



Perceptual learning and feature-based approaches to concepts – a critical discussion

Richard Stöckle-Schobel *

School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh, UK

Edited by:

Arnon Cahen, Ben Gurion University
in the Negev, Israel

Michela C. Tacca, Heinrich-Heine
University Düsseldorf, Germany

Reviewed by:

Philippe G. Schyns, University of
Glasgow, UK

Rob Goldstone, Indiana University,
USA

***Correspondence:**

Richard Stöckle-Schobel, School of
Philosophy, Psychology and Language
Sciences, University of Edinburgh,
Dugald-Stewart-Building, 3 Charles
Street, EH8 9AD Edinburgh, UK.
e-mail: s0959457@sms.ed.ac.uk

A central challenge for any theory of concept learning comes from Fodor's argument against the learning of concepts, which lies at the basis of contemporary computationalist accounts of the mind. Robert Goldstone and his colleagues propose a theory of perceptual learning that attempts to overcome Fodor's challenge. Its main component is the addition of a cognitive device at the interface of perception and conception, which slowly builds "cognitive symbols" out of perceptual stimuli. Two main mechanisms of concept creation are unitization and differentiation. In this paper, I will present and examine their theory, and will show that two problems hinder this reply to Fodor's challenge from being a successful answer to the challenge. To amend the theory, I will argue that one would need to say more about the input systems to unitization and differentiation, and be clearer on the representational format that they are able to operate upon. Until these issues have been addressed, the proposal does not deploy its full potential to threaten a Fodorian position.

Keywords: concept learning, perceptual learning, computationalism

Cognitive psychology has recently seen the development of several new models positing a perceptual basis for conceptual systems. The panoply of views ranges from proposals to eliminate the distinction between concepts and percepts altogether (Barsalou, 1999) over more modest appraisals of the relations between the two (Goldstone and Barsalou, 1998) to proposals for the creation of cognitive processes through experience with perceptual stimuli (Schyns and Rodet, 1997). As one important contribution to this line of research, Robert Goldstone's perceptual learning approach stands out and shall be at the center of our present investigation in the links between the perceptual and the conceptual.

Among the specific questions related to perception and learning, Goldstone and his colleagues and collaborators discuss the possibility and mechanisms of perceptual learning (Goldstone, 1998), the influence of perception on categorization (Landy and Goldstone, 2005), the role of features of objects in categorization (Schyns et al., 1998), and learning in early ontogeny (Goldstone et al., 2011). Their target is the more conservative fixed-feature approach – a form of computationalism, which holds that new concepts are constructed by using pre-existent, cognitively fixed features. One avid defender of this view is Fodor (1975, 1981, 2008), whose classic argument against concept learning especially affects perceptually based "empiricist" theories. The upshot of this argument is that learning concepts needs to be based on a vocabulary in which hypotheses about these concepts are formulated. But that vocabulary itself already needs to contain the concept that is just being "learned." Fodor infers from this that all concepts must be innate, or at least not learned. In what follows, I will call this "Fodor's Challenge for theories of concept learning." Taking its conclusion as an undesirable outcome for any theorist who wants to maintain a notion of genuine learning, one might ask the following question.

Assuming that cognition is at least a partly computational process, is there any way of having new symbols from perceptual origins entering the internal symbol system? As Goldstone answers this to the affirmative, I will discuss his proposal and point out two problems with it that need more consideration.

FODOR'S CHALLENGE FOR THEORIES OF CONCEPT LEARNING

In order to see where Goldstone and colleagues aim when they criticize the fixed feature position, I will first briefly set out this position and the challenge it poses to research in perceptual learning. Fodor goes through the following steps to reach the conclusion that concepts, and by that also features, cannot be learned.

First, Fodor construes learning mechanisms as "rational-causal processes" (Fodor, 1981, p. 273) – Being a rational process, learning is mediated by psychological states, such as beliefs. Also, regarding the possible constituents of thought, Fodor argues that what is not learned is innate or acquired in some other non-rational way.

A further premise of his argument concerns possession conditions for concepts: "A sufficient condition for having the concept C is: being able to think about something *as (a) C*" (Fodor, 2008, p. 138, original emphasis.). This means that the main act of concept use is using the concept in forming beliefs (or other types of thoughts) – as contrasted with using the concept to categorize new sensory experiences, or to act upon a thing in the world. For Fodor, thinking is prior to perceiving and acting in the order of concept use.

Now, from this position one needs to have a model of how a concept can enter this realm of thought. Fodor argues that the only available, empirically tested model for learning is the following: learning the concept C consists in forming hypotheses about C and testing them against the available evidence. Thus, learning

is a process of inductive inference. Forming a hypothesis about the concept *C* requires bringing the property expressed by *C* before one's mind. One needs to think about a piece of evidence "x" as (a) *C* to (dis-)confirm the hypothesis about *C*. To learn which things are green, one must judge something to be (or not be) green. This act of judging is a mental going-on for which one needs to be able to think about green things, or about a thing as a green thing. Now, what is already used for hypothesis formation is not learned in the application (confirmation or disconfirmation) of the hypothesis. *C* was already available to form the hypothesis, thus *C* was not learned.

