of undifferentiated encoding, it would be plausible to say that a subject's cognitive processes operate on anonymous (nervous) signals that, for the sake of a computational implementation, can be represented by numerical values, for numbers are free of any connotation. Given that nervous systems in animals and humans have a

---

[19] This implements Mark Baldwin's (1895) "circular reaction" and Ulric Neisser's "perceptual cycle": "the schema accepts information as it becomes available at sensory surfaces and is changed by that information; it directs movements and exploratory activities that make more information available, by which it is further modified" (Neisser 1976: 55).

distinct structure, it is further reasonable to assume that these numbers are arranged topologically such that numbers are stored in an array, which provides the context for schemata. An observer looking into the array may even be able to ascertain a link between its elements and some sensor activity and action performed by the subject. She may further observe that a particular stimulus translates into a certain value in the associated element in the number array, and that a certain value in another element translates into a specific action of the subject.

On the one hand, the introduction of the number array complies with the radical constructivist notion of cognitive closure (Riegler 2001), claiming that the nervous system is "a closed network of interacting neurons such that a change in the state of relative activity of a group of neurons always leads to a change in the state of relative activity of other groups of neurons" (Maturana 1974: 42).[20]

On the other hand, the proposed model does not represent an instance of Alan Newell and Herbert Simon's (1976) physical symbol system hypothesis because the elements on which cognitive processes operate are free of the semantics the programmer would otherwise impose on them. The latter representational account was subject to criticism, saying that symbols must not be mistaken for signals, and according to the Frame Problem criticism symbols have no grounding at all or no grounding other than in the conceptual worldview of the programmer (which I called the "PacMan syndrome", Riegler 2001). Furthermore, the model does not distinguish between high-level activities that require consciousness such as playing chess and other forms of cognitive activity such as perception, which require connectionist models. It therefore avoids the eliminativist's criticism (Churchland 1984) that "belief" is no genuinely scientific notion and will be rendered obsolete as soon as all cognitive processes can be mapped on to connectionist models. The advantage of the proposed model is that it remains explanatorily transparent at all levels of complexity. This is crucial for philosophy, for if we do not have a principled understanding of how cognition works, we shall have succeeded as engineers but not as philosophers (Perlis 2000).

Another advantage of the model is that embedding expectations in actions offers the flexibility necessary for cognitive processes. These expectations serve as "checkpoints"[21] for verifying whether the schema still successfully copes with recalcitrant experiences. Such constructivist-anticipatory schema processing (Riegler 1994) covers two aspects:

1. The dependence of knowledge on *a priori* formed concepts – for perceptual knowledge Wittgenstein called this "aspect-seeing" (Himmelfreundpointner 2013). The cognitive subject is conceptually engaged with her experiences. Non-epistemic seeing (Dretske 1969) is an empty notion since schemata can only be executed when their condition part matches the current context in the numerical array, and remain silent otherwise.

---

[20] Working over a number array means that all cognitive processes are closure operators. A more detailed analysis of this mathematical property has to be left for another occasion.

[21] Hetherington (2012) suggests "checks" to improve imperfect thinking: "Perhaps we need observations as 'checks' on what could otherwise become thoughts 'floating free' in our minds?"

2. The anticipatory nature of knowledge – traversing through the constitutive elements of a schema (i.e., exhibiting behavior) will bring about a desired goal state: “behavior is the process by which organisms control their input sensory data” (Powers 1973: xi).

KNOWLEDGE AND ACTION

One could object that concepts such as “apple” are very simple and in no way reflect human knowledge, for (a) human knowledge is not only propositional in nature, i.e., a “knowing-that,” but is also a practical “knowing-how” (and a variety of other versions usually subsumed under the label “knowing-wh” that can be mapped onto propositional knowledge); and (b), in general, the structure of human knowledge is quite complex.

