

# Are Sensory Concepts Learned by “Abstraction” from Experience?

Pär Sundström<sup>1</sup> 

Received: 18 August 2017 / Accepted: 7 April 2018 / Published online: 17 April 2018  
© The Author(s) 2018

**Abstract** In recent years, many philosophers and scientists have argued or accepted that it is impossible to learn primitive sensory concepts like “blue” and “red”. This paper defends a more qualified picture. I try to show that some received characterisations of “learning” are nonequivalent and point towards different learning-nonlearning distinctions. And, on some ways of specifying such a distinction, it might be correct that we do not and cannot “learn” a concept of blue. But on other ways of specifying such a distinction, we can and do sometimes “learn” a concept of blue from experiences of blue. The latter part of the argument connects with some traditional “abstractionist” views, and I defend the present claims in view of some widely circulated concerns about “abstracting” concepts from experience. I close with some reflections on how one might, in view of all this, think about “the learning-nonlearning distinction”.

## 1 Introduction

As I will use terms, to “have a concept” of something is to be able to think about that thing. To have a concept of horses is to be able to think about horses; to lack a concept of horses is to lack that ability.<sup>1</sup>

The question for this paper is whether one can learn sensory concepts from experiences of that to which these concepts apply; for example, whether one can

---

<sup>1</sup> Section 2.1 below specifies this a bit further.

---

✉ Pär Sundström  
par.sundstrom@umu.se

<sup>1</sup> Department of Historical, Philosophical and Religious Studies, Umeå University, 901 87, Umeå, Sweden

learn to think about blue from experiences of blue. It may be natural to suppose that this is not just possible but that we humans in fact sometimes do learn sensory concepts this way. However, the view that we do not ever learn a concept of blue from experiences of blue, and that this is not even possible, has in recent years been widely accepted in science and philosophy. The main root of this state of the art is a line of argument by Jerry Fodor (in, among other places, 1975, chap. 2, 1981, 2008, chap. 5).

Fodor defends to begin with this claim:

*PrimitiveNotLearned*: It is impossible to learn a primitive concept.

That a concept is *primitive* means for Fodor that it lacks analysis. For example, the concept “white horse” can plausibly be analysed in terms of the components “white” and “horse”. It is more controversial whether either of the concepts “horse” or “white” can be analysed in terms of more primitive components. If they cannot they are “primitive”.

To *learn* a concept is according to Fodor *one way* of coming to possess it; but just one. As Fodor uses terms, there is both (a) concept acquisition (and cognitive development more generally) that amounts to learning, and (b) concept acquisition (and other cognitive development) that does *not* amount to learning. Here is one passage that draws this distinction:

...concept learning is some sort of rational process. It contrasts sharply with kinds of concept acquisition where, for example, a concept is acquired by surgical implantation; or by swallowing a pill; or by hitting one’s head against a hard surface, etc. Intuitively, none of these is concept *learning*, but any or all of them might eventuate in concept attainment (2008, 135).

This passage appeals to an “intuitive” separation between concept *learning* and concept acquisition that is not learning. It also characterises concept learning as a “rational” process. In other passages, Fodor characterises concept learning as “computational”, “intentional”, “cognitive” and “in the domain of cognitive psychology” (1975, 93, 2008, 136–7, 139, 151–2), and concept acquisition that is not learning as “automatic”, “noninferential”, jumpy,<sup>2</sup> “brute-causal” (1981, 273), and as “merely triggered” by circumstances (1981, 275). These different characterisations will be central in what follows.

In addition to *PrimitiveNotLearned*, Fodor defends:

*ManyPrimitive*: Many of our concepts are primitive.

*ManyPrimitive* is vague, but the idea is that primitive concepts include sensory concepts like “blue” and “red” and much more besides; Fodor mentions in these contexts concepts like “horse”, “tree”, “carburettor”, and “umbrella” (2008, 129).

The conjunction of *PrimitiveNotLearned* and *ManyPrimitive* entails:

<sup>2</sup> The exact words, placed in the mouth of an interlocutor in a dialogue, are: “You say that we jump, by some or other ‘automatic’ process ... to our concepts. And you also say that these jumps don’t count as inferences” (2008, 164).

*ManyNotLearned*: Many of our concepts are such that it is impossible to learn them.

This conclusion, *ManyNotLearned*, is widely contested. However, critics of it have typically objected specifically to the idea that it is impossible to learn concepts like “horse” and “carburettor”. Many of them have accepted Fodor’s view that it is impossible to learn primitive sensory concepts, like “blue” and “red”.<sup>3</sup>

Such theorists come in two types. The first type is exemplified by Stephen Pinker. Pinker maintains that, contrary to *ManyPrimitive*, few of our concepts are primitive, and “carburettor”, for example, is not among these (2007, 107). But Pinker accepts that there are some primitive concepts, and that primitive concepts are innate (and therefore, presumably, not learned), and he thinks that “red” is “patently” among those (*ibid.*, 92–3).<sup>4</sup>

A second type of theorist accepts *ManyPrimitive* but urges, contrary to *PrimitiveNotLearned*, that primitive concepts come in two different types: *Some* primitive concepts, like “horse” and “carburettor”, can be learned. But others cannot. And the learning of the former builds upon the latter. This kind of view is held by Eric Margolis (1998), Stephen Laurence and Margolis (2002), Margolis and Laurence (2011), Susan Carey (2009, 2014, 2015), Daniel Weiskopf (2008), and Kim Sterelny (1989), all of whom more-or-less clearly suggest that the class of primitive concepts that cannot be learned includes sensory concepts like “blue”.<sup>5</sup>

<sup>3</sup> There is arguably no redundancy in ‘primitive sensory concept’. For example, “blue-and-red-striped” is a concept that is sensory but plausibly analysable into the components “red”, “blue” and “striped” and hence not primitive. The focus here will be on *primitive* sensory concepts but I will not always repeat the qualifier ‘primitive’.

<sup>4</sup> Jackendoff (1989) pursues a similar objection to Fodor’s argument, but is less explicit about which concepts he takes to be primitive and innate.

<sup>5</sup> To give the flavour of one such view: Laurence and Margolis develop a model of concept learning focussing on natural kind concepts like “dog” and “water”, which they take to be primitive. In their model, such concepts are learned on the basis of “sustaining mechanisms”, one type of which—“Perhaps the most important type” (2002, 38)—“implicates a *kind syndrome*”, which is “a collection of properties that is highly indicative of a kind yet is accessible in perceptual encounters. This may include things like the typical shape, motions, markings, sounds, colors, etc., associated with a kind” (*ibid.*). And Laurence and Margolis at least strongly suggest that, in order to learn a natural kind concept with the aid of this kind of sustaining mechanism, a subject needs to have *concepts* of the elements of the “kind syndrome”. For example, Margolis (1998) says that a subject who acquires a concept this way “has to *know* a collection of salient, relatively accessible properties that are highly indicative of the kind” (359; italics added).

There are other theorists who argue that sensory concepts cannot be learned but where Fodor’s argument is not a salient background for their view; see for example Aydede and Güzeldere (2005, sect. 3) and McGinn (2013).