Fodor's conclusion is that all concepts are either innate or non-rationally (brute-causally, see Fodor, 1981) acquired. This conclusion is supposed to affect theorizing about learning concepts in all areas of the cognitive sciences, from developmental psychology to artificial intelligence research, since it affects the theory choices one has to explain the phenomena of these disciplines. Consider artificial intelligence: an important research aim of AI is to develop systems with human-like intelligence – computer programs that play chess like grandmasters, robots that move like biological organisms, and the like. In order to arrive at a theoretical foundation for such systems, philosophers explore the possibility and extent of the computational theory of mind (CTM) hypothesis – roughly, that the human mind can be best described as a system that works like a computer that operates on symbolic representations. For a Fodor-type computationalist, the number of symbols would be predetermined by the system, and so the symbols would be innate. Given the additional constraint that each symbol of such a computational mind equals one concept, one has arrived at the point where Fodor's challenge and the computationalist program tie in. Landy and Goldstone (2005) describe such conceptions of CTM as being essentially linked to the idea that a fixed store of primitive, basic symbols is sufficient for successful cognition, and continue by saying that this classical version of CTM "entails a fixed set of primitives, or at least demands that any alterations to the primitive set are not cognitively interesting acts" (Landy and Goldstone, 2005, p. 346). Thus, we have characterized one stance toward Fodor's challenge, the fixed-features approach: it accepts the conclusion of Fodor's challenge, embraces the radical Concept Nativism that it entails, and denies any transformational effect on the cognitive system that would count as learning a new primitive concept.

By challenging the sufficiency of a fixed set of symbols for explaining human cognition, and by denying that changes to primitive symbols are not cognitively interesting, Landy and Goldstone set out their alternative to the Fodorian position and by that give us the second theory that we will presently take into account as a reply to Fodor's challenge. They want to argue for the creation of cognitive symbols from perceptual materials, and they want to argue for the possibility of manipulating "systems of high-level categories" (Landy and Goldstone, 2005, p. 346) to better fit the demands of the cogniser. The question motivating the present investigation thus is: Can Goldstone's theory of perceptual learning, and especially Landy and Goldstone's stance against fixed-feature languages, stand against Fodor's challenge and the acceptance of the fixed-feature approach, and can it give a mechanistically and computationally credible account of human concept learning?

PERCEPTUAL LEARNING AS A REPLY TO FODOR'S CHALLENGE

The question before us is whether it is possible to enrich a symbol system through the manipulation or introduction of perceptual information, or perceptual symbols. Learning features, like other forms of concept learning, can in an important sense be seen to hinge on the possibility of arriving at thoughts one was not able to hold or express before, and thus on having an alternative to Fodor's innateness conclusion by ways of providing an alternative empirical model for concept learning (rejecting the premise that hypothesis formation and testing is the only empirically available model for concept learning). The Fodorian CTM perspective, on Landy and Goldstone's (2005) account, denies this possibility, whereas several recent contributors to the debate have tried to develop models that support an affirmative answer. One main inspiration for this project comes from Gibson's (1963) theory of perceptual learning. A second major theoretical development was initiated by the work of Philippe Schyns on feature creation through experience with stimuli (Schyns and Murphy, 1991, 1994; Schyns and Rodet, 1997), leading up to the unified account of Schyns et al. (1998). The idea that the learning of a novel vocabulary of features yields new categorizations, which will be introduced below as a part of Landy and Goldstone's (2005) account, is rooted in the groundbreaking work of Schyns and his colleagues.

With their proposal, Landy and Goldstone mainly challenge Fodor's assumption that the primary use of concepts is in forming thoughts, as opposed to using concepts in dealing with the world via reacting on (sensory) inputs and acting in it/producing (behavioral) outputs. Grounding concept use and concept learning in perception does however not preclude the use of new perceptual concepts in higher cognitive activities – this is an important point made by Landy and Goldstone, e.g., in their discussion of changes in scientific reasoning through perceptual changes¹. It is worth dwelling on this aspect of Goldstone's theory before turning to the core of Landy and Goldstone's (2005) proposal. Goldstone sees the possibility of what he calls perceptual learning, following Gibson (1963):

Any relatively permanent and consistent change in the perception of a stimulus array, following practice or experience with this array, will be considered as perceptual learning. (Gibson, 1963, p. 29)

On this definition, perceptual learning is a sensory as well as a cognitive process: the changes in focus, or attentional center, to give two examples, in seeing something are at the same time changes in the categories pertaining to the perceived object. The repeated sensory contact with a certain class of objects will bring about a change in the way one thinks about these objects, which will in turn influence its perception, i.e., the sensory processes.

Goldstone explicitly wants to trace the ties between these perceptual changes and the possible conceptual changes that accompany them. He holds that one traditionally neglected aspect of the relation between perception and conception is the influence that

¹Landy and Goldstone (2005) discuss changes in ontology, cognitive properties of groups of scientists, and changes in scientific practice through new perceptual capacities as cases in point.

the conceptual system has on perception. In categorical perception, the learned categories influence the performance in perceptual tasks. Especially in the sciences, there are multiple examples for this. Mathematicians can name several properties of a function just by looking at its graph. Similarly, after studying the geological categories and training to differentiate various stone samples, geologists have a sharper grasp of the differences between stone types and are able to name them much faster than any layperson could (Goldstone, 1994; Goldstone and Hendrickson, 2010; Goldstone et al., 2012). This is also the second point Goldstone and Barsalou (1998) stress:

(...) perception's usefulness in grounding concepts comes from several sources. First, perception provides a wealth of information to guide conceptualization. Second, perceptual processes themselves can change as a result of concept development and use. Third, many of the constraints manifested by our perceptual systems are also found in our conceptual systems. (Goldstone and Barsalou, 1998, p. 232)

The first statement of this quote, that our perception can be a source of information for our conceptual system, does not sound very controversial since it is not very informative and specific in itself. In what way does perception inform conception? Even on Fodor's account, perception informs conception in so far as perceiving an object x can cause the triggering of the accompanying concept X . For Goldstone, and especially for Barsalou (1999), there needs to be a more detailed description of the way in which perceptual information touches upon our concepts; a description which probably even does away with the distinction between perception and conception. The second point has bearing on the present question in so far as it is the converse of the claim that Landy and Goldstone (2005) put forward to challenge Fodor: if both of these directions of influence were part of the actual workings of the human mind, then the strongly computationalist position would either lose a lot of its plausibility, or would have to be reformulated to accommodate these interrelations. Such an accommodation would however run against the self-proclaimed Rationalist position that Fodor adopts. Finally, the third point is especially important for Barsalou's (1999) project, but beyond the scope of the current investigation.