Such criticism would not take into consideration the conditional nature of schemata and the fact that in the action part of schemata every other element can be embedded by reference. Such nested embeddedness creates the level of complexity needed to match the complexity of human knowledge. Embedding elements means that executing an action sequence consists of both the execution of a single action and the reference to other elements in any of the libraries. The referenced elements in turn can refer to yet other elements, which creates an arbitrarily deep hierarchical structure that shares and reuses elements in a variety of different schemata. Therefore there is no upper limit to the complexity of behavior that can be exhibited by the proposed cognitive model. This can be easily demonstrated quantitatively. In thought experiments with creatures possessing input, internal and output states, Valentino Braitenberg (1984)[22] demonstrated that for an observer there is an opaque relationship between triggering conditions and observable behavior. The essence of his “law of uphill analysis and downhill synthesis” is that complex behavior can be generated by a simple structure. An observer may need to take the “intentional stance” to predict the intricate behavior of a creature because she is unaware of the underlying simple links that connect the creature’s input sensors with its output effectors. Heinz von Foerster (2003: 312) showed that a fairly simple architecture of four sensors, four effectors and four internal states suffices to display about 10^{126} different behaviors while one would need only 32 wires to connect input, internal and output states. Clearly there is a gross mismatch.

As for criticism (a), the conditional nature of schemata makes the suggested architecture align with the idea that knowledge and beliefs are dispositional (Ryle 1949), i.e., the subject has particular behavioral dispositions pertaining to the content of her knowledge. This does not necessarily cover only visible behavior in the sense that behaviorists made it a requirement. In fact, James’s claim that “cognition is incomplete until discharged in act” (James 1897: 85) should cover non-observable behavior as well – cf. Peter Godfrey-Smith, who considers cognition complete “when the agent has worked out a way to accommodate past experience in a way that predicts what will happen next” (Godfrey-Smith 2014: 60). The importance of behavior that does not necessarily need to be outward but that produces private mental episodes instead was already anticipated by Ryle, who referred to it as “a propensity not only to make certain

---

[22] Decades before Braitenberg, such scenarios were actually implemented in the cybernetic experiments of William Grey Walter (1951).

theoretical moves but also to make certain executive and imaginative moves, as well as to have certain feelings" (Ryle 1949: 135).

DO ANIMALS HAVE BELIEFS?

In all likelihood, animals do not possess the same intellectual capacities as humans but, as Eric Margolis & Stephen Laurence (2011) aptly observe, "that doesn't mean they are as dumb as thermometers." This comment refers to the claim of some philosophers, in particular Donald Davidson (1982), who claim that the discriminatory capacities of certain animals must not be interpreted as the mastering of concepts. Let us consider some animal examples discussed by Achim Stephan (1999) on the backdrop of the proposed model.

(1) A certain species of wasps take the food for their grubs to the entrance of their burrow, check the burrow for parasites and only then drag the food into the burrow. If the food is moved away from the entrance while the animal inspects the burrow, the wasp starts with the procedure over again, obviously disregarding the results of its previous inspection. Such rigid behavior can be implemented very easily with the suggested architecture for it consists of a very few expectations and actions, only: a_1: Move the food to the entrance; a_2: inspect the burrow; a_3: drag the food into the burrow; e_1: Is the food at the entrance? If yes → a_2, no → a_1; e_1: Has the burrow been inspected? If yes → a_3, no → e_1.

(2) Certain birds seem to learn that eating a monarch butterfly causes them to vomit and will from then on avoid eating anything that looks like that butterfly, including the non-poisonous viceroy butterfly. In the suggested architecture, learning this knowledge amounts to linking an appropriate concept (i.e., the sum of expectations that make it discriminate a monarch butterfly from other potential food) with the action of flying off.[23] In agreement with Stephan, neither in this nor the previous example does the animal hold any belief, here about the edibility of butterflies; its knowledge reflects mere conditional behavior.

(3) Another bird, the piping plover, is known for its "broken wing display" whenever an intruder bird threatens the piping plover's nest. Feigning a lame wing means it distracts the intruder from the nest. Clearly this does not work when a cow threatens to incidentally step on the nest so here the bird uses another technique, i.e., fluttering in the cow's face to make it walk another way. Compared to cases 1 and 2, this bird's behavioral reservoir is more complex but only in a quantitative sense: it needs to discriminate between more intruders and situations, and link the appropriate behavior with the respective concepts. So it would still be superficial to ascribe beliefs to it.