A recent paper by Laurence and Margolis (2012) makes up a slight and somewhat hedged exception to the widespread acceptance of the claim that it is impossible to learn primitive sensory concepts. Laurence and Margolis argue that *generic* sensory concepts, like the concept of “white” that applies to many different shades of white, (a) can be learned (via a process they call ‘abstraction’), and (b) are perhaps primitive (for the perhaps, see sect. 4.2). Meanwhile, they deny that such an account can be extended to concepts of specific sensory qualities; these, they argue, must be “innate” (13). Laurence and Margolis’ paper explores issues closely related to the present one. For reasons of space I will not spell out all the differences and similarities between the two, but some important differences and similarities will be footnoted along the way. Another paper that explores terrain nearby the present is Samet (1986), which criticises Fodor’s views with a focus on sensory concepts. However, Samet ultimately does not challenge

Thus, theorists of both types agree with Fodor that primitive sensory concepts like “red” and “blue” cannot be learned. They disagree with Fodor in maintaining that “carburettor” can be learned. And they disagree with each other about the *reason why* “carburettor” can be learned: the former theorist maintains that “carburettor” can be learned because it is nonprimitive; the latter maintains that “carburettor” can be learned despite being primitive.

I shall here leave aside concepts like “carburettor” and focus solely on whether primitive *sensory* concepts can be learned.

I will accept some established frames of the debate. To begin with I will—at least up until the last few paragraphs of the paper—rather uncritically use Fodor’s characterisations, cited above, as guides as to whether a given process of cognitive development amounts to learning. I think these characterisations are widely accepted. They are not questioned by any of the above-mentioned critics of Fodor’s argument, even though they are often highlighted.<sup>6</sup> And they are sometimes echoed by others.<sup>7</sup> I shall also, at one juncture, rely on an argument by Fodor for “computationalism” about the mind. This argument too is, I think, widely accepted and even more widely not-questioned both in this debate and beyond.

I will try to make plausible that sensory concepts are sometimes *acquired* from experience. I will also try to make plausible that this acquisition process has some but not all of the characteristics associated by Fodor—and others—with “learning”. And therefore, whether it is correct to say that this acquisition process amounts to learning depends on the exact precisification of the learning–nonlearning distinction.

It is an interesting possibility that the case of experience-based sensory concept acquisition is exceptional. It may well be that processes of cognitive development typically have all or none of the characteristics standardly associated with “learning”, and therefore fall more clearly on one side or the other of the learning–nonlearning divide. For example, Sect. 3 below discusses a case that arguably has all the characteristics standardly associated with “learning” (like being “rational”, “computational” or “cognitive”). The typical development of our mnemonic abilities arguably has none of them. I shall not defend these suggestions; I just note the possibility that experience-based sensory concept acquisition is an interestingly special case where a set of characteristics that are typically all present together or all absent together fails to be so. And, anyway, whether it is exceptional or not, I think this case, where a set of characteristics standardly associated with “learning” are (a) not all present and (b) not all absent provides an illuminating perspective on the distinction between learning and nonlearning.

---

Footnote 5 continued

any part of Fodor’s argument that is in focus here. Samet concedes that primitive sensory concepts cannot be learned, and challenges only Fodor’s supposition that *nonlearned* entails *innate*.

<sup>6</sup> For example, Laurence and Margolis highlight Fodor’s characterisation of concept learning as “rational” (2002, 29) and do not take issue with that characterisation.

<sup>7</sup> For example, Carey (2009, 12) characterises learning processes—she may have in mind *concept learning* processes specifically—as processes that “build representations of the world on the basis of computations of input that is itself representational”, echoing Fodor’s characterisation of learning processes as computational.

I shall give the name ‘abstraction\*’ to the process by which primitive sensory concepts can be, and are sometimes, acquired and in a sense “learned” from experience. I choose the ‘abstraction’-part because the process has some characteristics that have historically been associated with that label. I choose the ‘\*’-part because the process has some characteristics that have historically not been associated with ‘abstraction’. The historical connections are not central to my interests here, but I shall make a few remarks along the way about how the present picture compares with some historically held views.

My ambitions are limited. I mentioned some frames of the debate that I shall accept. I will also not argue for anything stronger than that primitive sensory concepts can be and sometimes are acquired and in a sense learned from experience. I will, for example, not argue for any of the following claims: (1) that sensory concepts can be acquired from experience without the prior possession of other concepts; (2) that sensory concepts in some sense form the “basis” of our acquisition of all our other concepts; or (3) that all our concepts are primitive sensory concepts or combinations of such. All the central claims to be made will be compatible with each of these claims. Some of the below may also be of use for someone who wishes to defend either of these claims. But it is not part of my present aim to defend any one of them.<sup>8</sup>

Section 2 specifies some central terms. Section 3 presents Fodor’s main argument for *PrimitiveNotLearned*; aspects of this argument will serve as an important foil in later sections. Sections 4 and 5 contain the core of the present argument. Section 4 argues that we can *acquire* a concept of blue from experiences of blue, and Sect. 5 argues that, drawing on some received ways of characterising “learning” and “nonlearning”, this concept acquisition is an instance of learning. Section 6 addresses an objection deriving from Geach and Sect. 7 a worry concerning attention and circularity. Section 8 summarises and spells out four different ways of thinking about “the learning-nonlearning distinction” in view of the preceding.

## 2 Terminology

### 2.1 “Having a Concept”

I said that to “have a concept” of something is to *be able to think* about that thing. That was a bit imprecise. More precisely, to have a concept of something at a time *t* is to be able to think about that thing at *t* *without first doing something else and without anything first happening to one*. To illustrate both the content and the naturalness of this “immediacy” restriction: There may be some mathematical property, *P*, that I can think about if—but only if—I first go through some mathematical training. If in my current state I am able go through that mathematical

<sup>8</sup> See Prinz 2002 for a recent defence of something like (3), Chalmers 2005 for a view that is, I think, not so remote from (2), and Locke 1689, Hume 1739, Russell 1911–1912, and Price 1953 for earlier views in this neighbourhood.

training there is a sense in which it is now correct to say that I “am able” to think about P: I am able to think about it in a sense that a cat, say, is not. But I am not in the present “immediate” sense able to think about P, because I am not able to do it without first doing something else. And therefore I do not now “have a concept” of P in the present sense of the phrase. Similarly, there may be something that I can think about if—but only if—I first have a certain experience. If I am now able to have that experience, there is a sense in which I am now able to think about that thing. But I do not have the thinking ability in the present “immediate” sense since I can employ it only if I first have the experience. And therefore I do not now “have a concept” of that thing, in the present sense.<sup>9</sup>

This specifies what it is to *be able* to think about something. It does not specify what it is to be able to *think*. I shall be minimally committal about this. I assume that there is a familiar and ordinary use of the term ‘thinking’ that clearly enough points towards a relatively natural cognitive kind that different philosophers and psychologists have offered competing or supplementing theories of. I have in mind the use of ‘thinking’ in sentences like, ‘I am now thinking about Tutankhamun’ and ‘the neighbour’s dog cannot think about Tutankhamun’. About this cognitive kind, some have held that nonhuman animals do not instantiate it (Davidson 1975) while others have held that some such animals do (Bermudez 2003; Carruthers 2004). Some have held that its instantiation in us involves the instantiation of sensory “ideas” (Hume 1739); others that it involves nonsensory symbols in a language of thought (Fodor 1975); and yet others that it involves acceptance of certain inferences or transitions (Peacocke 1992). To “have a concept”, as that phrase is understood here, is to be able to engage in *this* kind of cognitive act, however that kind should ultimately be understood.<sup>10</sup>