With these preliminaries set out, let's see how they form a frame for Landy and Goldstone's (2005) answer regarding Fodor's challenge. A short characterization of the main aims of argumentation, with references to extended presentations and discussions of these points, can be put as follows:

1. In learning about things we do not already understand, our cognitive system constructs specialized variable-feature languages that deal with these novel things (cf. Schyns and Murphy, 1991, 1994; Schyns and Rodet, 1997; Quinn et al., 2006).
2. The vocabulary of these languages consists of stimuli that we perceptually pick up and group as belonging to features, or feature dimensions (Schyns et al., 1998).
3. New features can be learned by applying the grouping mechanisms of unitization and differentiation, as the main players among other perceptual mechanisms (Goldstone, 1998, 2003; Goldstone and Landy, 2010).

It is generally assumed that concepts are the tools for, or the components of, thought. Thus, they are rather highly developed parts of our mental lives – conceptual thought is at the upper end of the scale of cognitive activity. Many things that we think about are very specific to a problem domain, like choosing a move in a chess game, while others are central to many modern human activities, like deciding which way to go to reach the nearest restaurant. Keeping with the computational tradition in the study of cognition, one can speak of different “vocabularies” or symbol stores for different tasks, with some being used for a more diverse range of activities than others.

Landy and Goldstone (2005) frame the debate as pertaining to languages of cognitive systems, which is not an uncommon level of discussion, given that Computationalism treats cognition as symbol-manipulation, and a number of symbols, combined with operations over these symbols, can with some right be called a “language.” In the context of this paper, I propose to call such a language a computational language (Language_C), to highlight that the sense of “language” is somewhat restricted as compared to a spoken language. Computationalists like Newell and Simon (1976) or Fodor and Pylyshyn (1981, 1988) favor a fixed Language_C, whose symbols are inherent in the cognitive system, and sufficient for any kind of cognitive activity within that system – there is no need to import new symbols, since the given stock is supposed to express any proposition that the system would need to process. Biederman's (1987) geon model is another example for such a fixed Language_C, with the added twist that he attempts to posit perceptual representations – representations of basic geometrical forms – as a part of the innate stock of symbols.

To counter this model, Landy and Goldstone present what they call a “variable-feature language” (Landy and Goldstone, 2005, p. 347): a Language_C that can be enriched with new primitive symbols, if new perceptual tasks require this. In Landy and Goldstone (2005), they characterize these enrichments as additions to particular sets of symbols, constrained by the category, or task, they are used for. In this, they follow Schyns and Murphy's (1991, 1994) major contribution to feature-based approaches to concept learning². This leads to changes in highly specialized vocabularies, and need not necessarily affect the foundations of the Language_C. Landy and Goldstone talk of special-purpose Languages_C and general-purpose Languages_C in the cognitive system. While the latter are not excluded by Landy and Goldstone to be innate, given that they are ubiquitous in the most basic cognitive functions, the former need to be learned on their account. This is because the tasks that they are needed for are highly specialized in one way or another: examples that their paper discusses are fine perceptual discriminations such as discriminating brightness and saturation, and scientific theorizing and theory construction. Landy and Goldstone (2005, p. 348) compare the cognitive symbol system to LEGO blocks: some objects can only be constructed in a very cumbersome manner if only using the standard blocks (think of sails for a pirate's ship), so adding LEGO sails to their

²I want to thank an anonymous reviewer for pointing this out to me. The functionality principle, that functional demands shape the perceptual processes of categorizing new stimuli and forming new featural discriminations, from Schyns and Murphy (1994), has been a cornerstone of recent work in this area.

repertoire facilitates that specific kind of building process. The disadvantage of these special parts, however, is that they cannot serve for much else except their originally intended function. This, again, echoes the constraints on special-purpose Languages_C: the concept “color saturation” only has a very limited set of tasks for which it is needed, whereas the concept “not” has a scope that’s equivalent to the generic LEGO blocks.

Now, in terms of the mechanisms of learning, Landy and Goldstone’s theory’s main component is the addition of a cognitive device at the interface of perception and conception, which slowly builds “cognitive symbols” out of perceptual stimuli. By adding these new symbols to the symbolic building blocks of thought, this device is the agent of concept learning and conceptual change. The main operations in this system are unitization and differentiation, two mechanisms which either unite previously separated conceptual elements, or split a vaguely bounded element class into finer groupings. In my present investigation, I will focus on these two mechanisms, since they are central to the argument by Landy and Goldstone (2005). When linking their theories with other, related work in the field, like in Goldstone et al. (2011), or Goldstone and Landy (2010), they also discuss other ways of learning. These include processes that Fodor would classify as brute-causal acquisition rather than genuine learning, which raises some questions I can only hint at in this investigation. Nevertheless, I will introduce these mechanisms later on.

Unitization can be described as a process of grouping several previously independent categories under one heading: “When elements co-vary together and their co-occurrence predicts an important categorization, the elements tend to be unitized” (Landy and Goldstone, 2005, p. 350). Here is an example of a process of unitization learning: suppose you learn what a cup is by seeing various different cups and not-quite-cuplike objects. Something qualifies as a cup if it consists of a cylindrical container and a handle to the side of the container. The contrast class of cuplike objects consists of other configurations of containers and handles, like a handle spanning the top or the bottom of the container (the former looking a bit like a bucket), or with a handle only connecting with the container at one point (looking like a horn attached to the cylinder), or even just unconnected handles and containers. The rules of unitization would incline you to unite the two featural elements (cylinder and handle) into a token of the concept if and only if they are in the right spatial configuration (handle on the side, both parts properly connected). Unitization allows you to conceive of the two parts as one object, and with that also to keep unfitting combinations, which do not satisfy the perceptual constraints, out of the class of cups.