(4) Vervet monkeys are known for having a range of alarm calls in function of different predators. The monkeys also seem to be able to detect false alarms that are issued by

---

[23] I will not go into the intricacies of discussing the details of learning here. However, it can be said that animal learning without beliefs works in analogy to evolution, i.e., by *canalized* random changes (cf. Riegler 1994). In Gregory Bateson's (1978) levels of learning, the wasps in the previous example exhibit "zero learning," which is not subject to correction, while the birds in this and the two following examples demonstrate Learning I, which is the "change in specificity of response" out of a set of alternatives.

inexperienced young members of their group or by adult monkeys that deliberately use the ensuing reaction of alarm calls to end fights. If the latter repeat their behavior, other members are increasingly unlikely to react to their false calls. While some philosophers, such as Colin Allen, whom Stephan cites in his paper, are ready to ascribe beliefs to these monkeys, I find it hard to agree that their behavior is qualitatively different from the animals in the previous cases. Again, learning would need to endow these animals with ever more complex concepts but in no cases would this require the ability to reflect on the content of their own knowledge. The expectations that trigger the flight action would need to include representations of those members who give false alarms such that each time the corresponding expectation is referred to in the "do nothing" schema, its height is increased, giving it more influence in the chain of conjunctions it is embedded in. One could perfectly imagine a Braitenberg creature performing exactly the same action as a vervet monkey and still being devoid of self-reflective mental activity.

(5) Language-related experiments with chimpanzees such as Lana and Kanzi showed that these animals are able to develop new combinations of signs for which they had no past experience. They could "form syntactically correct sentences, could recognize written symbols, could read and could complete incomplete sentences appropriately" (Bettoni 2007: 37). About Lana it could be said that "she can experience a recursive coordination of behavioural coordinations, through which she could *recursively* influence what she was experiencing" (ibid, my emphasis). Being able to make one's own knowledge structures subject to other mental processes is essentially different from the animals in cases 1 to 4. As von Glasersfeld stresses, there is a fundamental difference between the concept of an object and the concept of object permanence (Glasersfeld 1995: 70). Many animals derive the concept of an object (such as an apple) from their sensorimotor experiences. It requires the coordination of experience in many different modalities. Birds and monkeys may possess this knowledge concept (but not the wasps in the first example). To arrive at the concept of object permanence, however, the subject has to carry out a reflective abstraction and develop the *belief* "that the object in question 'exists' somewhere while it is not being experienced" (ibid).

These examples should be enough to conclude that in order to hold a belief one needs to transcend mere empirical abstractions and perform a reflective abstraction on one's own knowledge structures. Here I agree with Davidson (1982), who claims that animals cannot have beliefs. They may have excellent discriminatory capabilities that enable them to distinguish between various types of enemies but they have no second-order insight into the nature of their beliefs. However, in contrast to Davidson, this does not necessarily require language, as the ability to perform reflective abstractions precedes linguistic communication – cf. Ruth Marcus's (1995: 127) observation that animal psychologists seem to have an "obsession with linguistic behavior for defining belief" and that the "propositional attitude of believing is not confined to a disposition to verbal behavior" (ibid: 128).

THE RADICAL CONSTRUCTIVIST ACCOUNT: BELIEF ENTAILS KNOWLEDGE

Defining knowledge in terms of concepts and schemata is not only opposed to the definition of knowledge in terms of beliefs but also forces us to reverse their relation and make knowledge fundamental to beliefs (cf. Tim Williamson's 2000 proposal "knowledge first") .

Like animals, young human children may not be attributed beliefs for they lack the capacity to reflect on their knowledge. A particularly interesting experiment to this effect shows that while young children do have recollections of events and procedures they were part of, they cannot reflect on them in language when tested months later. Their verbal descriptions of an event are "frozen in time, reflecting their verbal skill at the time of encoding, rather than at the time of the test" (Simcock & Hayne 2002: 229).

These observations suggest that belief is a conscious reflection in language on fundamental elements in one's knowledge, such as comparing an expectation with the concept it is contained in. For example, if the subject's concept of "apple" contains the expectations e1 "is green" and e2 "is about 7cm", the subject can compare either expectation with the concept "apple" to arrive at the beliefs "apples are green" and "apples are about 7cm in diameter."[24] These beliefs are clearly related to propositions. But they entail knowledge: in order to have a belief you first have to have the appropriate knowledge to form a belief about.

To illustrate the difference between knowledge and belief, consider Francisco Varela's (1992: 249) example of a kingfisher that could be said to use Snell's law of refraction to determine the angle it should dive into the water to catch the fish it sees in the water. Such a belief merely reflects the perspective of the observer. If the bird was able to reflect on its own doing, it could attribute such a belief to itself. But so far it seems reasonable to assume that this is not case as the bird is not capable of self-reflection (nor is it well-versed in physics).