From this understanding of ‘having a concept’, there is a straightforward extension to related phrases. Thus, to “*acquire* a concept” of something is to acquire an ability to think about that thing, to “*learn* a concept” is to learn an ability to think about something, and so on. I believe, and will assume, that the just-explained understanding of such phrases agrees with how they are understood in all the literature discussed in this paper.<sup>11</sup>

<sup>9</sup> There is a close connection, I believe, between the present “immediate” sense of ‘*being able* to think’ and “basic action” as understood by Danto (1963, 1965): I believe it is right to say that a subject S “is able” to think about x at a time t in the present sense if and only if thinking about x is part of S’s repertoire of “basic actions” at t.

<sup>10</sup> I will not worry about the distinction between being able to think about something “*de re*” and “*de dicto*”. When I say that someone is able to think about something, x, I mean that the subject is able to think about x *either* “*de re*” or “*de dicto*”, on *any* way of understanding that distinction (for discussions of different such ways, see Hawthorne and Manley 2012 and the essays in Jeshion 2010).

<sup>11</sup> The present understanding of these *phrases* is compatible with different ways of understanding the term ‘concept’. To illustrate, take two theorists, one of whom uses ‘concept’ to mean *primitive symbol in the language of thought* while the other uses ‘concept’ to mean *primitive Fregean Sinn*. Despite their different understandings of the term ‘concept’ (and perhaps different theories about what thinking something involves), these two theorists may well agree in understanding the phrase ‘having a concept’ to be equivalent with ‘being able to think something’. I am inclined to think that Fodor (1975) and Peacocke (1992) exemplify this situation.

Philosophers often highlight the possibility of having multiple abilities to think about one and the same thing (compare Frege 1892). But issues about the individuation of thinking abilities will be of marginal importance in what follows. I will therefore allow myself locutions that are arguably inappropriate, like ‘acquiring *the* concept “so-and-so”’. These locutions can be paraphrased in terms like, ‘acquiring *a certain, recognisable* ability to think about so-and-so’.

## 2.2 “Blue”

My question is whether one can learn “sensory concepts” from experiences of that to which they apply. We can distinguish at least two types of sensory concept: (1) concepts of *specific* sensory qualities, like the particular shade of blue that I right now experience, and (2) concepts of *generic* sensory qualities, like the concept “blue” that encompasses many different specific shades. I will argue that it is *most clearly* the case that we can acquire—and in a sense learn—concepts of *specific* sensory qualities from experiences of that to which the concepts apply. For reasons of space, I must set aside a discussion of whether this conclusion can be generalised to other types of sensory concept. Since it would clutter the exposition to keep track of different sensory quality concepts throughout, I will in many contexts use imprecise talk like “sensory concept” and “concept of blue”.

## 3 Fodor’s Argument for *PrimitiveNotLearned*

To repeat: Fodor and many others maintain that primitive concepts cannot be learned. Yet others maintain that this is true at least for primitive *sensory* concepts, like “blue”. Why might one believe either of these things? A central line of thought in Fodor’s work can be summarised thus:

- (1) If concepts are learned, they are learned by a “process of projecting and confirming hypotheses about what the things that the concept applies to have in common” (2008, 132).
- (2) But concepts—or at least primitive concepts—cannot be learned in this way.
- (3) Therefore concepts—or at least primitive concepts—cannot be learned.

This summary states two different arguments: a weaker one with a conclusion restricted to primitive concepts and a stronger one with a conclusion about concepts generally. We will get to the difference between these arguments shortly.<sup>12</sup>

Fodor develops the argument with a focus on the following “concept learning paradigm” in psychology. A test subject is provided stimuli consisting of different combinations of shapes and colours, for example red triangles, green triangles, black

<sup>12</sup> Here is a summary of the stronger argument in Fodor’s own words: “If the mechanism of concept learning is the projection and confirmation of hypotheses (and what else *could* it be), then there is a sense in which there can be no such thing as learning a new concept” (Fodor 1975, 95, see also 1981, 269–270, and 2008, 139). The text below cites a passage where the more restricted argument occurs. Note that, if the stronger argument is right then the premise *ManyPrimitive* in Sect. 1 is not needed to secure the conclusion, *ManyNotLearned*; that conclusion is entailed by the present, stronger argument alone.

squares etc. The subject's task is to figure out what it is for something to be "flurg" on the basis of data points like: "this one is flurg", "that one is not". In a given case, being "flurg" may be the same as being *green or triangular*.

It may be right—relatively uncontroversial even—that when a subject in such an experiment comes to understand what it is to be "flurg" then there is *something* that she has *learned*. It seems clear that there is something that she has come to know that she did not know before, and it may well be right that her acquisition of new knowledge is a learning process. In any case, it seems clear that she has acquired new knowledge through a *reasoning* processes of this kind: "*To be flurg cannot be the same as being red-and-triangular since the green triangle was flurg. It cannot be the same as being triangular since the green square was flurg. But the available data allow that being flurg is the same as being green-or-triangular*".

It is however debatable whether there is a *concept* that one learns—or even, typically, acquires—in this kind of experiment. *If* there is a concept that one typically acquires it is presumably the concept "flurg". But that concept is plausibly one and the same as the concept "green or triangular", and subjects in this experiment typically *can* think *green or triangular* even at the outset. And if they can do that, they do not *acquire* that concept in the course of the experiment.

If they do not learn a concept, what do the test subjects learn in this experiment? One possibility is that they learn a new linguistic *term* ('flurg') for something that they could already *think*—and, for that matter, *say* in other terms. That would be to say that subjects increase their *vocabularies* but not their mental or linguistic "expressive power".

I believe Fodor was always favourable to this conclusion: that subjects never learn—and typically do not acquire—*any concept* in the "concept learning paradigm", and that this illustrates that concepts generally cannot be learned.<sup>13</sup> He was not consistently adamant about this strong conclusion. Fodor sometimes concedes that *complex* concepts, like "flurg"/"green or triangular", might in some sense be acquired, or even learned, in the experiment. For example, in one passage he floats the idea that a concept that is at the outset available only for "hypothesis formation" comes to be available for a wider range of cognitive tasks, like reasoning, organisation of memory and perceptual identification and that there is therefore a sense in which the concept may be said to be "acquired".<sup>14</sup>

But whether or not the complex concept "flurg"/"green or triangular" may be in some sense acquired or learned in the experiment, Fodor was always adamant that this can plausibly be maintained *at most* for such complex concepts. There is no plausibility at all in maintaining it for the case of *primitive* concepts like "green". The reason is that these are concepts one needs in order to at all go through the

<sup>13</sup> On the "typically do not acquire": It is hard to rule out that some concept is sometimes acquired in this experiment. If one can in principle learn Latin by eating a pill (compare Sect. 1 above), then presumably the "concept learning experiment" could in principle cause some new thinking ability to click into place in the brain (that could be a thinking ability unrelated to the content of the experiment, like an ability to think about some mathematical property). And if this is possible in principle it may be hard to rule out that it has in fact happened, say, once in human history.