It may be necessary to distinguish two kinds of uniting learning cases: associative chunking and perceptual unitization³. Associative chunking is the process through which two elements that co-occur regularly become associated: if one is accustomed to getting a glass of water with an ordered cup of coffee, then being served just a cup of coffee will create the expectation of a glass of water that has to follow: “drinking coffee” as an activity-concept has these two elements.

Let me consider a case in which two feature dimensions are reliably correlated so that the occurrence of either one is a reliable sign of the occurrence of the higher-order phenomenon. Would such a case be more aptly described as unitization or as associative chunking?

Suppose that fire fighters always take big red vehicles that sound off a siren alternating between the first and the fourth tone of a scale (say C and F), and that all other emergency sirens use a different interval, say the prime and the fifth (C and G). Upon learning about the visual properties of fire engines, one might form a concept “fire engine” that is related to big red vehicles. Having also learned that the peculiar siren sound of the prime and the fourth is the fire fighter siren (having formed a concept “fire fighter siren”), one has formed the basis for putting together those two stimuli as the two most reliable signs for the presence of a fire engine. Thus, either of the stimuli can be used to trigger the concept “fire engine,” despite the lack of the other. While more elaborate than the coffee-and-water case, from Goldstone’s perspective this would still count as chunking, since the co-occurrence is not based on spatial, but on causal and temporal contiguity, which supports the formation of two separate feature elements that later get a common “heading” “fire engine.”

The difference between unitization and chunking can be made clearer by another example, this time adapted from Landy and Goldstone (2005, p. 352): in the same way that a photograph of a group of people is a combination of pictures of the individuals, a unitized concept is a combination of features that stand in certain cognitively interesting, complex relations – possibly with spatial configuration as the main combinatorial criterion (again, as in the photograph case). In light of the perceptual constraints on unitization, which are the main point of difference to associative chunking, it would be more prudent to not expect unitization over different sense modalities such as visual and sound perception (as in the fire engine example) – at least until robust experimental data supports this idea.

The second mechanism for concept learning is dimension differentiation, “by which dimensions that are originally psychologically fused together become separated and isolated” (Goldstone, 2003, p. 249). Especially in differentiating dimensions, perceptual constraints influence the process: while it is easy for adult perceivers to separate the properties “size” and “brightness,” it is much more difficult for the non-specialist to separate other fused dimensions such as brightness and hue. Differentiation might also be at work in separating non-dimension features, as in the fire engine example above. Suppose one never has paid much attention to siren sounds, and so has never noticed the difference between the fire fighters’ siren and all other sirens – one has one single concept “emergency siren.” Upon learning about the tonal difference between the two intervals, probably in a music class, one might start noticing the difference, and thus differentiate one’s concept into “fourth interval siren” and “fifth interval siren,” and then even relate these to the appropriate kinds of emergencies (“there’s been a robbery next door, I’m quite sure that I’ll soon hear the fifth interval police siren”).

To make the differences between the perceptual learning perspective and classical computationalism clearer, here is another rephrasing with an example. For the fixed feature approach, new

³The distinction was pointed out to me by Robert Goldstone.

mental representations are new combinations of previously available primitive elements. Associative chunking, as in “glass of water” and “cup of coffee” as components of the concept “things that I drink when having coffee,” requires the availability of the components that are combined. In the variable-feature approach, new representations need not necessarily be primitive elements, or psychologically pre-available elements. Rather, they can be stimulus elements with “no parsing in terms of psychological primitives” (Landy and Goldstone, 2005, p. 350) – so, what Landy and Goldstone want to argue for is the import of perceptual tokens into the cognitive system. As with LEGO blocks, constructing the concept “cup” from perception is like designing a LEGO cup, with the restrictions that this brings to the use of the concept (you can mainly use the cup to use it as a cup, and not to, e.g., build a LEGO house from LEGO cups).

One already alluded to example comes from Burns and Shepp’s (1988) study on color vision. Their main idea is that the three defining features of any given color – its brightness (value), saturation (chroma), and hue – are difficult to selectively attend to for an untrained observer since color perception is the perception of quite holistic stimuli. If this were the case, then one would expect that test subjects would have difficulty separating these dimensions when comparing a range of samples. This is just what their experiments demonstrated. In their study, Burns and Shepp also found that differentiating brightness and saturation is easier for trained individuals, such as artists. Landy and Goldstone take this as evidence for the creation of new feature detectors: if there was only one detector for “color” before the training, and the subjects were able to differentiate the brightness and saturation of a range of colors after their training, then the perceptual task must have been the cause for feature learning, and for the creation of new perceptual, discriminatory capacities. And surely, if a person did not know the difference between the brightness and the saturation of a color before, and could make a discriminatory judgment after the study, then a new concept has been learned.

A final, important aspect of Goldstone’s proposal is that his and his colleagues’ studies do not rely on predetermined, fixed stimuli sets, but on totally novel ones that often cannot readily be parsed into already known structural elements. An example can be found in Schyns and Murphy’s (1994) “Martian rocks” studies. The study employed various black blobs with several kinds of round or prolate appendices, without any indication of possible fragments, or parts of the whole object.

In their argument for using such alternative materials, Schyns et al. (1998) express that they want to exclude the possibility of using known categories in their experimental tasks, and that they want to get a better understanding of the ways in which totally new categories are learned. If a shape is almost certain to not represent any possibly innate, fundamental shape primitive, then it should be very likely that learning to pick out that shape is a case of learning something new.

With alternative materials, many different interpretations are possible, there are multiple features that could be encoded, and the analog format (as opposed to the digital signs one also finds in fixed feature experiments) make it possible to study something akin to real-life concept learning, where the interesting or learn-worthy features also are not as plainly recognizable.