The distinction between unreflective and reflective knowledge was proposed by Ernest Sosa, who distinguished two general varieties of knowledge: animal knowledge about one's environment, past or experience "with little or no benefit of reflection or understanding" (Sosa 1991: 240), and reflective knowledge linked to the "understanding of its place in a wider whole that includes one's belief and knowledge of it and how these come about" (ibid). While for Sosa this distinction serves mainly as a vehicle to relieve animal knowledge of the need to be epistemically justified (see also next section), I venture to make it the criterion for telling knowledge from belief.

KNOWLEDGE-THAT AND KNOWLEDGE-HOW

In philosophy there has been a long-standing discussion about the relationship between knowing-that and knowing-how. It is widely believed that for the former we can form

[24] The fact that such an expectation contains the aspects of both extension and color should suffice to show that the radical constructivist account could not be bothered less with the primary/secondary quality distinction.

epistemic justifications.[25] For example, we have good empirical reasons to assume that apples are green. But many harbor doubts about whether we can do this for the latter:

> "An intelligently performed action need not be preceded and powered, let alone at all constituted, by consideration of some proposition, let alone by knowledge of the proposition's being true." (Hetherington 2006: 71)

and thus declare them independent from each other. Such anti-intellectualism is contrasted with the intellectualist view that "knowing-how is reducible to a set of knowings-that" (Ryle 2009: 230) or radical anti-intellectualism, which reverses the dependency: "Effective possession of a piece of knowledge-that involves knowing how to use that knowledge, when required, for the solution of other theoretical or practical problems" (ibid: 235).

There are good intuitive reasons to oppose intellectualism: "Why assume that bicycle riding, for example, is a symbol-processing task?" (Touretzky & Pomerleau 1994: 351). When teaching others, we can only rudimentarily instruct the learner in what to do, issuing propositional commands such as "push the pedals" and "keep your balance."

So it seems appropriate to accept Stephen Hetherington's "knowledge-as-ability," which equates the ability to do something with the knowledge-how to do it: "knowledge that p is the ability – the know-how – to respond, to reply, to represent, or to reason accurately that p" (Hetherington 2006: 77). On this view, knowledge-that is no longer an intellectual relation to a proposition but rather an ability (Fantl 2012) and hence integrates with knowing-how.

Also the proposed computational model provides an integrative theory of knowledge, which rejects the idea that higher-level intelligence is best expressed in symbolic form while for lower-level intelligence, a connectionist approach is best suited, i.e.,

> "a model of cognition in which conscious, deliberate symbol manipulation is the top-level, sub-conscious symbol processing the intermediate level, and specialized nonsymbolic modules appear at the lowest level." (Touretzky & Pomerleau 1994: 351)

It is difficult to see why such a tripartite structure should be advantageous. Quite on the contrary, it is far from obvious how the boundaries between these levels should be defined: When does conscious cognition become unconscious? etc.

The present model, by contrast, not only dispenses with any such separation, it also makes the intellectualism/anti-intellectualism debate superfluous, for knowledge-that and knowledge-how are just ends of a continuum rather than qualitatively different concepts. This continuum allows us to compare how learning knowledge-that and learning knowledge-how relate to each other.

Learning an ability such as riding the bicycle is the continuous transition from knowing-that (i.e., propositional instructions such as "push the pedals", "keep your balance" etc.) to knowing-how (i.e., everything that is necessary for staying in the saddle and moving forward). This is accomplished by adding expectations to the action sequence that

---

[25] "Epistemic justification is traditionally associated with being able to generate reasons in defense of one's beliefs, but in many instances of knowledge, one does not seem to be in a position to provide anything like a defense of one's beliefs" (Foley 2004: 59).

carries out bicycle riding in order to fine-tune ones movements on the bicycle. In other words, learning how to do something amounts to extending the initial coarse schema – which, for example, was inferred from the propositions in a textbook or the verbal instructions of a peer – by adding expectations that smooth the execution of the schema. On these assumptions, it follows that at some point the knowledge pertaining to the ability to ride a bicycle becomes so intricate and hierarchically nested (and here lies a fundamental difference between following simple rules of the sort Ryle (1971) thought of and traversing through the structure of nested expectations) that any effort to form a belief by reflecting on ("introspecting") the components of schemata becomes an impossible cognitive task. In other words, we can no longer offer epistemic justifications for our doing because it hides behind an impenetrable layer of complexity.