<sup>14</sup> 1981, 270–1. Fodor says he floats this idea "partly because it saves the Empiricist story from the charge of utter logical vacuity".



reasoning that the learning process consists in; they are the concepts one *reasons with* when one considers what is compatible and not compatible with one's data.

I am happy to accept that no concept—not “flurg”, nor “green”—is ever learned, and that none is typically acquired in the kind of process of “projecting and confirming hypotheses” that Fodor describes. This is to say that I am happy to accept premise (2) of Fodor's argument, both the strong and the weak version. To the extent that I shall question Fodor's view, I shall question premise (1), that this kind of projecting and confirming of hypotheses is the *only* way in which one can learn a concept. The next two sections will develop this.

#### 4 We Can and Do Sometimes Acquire a Concept of Blue from Experience

I shall assume that it is possible to *acquire* a concept of blue from experiences of blue. And I shall assume that human subjects in fact sometimes do acquire a concept of blue that way.

I think these assumptions are extremely plausible. The most convincing cases concern, I believe, acquisition of *specific* sensory quality concepts, like concepts of specific shades of colour. Take some specific shade of blue that I have never experienced, do not know whether anybody else has experienced, whether it has been named, whether anything instantiates it, or what its physical correlate is. It is plausible that, unless something first happens to me or I first do something else, I cannot now think about that specific shade. Having never experienced it I cannot think about demonstratively as “that shade”; not knowing whether it has been named I cannot think about as “the shade called so-and-so”; not knowing whether anybody else has experienced it I cannot think about it as “the shade that so-and-so experienced then and there”; not knowing whether anything instantiates it I cannot think about it as “the shade I would experience if I perceived such-and-such”; and not knowing its physical correlate I cannot think about it as “the shade with such-and-such physical correlate”. And, even if I have general grasp of colour space and its dimensions of difference and similarity, I plausibly cannot on that basis alone cognitively “determine” that particular shade and isolate it, in thought, from other nearby shades.

However, it is plausible that I could now *experience* that shade. And it is plausible that, if I did experience the shade then I could think about it by drawing on that experience. For example, as a result of having the experience I might be able to think that (as I may naturally put it) “ripe tomatoes do not have *that* shade of colour”. Moreover, it seems plausible that this kind of thing does in fact happen now and then. I believe it has happened to me for example.

Besides being plausible, the present assumptions are common ground, at least within a sociological frame that is significant for present purposes. For example, Fodor clearly accepts that we sometimes acquire a concept of blue from experiences of blue. He says:

it is not part of the Rationalist program to claim that your repertoire of primitive concepts is independent of the character of your experience. On the contrary ... there are presumably triggering experiences that are nomologically necessary for the acquisition of sensory-primitive concepts like RED (1981, 274; see also 2008, 145).

Similarly, Pinker accepts that “patently atomic” concepts, like “red”, are “triggered directly by the eyes and visual system” (2007, 93), which I take to mean that these concepts are at least sometimes *acquired* as a consequence of visual experiences. And, as we shall see in Sect. 6 below, Peter Geach, in the context of his criticism of “abstractionism”, seems also to accept that a concept of blue can be acquired from an experience of blue.<sup>15</sup>

I will on the above grounds assume that some sensory concepts—at least concepts of *specific* sensory qualities—can be and are sometimes acquired from experience.

As I said, I call the process whereby a concept of (a specific shade of) blue is acquired from experience ‘abstraction\*’. The process has some characteristics that have historically been associated with the label ‘abstraction’. For one thing, abstraction\* is a process in which a sensory concept is acquired from experience of that to which the concept applies. Abstraction\* is also, and more specifically, a process in which, typically, an ability to think about something, *x*, is in a certain sense “distilled” from an experience of *x* and *something else*. In a given case, an

<sup>15</sup> Note that what I have assumed in this section is considerably weaker than what (for example) Fodor supposes in the just-quoted passage above. Fodor supposes that certain experiences are “nomologically necessary” for the acquisition of primitive sensory concepts. (Pinker may be taken to suggest something similar.) I have assumed only that it is *possible* to acquire such a concept from experience and that we *sometimes* do so. That is compatible with a “pluralist” view according to which the relevant concepts can be—and are in fact—acquired in multiple ways. They may sometimes be acquired from experience and sometimes acquired in other ways, e.g., by picking up a linguistic term, like ‘blue’, and using it deferentially (see Kaplan 1989, 603–4; Sundström 2004; Tye 2009, sect. 3.6; Ball 2009 for defences of the latter possibility). The view that we can and sometimes do acquire abilities to think about blue from experiences of blue is, I think, very widely accepted. Consider for example the case of John McDowell. McDowell is often interpreted as, and criticised for, maintaining a “conceptualist” view that precludes that concepts are sometimes acquired from experience (for example, Heck 2000, 490–8; Peacocke 2001, sects. 2 and 3; Roskies 2008, 2010; Levine 2010, 191, all criticise McDowell for this). There are certainly passages in McDowell that encourage this interpretation (see for example 1994, 10). But the high plausibility of the claim that we can and sometimes do acquire thinking abilities from experience provides some reason to doubt that interpretation. And there are passages in McDowell that reinforce such doubt. McDowell familiarly argues (in response to Evans) that we have or can form concepts of specific shades of colour. But he concedes (to Evans) that we do not have such concepts in *advance* of, and *independently* of our colour experiences: “In the throes of an experience of the kind that putatively transcends one’s conceptual powers ... one can give linguistic expression to a concept that is exactly as fine-grained as the experience, by uttering a phrase like ‘that shade’, in which the demonstrative exploits the presence of the sample. ... What is in play here is a recognitional ability, possibly quite short-lived, that *sets in with the experience*” (56–7; italics added; see also p. 58). It is also worth noting that the present claim (that we can and sometimes do acquire an ability to think about blue from experiences of blue) amounts to no more than a limited form of “nonconceptualism” about experience. Specifically, it amounts to no more than the “state nonconceptualist” claim that one can be in a certain experiential state without having a certain conceptual ability. This claim does not—or not obviously—entail or presuppose the “content nonconceptualist” claim that the contents of experiences differ from the contents of cognition (compare Speaks 2005, 2015, chap. 11, and Byrne 2005, sect. 1).

ability to think about a specific shade of blue may be acquired from an experience of something—a ball, say—as having that specific shade *and such-and-such shape*.

But abstraction\* also has some characteristics that have not historically been associated with ‘abstraction’. For example, abstraction\* is a process in which a concept can be acquired from a *single* experience. This is in contrast with for example Russell’s idea that “abstraction” builds upon many experiences:

When we see a white patch, we are acquainted, in the first instance, with the particular patch; but by seeing many white patches, we easily learn to abstract the whiteness which they all have in common, and in learning to do this we are learning to be acquainted with whiteness (1912, 101).