With this picture of Landy and Goldstone’s (2005) reply to Fodor’s challenge in mind, let me now confront the question if their proposal can stand up to the challenge.

CONCEPT LEARNING OR CONCEPTUAL CHANGE?

Given that the perceptual learning approach can do the things described above, does it actually answer Fodor’s Challenge for concept learning? I want to argue that it doesn’t, because of two problems. Goldstone and colleagues have thus far left decisive questions pertaining to the central elements of their account, namely the details of the integration of perceptual symbols into the representational system, and the role of features and stimuli in that process, unanswered.

First and foremost, it is not clear whether the model provides prospects for concept learning at all. One might agree that the phenomena of unitization and differentiation are a form of learning, since they are mechanisms of restructuring previously available categories, and thereby they are means of grouping information in new ways that might lead to new beliefs. Consider somebody who finds out that two animals which she knew to be dogs belonged to two different breeds, say Labrador Retrievers and Dalmatians. This could clearly count as learning something one had not previously understood. But does it really count as introducing two new “psychological primitives”? An alternative view along fixed-feature lines would be to grant that “Labrador Retriever” and “Dalmatian” are indeed new, but only as names for two objects that had already been processed in thought in a different way, say as p and q (the letters standing for the symbols standing for the individual dog tokens). So, what has been added were not new symbols, but rather new labels for old symbols, or new beliefs about these symbols, as in “ p is a token of ‘Labrador Retriever.’”

It has been suggested to me to look at a more perceptually taxing kind of differentiation process, since this might support Landy and Goldstone’s position. Suppose that, in a psychological experiment, a subject is rewarded for identifying tokens of pacman shapes with a 92° “mouth” angle, and not rewarded if she chooses pacman shapes with a 90° “mouth” angle. Wouldn’t one want to say that learning to appropriately keep those two shapes apart counts as learning a new concept? While the example is intriguing, and representative for a class of psychological experiments on categorization, I would argue that it does not count as a case of learning a new concept of, say, “92° pacman.” Since the goal of the subject lies in getting the reward, it seems more appropriate to speak of the concept “choice that gets me a reward” and the perceptual input that is related to a token of the concept – a choice of a given answer, say “A” or “B” (if “A” and “B” stand for the answers related to the respective 92°/90° pacman). If one carries this thought further, the discriminatory input does not become involved in the conceptual content of the concept that is applied in the task – a 92° pacman is not a token of “choice that gets me the reward,” but it is a prompt to apply the concept by acting in a certain way. This is not to say that fine perceptual discriminations can never become conceptually relevant, or the topic of conceptual development: examples like Smith and Kemler’s (1978) study of changes in the integrality of dimensions such as color and shape surely count as evidence to the contrary. The theoretical status of these developmental changes, however, is exactly the topic of this paper,

even though I presently cannot go into a detailed treatment of the developmental literature for lack of space.

Returning to the dog example, the tricky question for the fixed-feature computationalist at this point, therefore, is not whether the new dog breeds were learned, but rather how the symbol, the name, and the object that the symbol and the name denote are causally related. This is a thorny question for philosophers like Fodor (1998), as they have to defend a very specific type of metaphysical theory of causation to make their analysis stick (see, e.g., Cowie, 1998 for discussion). Perceptual learning would evade this metaphysical question (“How can an innate symbol refer to anything that was encountered perceptually?”), but at the price of creating a psychological one (“How is it possible to import new cognitive symbols from perceptual origins into a $Language_C$?”), which will create the second worry that is identified further below.

At this point, one might be tempted to postulate that there are several kinds of learning: in one kind of learning experience, some genuinely new primitive psychological token (be it a new feature, or symbol, or whatever kind of enrichment one might be interested in) is incorporated into the cognitive system. A case in point would be turning a perceptual stimulus into a cognitively usable symbol that can be used for category judgments, forming thoughts, or other conceptual tasks. Here, something that has not been part of the $Language_C$ would be transferred into that same language. Fodor’s challenge is concerned with this kind of learning.

In another kind of learning experience, the available pieces of information get re-ordered, linked to other bits of information, or get categorized in a finer raster. Strictly speaking, nothing new enters the store of cognitive symbols, but the differentiation between different kinds of already available symbols will be finer, or coarser, depending on the type of change. The question remains if Landy and Goldstone would be happy with “only” providing a model for the second type of learning, since the aim of their article, in their own words, clearly was to give a model for the first type:

our alternative to fixed-primitive languages involves not giving up computationalism, but enriching it with mechanisms which allow the construction of new psychological primitives that are not just combinations of other known categories. (Landy and Goldstone, 2005, p. 347)

On one interpretation of the perceptual learning approach, the main processes of unitization and differentiation seem to fail to introduce new concepts, since they only operate on existing concepts, which are modified to be either more general or more specific regarding certain features of a given category. As Landy and Goldstone (2005) openly state, “feature creation simply involves alterations to the organization of stimulus elements into features” (Landy and Goldstone, 2005, p. 349). But a more strict computationalist, or a Nativist, could easily argue that this process does not strictly speaking add any new information to the cognitive system, as, e.g., Fodor (2008) does. Rearranging old concepts, on this view, cannot be counted as learning, since there is no new information added, but only a regrouping of old concepts. Like in the dog case above, there would only be the addition of new labels for objects that have previously been parts of the $Language_C$. If, however, one wants to object to this analysis, and maintain that unitization and

differentiation mainly work on percepts, then the worries raised in the next part of this paper will apply.

A variant of Fodor’s hypothesis-testing paradox can be formulated that transfers the point into the feature-based learning Goldstone endorses: in order to categorize a stimulus as being evidence for/being a token of a certain psychological feature, one needs to know what feature that is – in order to perceive a sound as a fire fighter siren, one needs to know what a fire fighter siren sounds like (fourth interval). And to know that, one surely needs a feature category that is available before having a stimulus to categorize accordingly. If there is not more to what perceptual learning can do for our understanding of what concept learning should be, then it does not have an explanatory advantage to Fodor’s Nativism in this regard, and need not even be incompatible with his large and fixed basic vocabulary of the mind – the primitive symbols being given, while more elaborate concepts might well be composed by mechanisms like unitization and differentiation.