Learning a proposition or a fact in general goes the opposite direction. The learner starts with actions that manifest themselves in a propositional concept. Ernst Mach argues that a chemist can recognize sodium "on the presupposition that a definite number of tests which he has in mind would give the results which he expects" (Mach 1986: 380). An untrained chemist may start with some simple schemata consisting of a single expectation (such as "does it have a silver sheen?"). In the course of her formative dealing with the substance, the action sequence of the schema will get increasingly more complex by adding actions that present the chemist with the required perception, the veracity of which needs to be fulfilled in a subsequent expectation before the next action changes the perceptual context again for the next expectation. In the proposed model, such learning is accomplished by adding expectations that leads to the ever more complete insight that one is dealing with a piece of sodium: "The concept 'sodium,' accordingly, is made up of a certain series of sensory characteristics which make their appearance upon the performance of certain definite manual, instrumental, and technical operations which may be very complicated in character" (ibid).

Some of the context-changing actions are older than others, as they were acquired before for other purposes. Stored in libraries, they represent entrenched structures in the cognitive apparatus, i.e., structures that if they were changed would fundamentally change many schemata that refer to them. In a sense this is what Quine attributed to the web of beliefs, which, too, is heterogeneous in this respect: "some claims are more central to our conceptual network and unlikely to be shifted in response to any specific unexpected sensory input" (Godfrey-Smith 2014: 55f). Hence these older elements are less prone to ever being changed and their entrenchment constitutes what we refer to as "habits of thought" in common language. However, the differences amount to a difference in degree rather than in kind and depend on the number of times an element is referred to in another element. Since a human being first constructs the objects that populate her experiential reality before she goes on to increasingly more sophisticated activities such as coping with social relations and mathematical problem solving, the concepts pertaining to the "reality of objects" and the "existence of objective truth" are more deeply entrenched than those relating to the reality of human relationships and scientific problems (Riegler 2001). While the latter can be modified (for example after psychotherapeutic intervention such as "habitual reframing"), the former stand out as anchor points in the network of schemata: "The most violent revolutions in an individual's [concepts] leave most of his old order standing. Time and space, cause and effect, nature and history, and one's own biography remain untouched" (James 1987: 513; see also footnote 17).[26]

CONCLUSION

It has become part of our language to speak of a knowledge-based society as if knowledge was that which makes humans different from animals. On what I think is a plausible reading of my arguments this can hardly be a decisive distinction, for knowledge amounts to mere abstractions from experience. That human science is a fine example of such empirical abstractions is beyond dispute. However, it seems appropriate to emphasize that the crucial single distinction that makes us human is our ability to generate reflective abstractions from "material actually found in experience or […] in a thought experiment with imaginary material" (Glasersfeld 1995: 70) – and these reflections are beliefs, which makes society *belief*-based, i.e., a belief system.

My argumentation has also built on the suggestion that many philosophical concepts should no longer be treated as incomparable qualities but as quantities on continua instead. I made the case for knowing-that and knowing-how. History of science tells us that physics made a major breakthrough by no longer treating "warm" and "cold" as qualities but rather as quantities on a temperature scale (Crosby 1997: 66f). Such a quantitative methodology will make the customary appeals to intuition more stringent, as it opens the doors to arguments based on computational structures and hence to a more precise way of seeing connections.

We are just at the beginning of explicating all the algorithmic details of a formal-computational model of knowledge and belief. One could argue, as pointed out by Eric Schwitzgebel in personal communication, that to model fully the massive complexity of human knowledge and agency would require an intractable nest of formal models. In practice this seems impossible and one would need to resort to simplification. This is certainly correct. But then, again, it is a question of quantitative difference: How much worse is the situation for a philosophy that relies on intuition alone, where implications of claims must be "manually" checked? With computer models, philosophers have the luxury of letting a machine think through the implications much faster and more reliably.

Accordingly, this paper expresses the conviction that philosophy could become more encompassing by amending its appeals to intuition with computational modeling, i.e., by turning to a methodological pluralism in which intuitions play the crucial role of connecting philosophy with the human.

---

[26] One could assume this to be the reason why some philosophers may find it easier to accept the realist stance than the anti-realist one.

ACKNOWLEDGEMENT

My gratitude goes to Eric Schwitzgebel and an anonymous reviewer who took the time to assess my paper critically and to provide me with helpful comments.