Further, abstraction\* is a process by which concepts of *specific* shades of colour can be acquired. By contrast, it is natural to understand Russell in the quote above as having in mind a process that outputs “acquaintance” with *generic* whiteness (*on the basis* of acquaintance with “patches” that exemplify different shades of white).<sup>16</sup> Note 18 below registers yet another difference between abstraction\* and abstraction as traditionally understood.

## 5 Is Acquiring a Concept of Blue from Experience an Instance of Learning?

Given that one can, and that we sometimes do, *acquire* a concept of blue from experience, the central question for us is whether this process, abstraction\*, is or is not a matter of *learning* the relevant concepts. The present guidelines on this issue will, to repeat, be Fodor’s characterisations of *learning* processes as “rational”, “computational”, “intentional”, “cognitive”, and “in the domain of cognitive psychology”, and *nonlearning* processes as “automatic”, “noninferential”, jumpy, “brute-causal”, and “merely triggered” by circumstances.

I think these characterisations suggest *a* way of distinguishing “learning” from “nonlearning” on which abstraction\* is *not* a learning process.

To begin with, it is not clear that abstraction\* includes any thinking or reasoning. In any event, it is not a process that is made up *solely* by thinking or reasoning: it begins with an *experience* and ends with a *capacity* for thought. To this extent and in this sense, abstraction\* is nonrational and noncognitive. It stands in contrast, for example, to the learning process that takes place in the “concept learning paradigm” described in Sect. 3 above. Recall that Sect. 3 never questioned that this process is a learning process; it just discussed what its *end result* is—possession of a new concept or possession of a new term for what the subject could already think and say. But whatever the end result is, the process of reaching it is, as was noted, a reasoning process.

<sup>16</sup> “Abstraction” as understood by Laurence and Margolis (2012) resembles Russell’s “abstraction” in building upon multiple experiences. Further (as I mentioned in note 5 above), it is a process that, they claim, results in possession of *generic* sensory concepts but cannot produce concepts of *specific* sensory qualities.

Somewhat connected with the nonrational and noncognitive character of abstraction\* is the fact that this process is not—or not solely—a personal-level *activity*. Gaining the ability to think about blue from experiences of blue is, at least in part, something that just happens to one, at the personal level.<sup>17</sup> Of course, to *use* this ability is to engage in a personal-level action, namely thinking. But *gaining* the ability is not a personal-level action. As a consequence of having the experience, one comes to possess the new thinking ability just as passively as one comes to have the experience itself. To this extent and in this sense, the process is an “automatic jump” and it stands again in contrast to the learning process of the “concept learning paradigm”.<sup>18</sup>

Given these nonrational, noncognitive, and automatic-jumpy features, there is ground for taking abstraction\* to be an instance of nonlearning, given Fodor’s characterisations of learning and nonlearning. However, there are also grounds for regarding abstraction\* as a learning process, if we focus on certain other of Fodor’s characterisations.

To begin with, there is reason to regard abstraction\* as a computational process. At least, there is reason to do so if we accept Fodor’s main line of argument for computationalism, which I will now briefly review.

Fodor is impressed by the fact that we can and often do explain action and cognition by appealing to psychological laws that are *general* both insofar as they quantify over *subjects* of mental states and insofar as they abstract from the *contents* of mental states. As an example of such a law, Fodor cites: “If  $x$  wants that  $P$  and  $x$  believes that  $\neg P$  unless  $Q$ , and  $x$  believes that it is within his power to bring it about that  $Q$ , then *ceteris paribus*  $x$  tries to bring it about that  $Q$ ”, where  $P$  and  $Q$  are schematic letters that may mean anything (1987, 13). Fodor sometimes refers to such laws as “intentional generalizations” (1994, 4).

How do processes that fall under intentional generalisations work? How does one for example transition from (a) a state of wanting that  $P$  and believing that  $\neg P$  unless  $Q$  and that is within one’s power to bring it about that  $Q$  to (b) a state of trying to bringing it about that  $Q$ ? According to Fodor, the computational theory of the mind offers the best—or only—available answer to these questions. This theory says that the transition from (a) to (b) is a process of symbol-transformation that is sensitive only to “syntactic” properties of symbols and not to the meanings or contents of these symbols, and yet reliably “respects” these meanings or contents. To illustrate, suppose that two symbols,  $O$  and  $E$ , are involved in a state of: wanting something that is symbolised by  $O$  and believing that what is symbolised by  $O$  will not come about unless what is symbolised by  $E$  comes about and that it is within one’s power

<sup>17</sup> “At the personal level” is meant to allow that the process involves active construction at the subpersonal level. “At least in part” is meant to allow that some personal-level action is a necessary *part* of the process. For example, it allows that active, personal-level focus of attention is a necessary part of the process as it occurs in us humans. See Sect. 7 below for some discussion of this issue.

<sup>18</sup> Here we can again take note of a difference—or lack of clear similarity at least—between abstraction\* and what has historically been associated with ‘abstraction’: the quote from Russell in the preceding section may be taken to suggest that “abstraction” as he conceived it is largely or entirely a personal-level activity, perhaps a kind of reasoning process. Again, “abstraction” as understood by Lawrence and Margolis (2012) is similar to Russell’s “abstraction” in this respect.

to bring about what *E* symbolises. And suppose that *O* and *E* can be identified in the mind (as can letters on paper) by their shapes. The relevant mind may then be set up in such a way that, by forces that operate solely on the shapes of *O* and *E*, it generates a state that involves the symbol *E* and amounts to trying to bringing about what *E* symbolises. If *O* in fact means that an omelette be made, and *E* that one breaks an egg, this process will have amounted to a transition *from* (1) a state of wanting that an omelette be made and believing that an omelette will not be made unless one breaks an egg and that it is within one’s power to break an egg *to* (2) trying to bringing it about that one breaks an egg. It will have amounted to this even though the mechanism driving this causal process was “blind” to these meanings. A parallel story can be told should *O* and *E* be identified in the mind by, say, frequency of a certain kinds of neural firing rather than by shape properties.

Let us now, against the background of this argument, consider the process, abstraction\*, of acquiring a concept of blue from an experience of blue, as this process now and then occurs in us. It seems clear that this process falls under a psychological law with the kind of generality—expressible by quantification over both subjects and contents of mental states—that tends to impress Fodor. That is just to say that the present focus on *blue* is arbitrary, as of course it is. It is not only the case that we humans can and sometimes do acquire a concept of *blue* from experiences of *blue*. We can and do sometimes also acquire a concept of *yellow* from experiences of *yellow*. More generally, it seems clear that some kind of “intentional generalisation” along roughly the following lines is true: *We humans are such that, for all x’s in a somewhat extensive domain: if one of us lacks a concept of x and has an experience of x then that person acquires an ability to think about x.* For present purposes it can be left open exactly how extensive the relevant domain is, but it presumably includes at least many (specific) sensory qualities. We can also leave open whether some further condition should be added to the antecedent of the just-stated generalisation, for example something about direction of attention.<sup>19</sup> That would not make a difference that matters for present purposes. What matters for present purposes is just that *some* principle that in the above way quantifies over both subjects and contents of mental states is correct.<sup>20</sup>

It is worth noting that it is not *trivial* that we sometimes acquire sensory concepts from experience by a process, abstraction\*, that falls under this kind of “intentional generalisation”. A creature could be such that it tended not to acquire a concept of blue from experiences of blue, or a concept of yellow from experiences of yellow, but instead tended to acquire the former concept from experiences of white triangles and the latter from itches on the nose.<sup>21</sup> It is a nontrivial fact that we are not creatures of that kind.