Gauker (1998) addresses a similar worry concerning Schyns et al. (1998). Gauker in fact poses a double dilemma to any “concepts as (composed out of) features” approach. Suppose that concepts are composed of features, and that learning a concept involves learning a certain amount of the properties that are associated with the concept. Learning the concept “bird” might be linked to associating the concepts “has wings,” “has a beak,” “flying animal,” or any other combination of attributes, to the concept “bird.” Yet, this would require these attributes to be developmentally more basic than the concept “bird.” How could that be? Gauker sees only two possibilities, which pose a dilemma for Goldstone. Either one accepts that there is a developmental hierarchy of features and concepts. This option is in principle open to anybody, but has special advantages for fixed-feature theorists. One could postulate that there are certain primitive features at the basis of the more elaborated conceptual constructions that are learned in cases like learning the concept “bird.” These features are already parts of the cognitive system that did not need to be learned, and so would form an (in some way) innate basis for our more superordinate concepts. If these features were innate, or pre-specified, they would be fixed. As laid out above, Goldstone wants to have some room for flexible features, so only relying on this option does not seem to be a viable option, especially since fixed-feature theorists might just postulate a big enough or flexible enough primitive basis of concepts that could really ground any supposedly perceptually learnable concept. The other option is to deny that feature concepts need to be more primitive than the superordinate concepts – Gauker associates Schyns et al. with this view. The problem for this view is that it requires an explanation of exactly how “truly new features are created” (Gauker, 1998, p. 27) – features that do not have a previous history of, e.g., having been fused with other features, forming a less differentiated category.

Landy and Goldstone (2005) attempt to answer these kinds of criticism by pointing out the changes that they have observed in their studies and in similar studies. They cite evidence that “early perceptual devices can be systematically and physically altered by the environment to change their representational capacities” (p. 351) to support the claim that new features can be created. For example, in simulations by Rumelhart and Zipser (1985), connectionist systems were able to create new detectors for different kinds

of stimuli in a competitive learning task. But while the evidence might be supporting this claim, it certainly need not support the connected claim that such a change in representational capacities causes changes in the Language_C , and by this causes the learning of new features. A fixed feature theorist might be very happy with the first claim, linking it to the activation or triggering of a certain store of symbols that affects early perceptual devices: environmental influences would first cause changes in the (already present) symbolic system, which would in turn result in changes in perception. The change in representational capacities thus might just be a change in frequencies of triggering certain symbols. One might call this a form of learning, since there would be changes in the perceptual domain, but the corresponding changes in the conceptual domain – starting to use previously available symbols for hitherto unperformed perceptual tasks – would not be substantial enough to warrant the label “concept learning.”

Adding to this point, one can enlist another example of a learning system that Landy and Goldstone (2005) briefly discuss in their, and which they revisit in Goldstone and Landy (2010): the Pask device⁴. A Pask device is “an array of electrodes partially immersed in an aqueous solution of metallic salts” (Landy and Goldstone, 2005, p. 351) that will physiologically change when electric currents are applied. Now, changes in electrical configuration in the device come with changes in functionality – the device will start reacting discriminatively to two kinds of sound frequencies: a “new ear” for the circuit has been trained while it got constructed. From Fodor’s perspective, it would however be a mistake to call this learning. These changes have all the characteristics that brute-causal acquisition of concepts in humans also has, so by definition, they do not amount to concept learning. This issue is independent of the question whether the Pask device actually is a representational system – if a certain reaction to frequency A counts as representing that frequency. Following Prinz and Barsalou (2000), I am inclined to regard the Pask device as a representational system, and in that sense a fitting analogy to a cognitive system. Focusing on the question whether the Pask device’s development of an electronic ear is more like knowing the difference between smoke and steam after being hit on the head or more like learning from observation about the difference, I submit that it is decidedly more like the former, and thus not a case of learning in Fodor’s sense⁵.

Based on this reasoning, it seems right to focus on unitization and differentiation as the main players in feature learning when seen as concept learning, while acknowledging that Fodor’s kind of brute-causal acquisition – as demonstrated by human analogs to the Pask device’s “learning” – plays a transformative role in human cognition that might be seen as enabling concept learning.

⁴The example of the Pask device has been introduced into the feature creation literature in the section “Authors’ response” of Schyns et al. (1998), where it is used to link the ideas of perceptual learning and of emergent properties in learning.

⁵This investigation invariably leads to the question whether perceptual learning in humans should count as Fodor-type learning or as brute-causal acquisition. Discussing this point is beyond the scope of the current paper, but can be developed into a different argument against the methodological set-up of Fodor’s challenge – the distinction between brute-causal and rational acquisition might not be as helpful as Fodor would like it to be, and cases like the Pask device might work in favor of giving up the distinction altogether.

Up to this point, the perceptual learning approach has not succeeded in answering Fodor’s challenge, since the alternative fixed-feature theory has been shown to give equally powerful explanations of phenomena like changes in representational capacity, while not having the problem of having to explain how new cognitive symbols could be created from perceptual materials. Also, the perceptual learning approach has not given a full model for the latter task, and thereby only stands on a partial base of providing good explanations for the influence of the conceptual on the perceptual. There is however another set of conceptual problems that call for a resolution before the perceptual learning approach can get off the ground and before we can assess whether an inference to the best explanation would support the fixed-feature approach or the perceptual learning approach.