REFERENCES

Albert H. (1968) Traktat über kritische Vernunft. J. C. B. Mohr, Tübingen. English translation: Albert H. (1985) Treatise on critical reason. Princeton University Press, Princeton NJ.

Baldwin M. (1895) Consciousness and evolution. Science 2(34): 219–223.

Bateson G. (1978) Steps to an ecology of mind: Collected essays in anthropology, psychiatry, evolution and epistemology. Granada Publishing, London. Originally published in 1972.

Bettoni M. C. (2007) The Yerkish language: From operational methodology to chimpanzee communication. Constructivist Foundations 2(2–3): 32–38. Available at http://www.univie.ac.at/constructivism/journal/2/2–3/032.bettoni

Blanchette I. (2006) The effect of emotion on interpretation and logic in a conditional reasoning task. Memory & Cognition 34(5): 1112–1125.

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

Cariani P. (2012) Infinity and the observer: Radical constructivism and the foundations of mathematics. Constructivist Foundations 7(2): 116–125. Available at http://www.univie.ac.at/constructivism/journal/7/2/116.cariani

Churchland P. (1984) Matter and consciousness. MIT Press, Cambridge MA.

Crosby A. W. (1997) The measure of reality. Cambridge University Press, Cambridge.

Davidson D. (1982) Rational animals. Dialectica 36: 317–327.

Davidson D. (1999) The emergence of thought. Erkenntnis 51: 7–17.

Dörner D. (1996) The logic of failure. Metropolitan Books, New York.

Dretske F. (1969) Seeing and knowing. Routledge and Kegan Paul, London.

Dykstra Jr. D. (2005) Against realist instruction. Constructivist Foundations 1(1): 49–60. Available at http://www.univie.ac.at/constructivism/journal/1/1/049.dykstra

Fantl J. (2012) Knowledge how. In: Zalta E. N. (ed.) Stanford Encyclopedia of Philosophy. Available at http://plato.stanford.edu/entries/knowledge-how/

Feynman R. (1985) Surely you're joking, Mr. Feynman! W. W. Norton, New York.

Foerster H. von (1973) On constructing a reality. In: Preiser F. E. (ed.) Environmental Design Research Volume 2. Dowden, Hutchinson & Ross, Stroudberg: 35–46. Reprinted in: Foerster H. von (2003) Understanding understanding. Springer, New York: 211–228.

Foerster H. von (2003) For Niklas Luhmann: "How recursive is communication?" In: Foerster H. von, Understanding understanding. Springer, New York: 305–323.

Foley R. (2004) A trial separation between the theory of knowledge and the theory of justified belief. In: Greco J. (ed.) Ernest Sosa and his critics. Blackwell, Malden MA: 59–72.

Füllsack M. (2013) Constructivism and computation: Can computer-based modeling add to the case for constructivism? Constructivist Foundations 9(1): 7–16. Available at http://www.univie.ac.at/constructivism/journal/9/1/007.fuellsack

Gärdenfors P. (1988) Knowledge in flux. MIT Press, Cambridge MA.

Gettier E. L. (1963) Is justified true belief knowledge. Analysis 23: 121–123.