Now, insofar as there is reason to accept Fodor’s general argument that processes that fall under “intentional generalisations” are computational, there is reason to suppose that the process abstraction\* is computational. And if, as Fodor suggests, being computational is a distinctive characteristic of cognitive development that

<sup>19</sup> Compare note 17 above and Sect. 7 below.

<sup>20</sup> See Speaks 2015, chap. 16–7, for a defence of closely related claims.

<sup>21</sup> See Unger 1966 and Dennett 2005, chap. 5 for some related scenarios.

amounts to *learning*, that is in turn a reason to regard abstraction\* as a learning process.<sup>22</sup>

Furthermore, abstraction\* falls, I believe, in a distinctive sense “within the domain of cognitive psychology”. There is, I think, a certain “autonomy” to the explanation of someone’s ability to think about blue in terms of (a) that person’s having experienced blue, and (b) the fact that she is the kind of subject that tends to acquire an ability to think about blue from an experience of blue, and (c) the theory that this process is implemented by a “computational” mechanism. This explanation is, I think, “autonomous” in the sense that the process will not be much further illuminated by an account of the neural realisation of the relevant computational process. Suppose for example we were to find out that the postulated “syntactic” features of the symbols involved in this process are *numbers of pulses* rather than, say, *frequencies of firing* of certain neurons. That would add detail to the already given computational theory, but it would not provide more illumination than that. Something similar is true, I think, if we do not assume computationalism but explain someone’s ability to think about blue just in terms of (a) and (b).

Consider for contrast the imaginary case of some subject that is disposed to acquire an ability to think about blue from an experience of a white triangle. Here an account of neural realisation might well provide a greater degree of illumination. We could admittedly make some safe assumptions about *why* the subject has the relevant disposition, even without knowing details of neurology. Assuming that the mental at least nomologically supervenes on the physical, we could safely suppose that some kind of “hard wire” connects, in this subject, (1) the neural basis of the state of experiencing a white triangle with (2) the neural basis of the state of being able to think about blue. But the details of neurology could still explain *why* that hard wire exists. For example, its existence could be explained by the *proximity* of these two neural states. (It is easier for a “wayward” wire to accidentally grow between proximal regions than between remote regions; to that extent, proximity can explain the existence of such a wire.) It is hard, I think, to imagine a similar degree of illumination from a neural account added to the explanation of an ability

<sup>22</sup> A computational account of the abstraction\* could be developed in several ways. One could propose (a) that an experience of blue as it occurs in us involves a symbol of a certain syntactic type, B, with a content, B, and that our minds are set up in such a way that, when we have an experience involving B (and perhaps some additional conditions are satisfied) then a symbol of the same syntactic type, B, and with the same content, B, becomes available for use in *thought*. (This proposal involves the “content conceptualist” supposition that an *experience* of blue and a *thinking* about blue can have the same content. As note 15 above observes, that conceptualist claim is arguably compatible with the *state nonconceptualist* claim that one can experience blue without being able to think about blue.) Alternatively, one could propose (b) that an experience of blue as it occurs in us involves a symbol of the syntactic type B with the content  $B_E$ , and that when we have an experience involving B (and perhaps some further conditions are satisfied) then a symbol with the same syntactic type, B, becomes available for use in thought but that a *thinking* with B has—perhaps on account of the distinctive “inferential role” of B as involved in thought—a “conceptual” content,  $B_C$ , distinct from  $B_E$ . (This proposal naturally combines with the rejection of “content conceptualism”.) Again, one could within the computationalist frame propose (c) that an experience of blue as it occurs in us involves a symbol of a certain the syntactic type  $B_E$ , and that when we have an experience involving  $B_E$  (and perhaps some further conditions are satisfied) then a symbol of a different syntactic type,  $B_C$ , becomes available for use in thought; at least, one can suppose this within the computational frame as long as there is some systematic way of deriving symbol-types involved in thinking from symbol-types involved in experience.

to think about blue in terms of (a) an experience of blue and (b) the relevant intentional generalisation, perhaps together with (c) a computational implementation hypothesis.

If this is correct, then there is a sense in which abstraction\* is more exclusively “within the domain of cognitive psychology” than is the process, of some possible subject, of acquiring a concept of blue from an experience of a white triangle. And if being “within the domain of cognitive psychology” is, as Fodor suggests, distinctive of learning processes, that is in turn another reason to regard abstraction\* as a learning process.

## 6 Geach on “Abstractionism”

It is often supposed that Geach (1957) presents an important case against the idea that sensory concepts can be learned by “abstraction” from experience. I shall now consider whether Geach provides any serious concern for the proposal developed here.

Geach takes it as given that concepts essentially form a network of capacities to make various judgments. The view is illustrated with an analogy:

the ability to make an appropriate [chess] move ... presupposes ... the capacities to carry out the moves and captures that are lawful for the pawns and the various pieces. As these skills are related to the chess-move, so concepts are related to the act of judgment (13).

Applied to a sensory concept like “red”, Geach’s idea is, I think, that to have this concept is not merely to be able to form judgments like, “this is red”. It is a more complex ability that involves also being able to form judgments like, “this is *not* red” and “something could be red and triangular”.

Geach claims that this kind of complex ability cannot be learned from experience. He does not in fact provide much argument for this claim. He appears to regard it as more-or-less obvious. The two passages that contain his most developed thoughts behind this assertion read as follows:

if I possess the concept *red*, then I can perform acts of judgment expressible in sentences containing the word “red”. This ability, however, certainly cannot be learned by any kind of attention to red patches for any length of time; even if after a course of attending to red patches the ability turned out to be present, we should still be justified in refusing to say it had been *learned* that way. We can say this just as confidently as we can say that the ordinary use of the word ‘red’ cannot be learned by hearing the word ‘red’ uttered ceremonially in the presence of a red object—simply on this account, that such a ceremony is not the ordinary use of the word ‘red’ (1957, 34).

it is integral to the use of a general term that we are not confined to using it in situations including some object to which the term applies; we can use the

terms ‘black’ and ‘cat’ in situations not including any black object or any cat. How could this part of the use be got by abstraction? (ibid., 34–5).<sup>23</sup>

However, *even if* one grants Geach’s view that (1) having a given concept is to have a network of capacities to make a variety of judgments, one cannot, I think, derive from these passages a strong challenge to what has been defended above, namely that (2) drawing on at least some received characterisations of “learning” and “nonlearning”, we can and do sometimes *learn* sensory concepts on the basis of experience.