THE NOTIONS “FEATURE” AND “STIMULUS”

A second worry that directly follows from the first one has to do with the notion of features and stimuli in concept learning. In Goldstone’s theory, concepts are (created out of) features; features are created from stimuli. Stimuli are in a format that is supposedly compatible to the symbolic vocabulary of cognition. Compatibility is a decisive criterion here, and can be defined as follows:

Compatibility =_(Def.) A set of symbols S is compatible to a cognitive mechanism M iff inputting S into M yields a (symbolic) output S_O which can be used by further cognitive mechanisms. A set of symbols S is compatible to a second set of symbols S_2 iff S and S_2 can both serve as input for M individually, and iff symbols from S and S_2 in combination also yield a symbolic output S_O which can be used by further cognitive mechanisms.

Without a fulfilled compatibility criterion, it would not be possible to incorporate the perceptually based symbols into the previously available and exercised cognitive activities. Speaking in terms of Language_C , all symbols need to be combinable to form correct sentences from them⁶. Or again in Landy and Goldstone’s (2005) terms:

the claim that novel perceptual features can be learned sounds murky, or even mystical, without the clarification that the novel features are always drawn from a larger, more expressive, more primitive language embodying the physical and pre-conceptual constraints on what can be incorporated into features in the first place. (Landy and Goldstone, 2005, p. 348)

What is a stimulus then? In any dictionary of behavioral and cognitive science, one will find descriptions of perception in terms of proximal and distal stimuli. The distal stimulus of a visual experience might be the tree one sees, whereas the proximal stimulus would be the light reflection arriving at the eyes. At which level do, e.g., Landy and Goldstone (2005) individuate stimuli? This question is pertinent to our present investigation since Landy and Goldstone should address it to make clear where exactly they would see the origins of perceptual experiences, and with that the origins

⁶This point is related, though not identical, to the issue of the compositionality of thought, raised, e.g., by Fodor and Pylyshyn (1988).

of perceptually based concepts: are they in the world (i.e., distal stimuli) or are they in sensory activations (i.e., proximal stimuli)?

Just looking at their (Landy and Goldstone, 2005) paper, they discuss examples of roughly half-moon shaped figures combined of five segments and call these objects stimuli. In another example, they talk about “pieces of physical information [that are related to, or] packaged together in the same psychological feature” (p. 349) and use this as a synonym for “feature.” These are exemplary – or metaphorical, respectively – descriptions of what features or stimuli should be, and yet, these are the most concrete mentions of those terms. A look at earlier renderings of the theory might help. Schyns et al. (1998) commit themselves to the following characterization of the meaning of the term “feature”:

The term “*feature* will refer to any elementary property of a distal stimulus that is an element of cognition, an atom of psychological processing. This does not imply that people are consciously aware of these properties. Instead, features are identified by their functional role in cognition; for example, they allow new categorizations and perceptions to occur.” (Schyns et al., 1998, p. 1, original emphasis)

Here, first, features are described as “elementary properties” of distal stimuli. They are also implied to be “elements of cognition,” i.e., Schyns et al. (1998) postulate a transition of perceptual properties into cognitive functions. The second point concerns a property’s role in cognition: it is supposed to be functional. By this description, Schyns and colleagues want to counter the objection that some feature of an object might not enter a perceiver’s conscious awareness and thus should not count as an element of cognition. While the wording in the quote above suggests a definition in terms of distal stimuli, it also alludes to the psychological role of features, which is even more obvious in Goldstone (2003):

A psychological feature [...] is a set of stimulus elements that are responded to together, as an integrated unit. That is, a feature is a package of stimulus elements that [...] reflects the subjective organization of the whole stimulus into components. (Goldstone, 2003, p. 242)

So, I suggest we should understand Landy and Goldstone (2005) as taking proximal stimuli as the larger background language. Still, there is the unanswered question how these perceptual signals are transferred into language-like symbols that can be used in the same cognitive operations as either innate or previously acquired symbols. How are these vocabularies matched to each other? Let me push the language analogy a little further with an example: suppose the Language_C is like English – approximately every English word corresponds to a mental symbol. Now suppose that the cognitive vocabulary gets enriched with a number of specialized features developed from a perceptual task, like learning chess moves. Perceptually learning a chess move, as opposed to learning it from a written description, might work as follows in the case of the rook’s permissible movements on the board: straight lines along the horizontal and the vertical axis, but not on diagonals. The correct movements can be observed by watching rooks in a large sample of chess moves, or video clips from chess matches, also with a variety of token rooks (made from different materials, or shaped in a variety of ways), and one might even learn how to tell

whether a chess piece is a rook or a queen. The interesting question then is: would a thought about a situation in a chess match be multimodal – would it involve the perceptually learned symbol for the rook as well as the previously available non-domain-specific vocabulary? Let’s take “[]” as a replacement for the perceptual symbol related to the rook moving one square to the left, just for this example, and phrase a thought like “If the rook moves one square to the left, the players will stop playing” multimodally: “If [], the players will stop playing” (given, e.g., that the result is a checkmate). Is it possible to infer the consequent of this conditional from being presented with a representation of []? While this last question might be for further empirical studies to decide, it already hints at the more general worry about perceptual tokens of some sort and their role in cognitive operations: given that originally, a certain cognitive function is performed by a mechanism using symbols of a (possibly innate) Language_C, how can the mechanism adapt to new symbols being introduced into it and filling that cognitive function? That is, how can a perceptual symbol store and an innately fixed symbol store become compatible, as defined above? Landy and Goldstone do not offer a model for this, and so I conclude that, as it stands, the construction of variable-feature language has not been sufficiently based on a model of transferring perceptual symbols into conceptual systems.