Glasersfeld E. von (1974) Piaget and the radical constructivist epistemology. In: Smock C. D. & Glasersfeld E. von (eds.) Epistemology and education. Follow Through Publications, Athens GA: 1–24. Available at http://www.univie.ac.at/constructivism/EvG/papers/034.pdf
Glasersfeld E. von (1983) On the concept of interpretation. Poetics 12(2/3): 207–218. Available at http://www.univie.ac.at/constructivism/EvG/papers/082.pdf
Glasersfeld E. von (1988) The reluctance to change a way of thinking. Irish Journal of Psychology 9(1): 83–90. Available at http://www.vonglasersfeld.com/111
Glasersfeld E. von (1995) Radical constructivism. Falmer Press, London.
Glasersfeld E. von (2006) A constructivist approach to experiential foundations of mathematical concepts revisited. Constructivist Foundations 1(2): 61–72. Available at http://www.univie.ac.at/constructivism/journal/1/2/061.glasersfeld
Glasersfeld E. von (2008) Who conceives of society? Constructivist Foundations 3(2): 59–64 & 100–104. Available at http://www.univie.ac.at/constructivism/journal/
Glasersfeld E. von (2009) Relativism, fascism, and the question of ethics in constructivism. Constructivist Foundations 4(3): 117–120. Available at http://www.univie.ac.at/constructivism/journal/4/3/117.glasersfeld
Glasersfeld E. von & Ackermann E. K. (2011) Reflections on the concept of experience and the role of consciousness. Unfinished fragments. Constructivist Foundations 6(2): 193–203. Available at http://www.univie.ac.at/constructivism/journal/6/2/193.glasersfeld
Godfrey-Smith P. (2014) Quine and pragmatism. In: Harman G. & LePore E. (eds.) A Companion to W. V. O. Quine. Wiley, Chichester UK: 54–68.
Goldman A. I. (1979) What is justified belief? In" Pappas G. S. (ed.) Justification and knowledge: New studies in epistemology. Reidel, Dordrecht: 1–23.
Hendricks V. F. (2006) Mainstream and formal epistemology. Cambridge University Press, New York.
Hetherington S. (2006) How to know (that knowledge-that is knowledge-how). In: Hetherington S. (ed.) Epistemology futures. Oxford University Press, Oxford: 71–94.
Hetherington S. (2012) Knowledge. Internet Encyclopedia of Philosophy. Available at http://www.iep.utm.edu/knowledg/
Himmelfreundpointner T. (2013) Wittgenstein, Rorty and Mitterer: On aspects and descriptions. Constructivist Foundations 8(2): 210–215. Available at http://www.univie.ac.at/constructivism/journal/8/2/210.himmelfreundpointner
James W. (1890) Principles of psychology. Volume 2. Holt, New York.
James W. (1897) The will to believe, and other essays in popular philosophy. Longmans, Green, and Co., New York.
James W. (1987) Pragmatism: A new name for some old ways of thinking. In: Writings 1902–1910, edited by Bruce Kuklick. Library of America, New York: 479–624. Originally published in 1907.
Keil F. C. (2003) Folkscience: Coarse interpretations of a complex reality. Trends in Cognitive Sciences 7: 368–373.
Latour B. (1987) Science in action. Harvard University Press, Cambridge.
Ludwig K. (2006) Belief. In: Borchert D. M. (ed.) Encyclopedia of philosophy. Second edition. Thomson Gale, Farmington Hills MI: 532–535.
Mach E. (1986) Principles of the theory of heat. Historically and critically elucidated. Translated from the German second edition of 1900. Reidel, Dordrecht.
Marcus R. B. (1995) The anti-naturalism of some language centered accounts of belief. Dialectica 49(2–4): 113–129
Margolis E. & Laurence S. (2011) Concepts. In: Zalta E. N. (ed.) Stanford Encyclopedia of Philosophy. Available at http://plato.stanford.edu/entries/concepts/
Maturana H. R. (1974) Cognitive strategies. In: Foerster H. von (ed.) Cybernetics of cybernetics. Biological Computer Lab (BCL) Report 73–38, dated May 1974. University of Illinois, Urbana: 457–469.