Note to begin with that—as I indicated earlier—Geach seems specifically to aim for the conclusion that sensory concepts cannot be *learned* from experience; he does not question that they can be *acquired* from experience. This seems clear from the following passage, where Geach defends his a priori approach to the topic of abstractionism by appealing to a limitation in his aims:

the connexion between *learning* and *being able* to do something is not just an empirical one; for we should plainly not be willing to call *any* process that ended in ability to do something “learning” to do it. Of course we cannot work out in our heads how abilities are acquired; but there are conceivable ways of acquiring them to which we should unhesitatingly refuse to apply the term “learning”. If ... a boy could come to know Latin by submitting to a brain operation, he would not have *learned* Latin from the surgeon. Now abstractionism is a theory of how certain mental performances are *learned*; the discussion I am entering upon is designed to show that the processes alleged by abstractionists would not be a *learning* of how to make these mental acts (18–19).

Recall also Geach’s previously quoted words: “even if after a course of attending to red patches the ability turned out to be present, we should still be justified in refusing to say it had been *learned* that way”.

Moreover, there is, as Sect. 4 above argued, good reason for Geach to aim for this limited conclusion because it is very plausible that sensory concepts can be *acquired* from experience (indeed, they *can* presumably be acquired from anything, including pills) and even strong reason to think that sensory concepts sometimes *are* acquired thus.

But given this, it is clear that Geach’s thoughts do not put in question anything in the present proposal. Nothing in Geach’s passages challenge the grounds given in Sect. 5 for regarding the process of acquiring a concept of blue from experiences of

<sup>23</sup> Here and elsewhere, Geach goes back and forth between talk of (a) concept possession and (b) mastery of *words*. Geach does not assume that these are the same or equivalent. But he thinks it is safe enough to focus on concepts that are “exercised linguistically”, even though he does not assume that all concepts are (ibid., 12–3). Besides the just-quoted passages there is one further passage where Geach develops some motivation for the claim that a concept like “red” cannot be learned on the basis of experience. In this passage Geach argues that (1) a concept like “chromatic colour” cannot be learned on the basis of experience, that (2) this concept is “surely no less ‘simple’, no less a concept of something ‘directly given in experience’ than the concept *red*” (37), and that (3) it is therefore “clear ... that the concept *chromatic colour* and the concept *red* stand on the same footing” (38). I will not discuss this passage in detail here. But in brief, I believe that at least (2) is very doubtful. Geach does not articulate any support for either (1) or (2).



blue as a learning process. The passages do not put in question (a) that this process, abstraction\*, falls under an intentional generalisation and that there is to that extent reason to regard it as “computational”, or (b) that abstraction\* falls in a distinctive way “within the domain of cognitive psychology”, or (c) that being computational or in the domain of cognitive psychology are characteristics of learning processes.

## 7 Worries About Circularity and Attention

“Abstractionist” views of concept learning or acquisition have often been subject to worries about circularity and attention. The general character of such worries is that: the idea of abstracting a concept from experience involves vicious circularity because (1) the presumed abstraction would require *attention* to relevant features in experiences, but (2) this attention would in turn require application and therefore prior possession of the concept allegedly abstracted.<sup>24</sup>

I think there is no serious concern in this neighbourhood for the present proposal. To see this, it is helpful to first consider whether worries like the above concern concept *learning* specifically or concept *acquisition* more broadly. The answer to that must be, I think, that they concern the latter. (What would the alternative construal be? That *acquiring* a concept of blue from experience does not require concept-involving attention, but that this acquisition process is not an instance of *learning* because...?)

When that much is clear it is also, I think, clear that there is no serious concern of the present kind for the proposal made here. In Sect. 4 above, I tried to show that it is both extremely plausible and widely accepted that we can and sometimes do *acquire* sensory concepts from experiences. And what is actual is possible. We can therefore conclude that it is not both the case that (a) *this* process, abstraction\*, requires a certain kind of attention, and that (b) this kind of attention presupposes application and therefore prior possession of the concept that abstraction\* results in.

One may of course wish to know, if it not the case that both (a) and (b), which one is not the case? I do not think we need a fully worked-out answer to this question to be assured of our present conclusions, but I offer some brief thoughts.

If we focus on the human situation—concept acquisition or learning as *we* in fact go through it—I am inclined to think that abstraction\* requires some attention; minimally some attention—either stimulus-driven (bottom-up) or goal-driven (top-down)—to some region of experienced space. Unless I at least attend to a region of experienced space where a specific sensory quality appears I tend not to acquire from experience the ability to think such things as, “ripe tomatoes do not have *that* shade of colour”. However, I see little reason to suppose that this kind of attention should require possession of the relevant sensory quality concept. And there is strong reason to suppose that this is not the case since it is (to repeat) (1) extremely plausible that we can acquire a concept of blue from experiences of blue (Sect. 4),

<sup>24</sup> For variants of such concerns, see Carruthers 1992, 55–6; Lowe 1995, 161–2; Laurence and Margolis 2012, sect. 2; and McGinn 2013, note 9. A similar concern is floated in Gauker 2011, 27.

and (b) plausible that such concept acquisition requires that kind of attention (this paragraph).<sup>25</sup>

## 8 Concluding Remarks

I have set aside some large questions about whether sensory concepts provide the basis for all our thought. But I have noted that it is both plausible and widely accepted that we can and sometimes do acquire sensory concepts from experiences of that to which they apply (Sect. 4). I have also (in effect) observed that some received characterisations of “learning” and “nonlearning” are nonequivalent, and that, drawing on some of these characterisations—in particular, the characterisations of learning processes as “computational” and “in the domain of cognitive psychology”—there is reason to place the acquisition of a concept of blue from experiences of blue among learning process (Sect. 5). Some widely circulated concerns about concept acquisition by “abstraction” from experience do not provide reasons to revise that conclusion (Sects. 6 and 7). Taken together, this calls into question a currently wide-spread consensus that primitive sensory concepts cannot be learned (recall Sect. 1).

I have rather uncritically assumed Fodor’s characterisations as guides to what learning is. Let me close with some reflections, not constrained by that assumption, on how one might, in view of the above, think about “the learning-nonlearning distinction”. I will mention four possibilities.

First, one could propose a kind of *eliminativism* about learning. For purposes of education, research heuristics, or understanding cognition and cognitive architecture, it matters—so this thought goes—whether a given process of cognitive development is rational, cognitive, computational, active or passive at the personal level, and in the domain of cognitive psychology. However, the above shows that our concept of “learning” is a confused mix of dissociable elements. This concept is therefore best left behind by attempts to understand cognitive development in developmental psychology, education sciences, and philosophy of mind.

Alternatively, one could propose that, while the above shows that our present concept of “learning” is confused, we should conclude nothing stronger than that this concept needs some—possibly minor—refinement.

A third thought is that our concept of “learning” is perfectly fine. Specifically, it usefully collects a number of different features that almost never dissociate. Admittedly, the preceding shows that the concept does not clearly apply, or fail to apply, to abstraction\*. But as Sect. 1 noted, this case might be exceptional; it might be a hard case that makes bad law, and we should not on the basis of this one case abandon a concept that is otherwise useful.

A fourth thought is that the preceding illustrates limitations of a reductive approach to learning. We should not expect learning to be well captured in other

<sup>25</sup> Campbell (2002) provides a detailed defence of the related claim that conscious attention to an *object* is “a state more primitive than thought about an object, to which we can appeal in explaining how it is that we can think about the thing” (45; see pages 68–78 for a central part of Campbell’s defence of this claim; see also Dickie 2011, sect. 2, for a related argument).

terms, like “cognitive”, “rational” or “computational”. Learning is what it is and not another thing, and it should be approached in an appropriately nonreductive manner.