Mitterer J. (1992) Das Jenseits der Philosopie. Wider das dualistische Erkenntnisprinzip [The beyond of philosophy. Against the dualistic principle of cognition]. Passagen Verlag, Vienna. Reprinted in 2011 by Velbrück, Weilerswist.
Myers-Schulz B. & Schwitzgebel E. (2013) Knowing that p without believing that p. Noûs 47(2): 371–384
Neisser U. (1976) Cognition and reality. W. H. Freemanm San Francisco.
Newell A. & Simon H. A. (1976) Computer science as empirical inquiry: Symbols and search. Communications of the ACM 19(3): 113–126.
Perlis D. (2000) The role(s) of belief in AI. Minker J. (ed.) Logic-based AI. Kluwer: 361–374
Pfeifer N. & Douven I. (2014) Formal epistemology and the new paradigm psychology of reasoning. Review of Philosophy and Psychology 5(2):1 99–221.
Porr B. & Wörgötter F. (2005) Inside embodiment: What means embodiment to radical constructivists? Kybernetes 34: 105–117.
Powers W. T. (1973) Behavior. The control of perception. Aldine de Gruyter, New York.
Quale A. (2014) Ethics: A radical-constructivist approach. Constructivist Foundations 9(2): 256–261. Available at http://www.univie.ac.at/constructivism/journal/9/2/256.quale
Quine W. V. (1951) Two dogmas of empiricism. Philosophical Review 60: 20–43.
Quine W. V. (1981) Things and their place in theories. In: Theories and things. Belknap Press, Cambridge MA: 1–23.
Quine W. V. & Ullian J. (1978) The web of belief. Second edition. McGraw-Hill, New York.
Raskin J. D. (2011) Ernst von Glasersfeld and psychotherapeutic change. Constructivist Foundations 6(2): 235–238. Available at http://www.univie.ac.at/constructivism/journal/6/2/235.raskin
Riegler A. (1994) Constructivist artificial life: The constructivist-anticipatory principle and functional coupling. In: Hopf J. (ed.), Genetic algorithms within the framework of evolutionary computation. MPI-I-94–241. Max-Planck-Institut fur Informatik, Saarbrücken: 73–83.
Riegler A. (2001) Towards a radical constructivist understanding of science. Foundations of Science 6: 1–30. Available at http://www.univie.ac.at/constructivism/riegler/20
Riegler A. & Douven I. (2009) Extending the Hegselmann–Krause model III: From single beliefs to complex belief states. Episteme 6(2): 145–163. Available at http://www.univie.ac.at/constructivism/riegler/56
Riegler A., Stewart J. & Ziemke T. (2013) Computation, cognition and constructivism: Introduction to the special issue. Constructivist Foundations 9(1): 1–6. Available at http://www.univie.ac.at/constructivism/journal/9/1/001.riegler
Rosch E. (1978) Principles of categorization. In: Rosch E. & Lloyd B. (eds.) Cognition and categorization. Lawrence Erlbaum Associates: Hillsdale NJ: 27–48.
Russell B. (2005) The analysis of mind. Dover, Mineola NY. Originally published in 1921.
Ryle G. (1949) The concept of mind. University of Chicago Press, Chicago.
Ryle G. (2009) Knowing how and knowing that. In: Collected papers. Volume 2 Collected Essays 1929–1968. Routledge, New York: 222–235. Originally published in 1946.
Scholl A. (2010) Radical constructivism in communication science. Constructivist Foundations 6(1): 51–57. Available at http://www.univie.ac.at/constructivism/journal/6/1/051.scholl
Simcock G. & Hayne H. (2002) Breaking the barrier? Children fail to translate their preverbal memories into language. Psychological Science 13(3): 225–231.
Sosa E. (1991) Knowledge in perspective. Cambridge University Press, Cambridge: 131–145.

Stephan A. (1999) Are animals capable of concepts. Erkenntnis 51(1): 583–596.
Stove D. (1991) The Plato cult and other philosophical follies. Basil Blackwell, Oxford.
Thagard P. (1988) Computational philosophy of science. MIT Press, Cambridge MA.
Touretzky D. S. & Pomerleau D. A. (1994) Reconstructing physical symbol systems. Cognitive Science 18: 345–353.
Ulrich C., Tillema E. S., Hackenberg A. J. & Norton A. (2014) Constructivist model building: Empirical examples from mathematics education. Constructivist Foundations 9(3): 328–339. Available at http://www.univie.ac.at/constructivism/journal/9/3/328.ulrich
Vaihinger H. (1913) Die Philosophie des Als Ob [The philosophy of as if]. Second edition. Reuther & Reichard, Berlin.
Varela F. J. (1992) Whence perceptual meaning? A cartography of current ideas. In: Varela F. J. & Dupuy J. (eds.) Understanding origins. Contemporary views on the origin of life, mind and society. Kluwer, Dordrecht: 235–263.
Varela F. J., Thompson E. & Rosch E. (1991) The embodied mind: Cognitive science and human experience. MIT Press: Cambridge MA.
Walter W. G. (1951) A machine that learns. Scientific American 185(2): 60–63.
Weinberg J., Nichols S. & Stich S. (2001) Normativity and epistemic intuitions. Philosophical Topics 29: 429–460.
Williamson T. (2000) Knowledge and its limits. Oxford University Press, Oxford.
Wittgenstein L. (1922) Tractatus logico-philosophicus. Translated by C. K. Ogden. Routledge, London.
Wittgenstein L. (1953) Philosophical investigations. Basil Blackwell, Oxford.
Wittgenstein L. (1993) Philosophical occasions. Hackett, Indianapolis.
Wolpert L. (1992) The unnatural nature of science. Faber and Faber, London.