I leave it for further research to determine whether either of these conclusions is right.

**Acknowledgements** Many thanks to Ben Bayer, Dave Chalmers, Tim Crane, Jerry Fodor, Kalle Grill, Nat Hansen, Neil Mehta, Jonas Nilsson, Richard Stöckle-Schobel, Andreas Stokke, and Inge-Bert Täljedal for very useful comments and discussions. Initial work on this paper was funded by Riksbankens jubileumsfond (RJ).

**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

- Aydede, M., & Güzeldere, G. (2005). Cognitive architecture, concepts, and introspection: An information-theoretic solution to the problem of phenomenal consciousness. *Noûs*, 39(2), 197–255.
- Ball, D. (2009). There are no phenomenal concepts. *Mind*, 118(472), 935–962.
- Bermudez, J. L. (2003). *Thinking without words*. Oxford: Oxford UP.
- Byrne, A. (2005). Perception and conceptual content. In E. Sosa & M. Steup (Eds.), *Contemporary debates in epistemology* (pp. 231–250). Oxford: Blackwell.
- Campbell, J. (2002). *Reference and consciousness*. Oxford: Oxford UP.
- Carey, S. (2009). *The origin of concepts*. Oxford: Oxford UP.
- Carey, S. (2014). On learning new primitives in the language of thought: Reply to Rey. *Mind and Language*, 29(2), 133–166.
- Carey, S. (2015). Why theories of concepts should not ignore the problem of acquisition. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 415–454). Cambridge, MA: MIT Press.
- Carruthers, P. (1992). *Human knowledge and human nature: A new introduction to an ancient debate*. Oxford: Oxford UP.
- Carruthers, P. (2004). On being simple minded. *American Philosophical Quarterly*, 41(3), 205–220.
- Chalmers, D. (2005). The matrix as metaphysics. In C. Grau (Ed.), *Philosophers explore the matrix* (pp. 132–176). Oxford: Oxford UP.
- Danto, A. (1963). What we can do. *Journal of Philosophy*, 60(15), 435–445.
- Danto, A. (1965). Basic actions. *American Philosophical Quarterly*, 2(2), 141–148.
- Davidson, D. (1975). Thought and talk. In S. Guttenplan (Ed.), *Mind and language* (pp. 7–23). Oxford: Oxford UP.
- Dennett, D. (2005). *Sweet dreams: Philosophical obstacles to a science of consciousness*. Cambridge, MA: MIT Press.
- Dickie, I. (2011). Visual attention fixes demonstrative reference by eliminating referential luck. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and psychological essays* (pp. 292–322). Oxford: Oxford UP.
- Fodor, J. (1975). *The language of thought*. Cambridge, MA: Harvard UP.
- Fodor, J. (1981). The present status of the innateness controversy. In *RePresentations: Philosophical essays on the foundations of cognitive science* (pp. 257–316). Cambridge, MA: MIT Press.
- Fodor, J. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.
- Fodor, J. (1994). *The elm and the expert: Mentalese and its semantics*. Cambridge, MA: MIT Press.
- Fodor, J. (2008). *LOT 2: The language of thought revisited*. Oxford: Oxford UP.

- Frege, G. (1997 [1892]). On *sinn* and *bedeutung*. In M. Beaney (Ed.) *The Frege reader*, pp. 151–71. Oxford: Blackwell.
- Gauker, C. (2011). *Words and images: An essay on the origin of ideas*. Oxford: Oxford UP.
- Geach, P. (1957). *Mental acts*. London: Routledge and Kegan Paul.
- Hawthorne, J., & Manley, D. (2012). *The reference book*. Oxford: Oxford UP.
- Heck, R. (2000). Nonconceptual content and the “space of reasons”. *Philosophical Review*, 109(4), 483–523.
- Hume, D. (1978 [1739]). In P. Nidditch (Ed.), *A treatise of human nature*. Oxford: Oxford UP.
- Jackendoff, R. (1989). What is a concept, that a person may grasp it? *Mind and Language*, 4(1–2), 68–102.
- Jeshion, R. (2010). *New essays on singular thought*. Oxford: Oxford UP.
- Kaplan, D. (1989). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 565–614). Oxford: Oxford UP.
- Laurence, S., & Margolis, E. (2002). Radical concept nativism. *Cognition*, 86, 25–55.
- Laurence, S., & Margolis, E. (2012). Abstraction and the origin of general ideas. *Philosophers’ Imprint*, 12, 1–22.
- Levine, J. (2010). Demonstrative thought. *Mind and Language*, 25(2), 169–195.
- Locke, J. (1975 [1689]). In P. Nidditch (Ed.), *An essay concerning human understanding*. Oxford: Oxford UP.
- Lowe, J. (1995). *Locke on human understanding*. London: Routledge.
- Margolis, E. (1998). How to acquire a concept. *Mind and Language*, 13(3), 347–369.
- Margolis, E., & Laurence, S. (2011). Learning matters: The role of learning in concept acquisition. *Mind and Language*, 26(5), 507–539.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA.: Harvard UP.
- McGinn, C. (2013). The origin of ideas. On his blog *philospot*. Retrieved August 30, 2016. Website: <http://philospot.com/>.
- Peacocke, C. (1992). *A study of concepts*. Cambridge, MA.: MIT Press.
- Peacocke, C. (2001). Does perception have a nonconceptual content? *Journal of Philosophy*, 98(5), 239–264.
- Pinker, S. (2007). *The stuff of thought*. New York: Viking.
- Price, H. H. (1953). *Thinking and experience*. London: Hutchinson.
- Prinz, J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA.: MIT Press.
- Roskies, A. (2008). A new argument for nonconceptual content. *Philosophy and Phenomenological Research*, 76(3), 633–659.
- Roskies, A. (2010). That’ response doesn’t work: Against a demonstrative defense of conceptualism. *Noûs*, 44(1), 112–134.
- Russell, B. (1910–1911). Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society*, 11, 108–128.
- Russell, B. (1912). *The problems of philosophy*. London: Williams and Norgate.
- Samet, J. (1986). Troubles with Fodor’s nativism. *Midwest Studies in Philosophy*, 10(1), 575–594.
- Speaks, J. (2005). Is there a problem about nonconceptual content? *Philosophical Review*, 114(3), 359–398.
- Speaks, J. (2015). *The phenomenal and the representational*. Oxford: Oxford UP.
- Sundström, P. (2004). Lessons for Mary. In J. Marek & M. Reicher (Eds.), *Experience and analysis: Papers of the 27th international wittgenstein symposium* (pp. 367–369). Kirchberg am Wechsel: The Austrian Ludwig Wittgenstein Society.
- Tye, M. (2009). *Consciousness revisited: Materialism without phenomenal concepts*. Cambridge, MA: MIT Press.
- Unger, P. (1966). On experience and the development of the understanding. *American Philosophical Quarterly*, 3(1), 48–56.
- Weiskopf, D. (2008). The origins of concepts. *Philosophical Studies*, 140, 359–384.