Sticking to the notion of features as primarily relevant to building mental representations, one could bring the theory of feature detectors into play, as in Barlow (2001). This, specifically in visual perception, would be an (obvious) way out of the problem, yet with the (obvious) problem that the feature detectors would have to be tuned to some specific inputs, and then the question arises again: How did the feature detectors come into being, and how did they get the tuning they exhibit? Feature detectors are an instance of (possibly innate) processes that one might use to explain concept learning, under the assumption that concepts need not be built into a system as long as there are built in processes that will be able to import concepts into the system. The problem, however, remains the same: a nativist can always argue that built in processes need to be tuned to their inputs in some way. In the perceptual learning literature (broadly defined), one can find several references to Gestalt laws (Schyns and Murphy, 1994; Quinn et al., 2006; Bhatt and Quinn, 2011), and even proposals to explain that acquiring Gestalt laws is possible through neural network simulations (Gerganov et al., 2007). Yet, how should adherence to a given Gestalt law, e.g., good continuity, be possible for a cognitive system without having concepts that are able to express the law and that would help classifying perceptions according to the principle?

Appealing to maturation, or another form of innateness, would yield no new variable-feature language in the sense of Landy and Goldstone (2005), as its large explanatory trump is its independence of the specific type of stimuli the system is confronted with. A classic example in point of this could be imprinting in newborn ducks, as discussed by Fodor (1981): any moving object will trigger the “concept” “mother,” having the duck following the moving object. The important thing is that the duck’s sensory system is predetermined to follow the closest moving object.

If on the other hand the specific stimuli play a role, or matter in some sense for the creation of psychological primitives, how do they influence the creation of a feature detector?

We seem to have gone the circle back to the original question, and the appeal to a larger background language of stimuli has not advanced us very far. Maybe, looking at the problem from a different perspective, under the heading of “physical information” will clear things up. After all, Landy and Goldstone (2005), also refer to the materials from which to get new features as “physical information” (p. 349).

Now, the task is to disambiguate the notion “physical information” and to link it to the question of the proximity of stimuli. Either, physical information, or a bundle of features, is supposed to be informationally structured before or in being perceived. This would require a form of direct realism, or of direct perception: for example, a given object affords to be perceived as a tree, so we as perceivers pick up the right kind of information in order to treat it as a tree. Or, on the other hand, physical information is the cognitive content that has been extracted from the experience; this could be something like a representation of a tree. The retina has registered a certain image, sent it to the visual cortex in one way or another and there, a representation of the tree is formed, or accessed, or activated. In line with the above decision to talk of stimuli in terms of proximal stimuli, it seems sensible to choose the interpretation of physical information as cognitive content. This interpretation, however, just skips the interesting question, which is: “How does a representation of a new object come to be included in a cognitive system?”, and leaves the field open for any kind of Nativist reply to the effect that the representation was triggered by some process, or that the perceptual stimulus got paired with an arbitrary symbol from the wealth of symbols in the representational mind. To avoid this, any proponent of perceptual learning has to go the long way and show just how perceptual content enters cognition and by which means new symbols, or new bits of a feature-“language,” are added to the system. Landy and Goldstone do not offer a model for this form of feature learning, and so I have to conclude that – while it is not conceptually excluded that such a model is possible – their proposal still has some way to go before it can pose a fully developed alternative to their fixed-feature opponents.

CONCLUSION: PERSPECTIVES FOR AN AMENDMENT OF THE PERCEPTUAL LEARNING APPROACH

To amend the theory, one would need to say more about the input systems to unitization and differentiation, and be clearer on the representational format that they are able to operate upon. Specifically, the following questions are still unanswered.

How can a cognitive mechanism that was presumably first stocked with innate computational symbols grow to work with learned perceptual features as input to and vocabulary for its activity? And is it possible to mix symbols of different origins and formats (amodal/modal) – to have “multi-lingually” integrated cognition?

Until the issues raised in this article have been addressed, the proposal does not deploy its full potential to threaten a fixed-feature approach à la Fodor: even if both approaches can be

construed as having similar levels of explanatory power, one of them satisfies Fodor’s challenge while the other one does not yet overturn its empirical premise. The disadvantages that stem from the problems identified in this investigation weaken the perceptual learning approach’s appeal and thereby put its opponent in the stronger argumentative position for now.

After this discussion, one might also be tempted to conclude that the notion of a feature language, flexible, or variable, is misguided as it invariably brings the issue of translation into the debate. How to translate a stimulus (and which stimulus) into a mental symbol? Also, it suggests an ordered, or “grammatical” structure in the non-mental/physical world that is the object of perception. This is dangerous, because the world as appearing to us might not actually be best carved into the perceived (natural?) kinds, but into theoretical kinds that we only perceive through mediation. One does not see the chemical structure of object *x* without being somewhat of a trained chemist, if at all, or at least it is not clear in the Perceptual Learning approach whether the interaction between perception and cognition leads to such depths of theory-ladenness of perception (as in *seeing* chemical structure) as opposed to a quick-and-dirty inferential connection between perceiving certain visual properties of a chemical sample – e.g., observing a deep green flame when burning a sample of a chemical powder, and identifying it as copper(II)-sulfate (maybe a perceptual–cognitive process), as well as the sample’s being a salt (an inference from that observation). In making the distinction between perceptual–cognitive and inferential, the adherence of a given process to perceptual constraints might indicate that the process is of the former type, whereas observations following theoretical, or “conceptual,” rules can be properly classified as the latter⁷. Still, the distinction is not always a clearly cut one.

This does not just touch upon Landy and Goldstone’s (2005) proposal, but more generally on their still dominating opponents: if the language metaphor does not work for features, as one might conclude from the problems raised in the previous section, then why should one be inclined to see the strong analogy between computers – symbol crunchers – and human minds with brains and nervous systems underlying them (in one way or another) as necessary? Perhaps the mind only becomes symbolic by starting to use symbols, but does not reflect this symbol-mindedness in the elements of cognition. Dissociating materials, or vehicles of thought, on the one hand, and thought contents on the other hand, might be a prudent move until a clearer picture of the connections between vehicles and contents is available.

ACKNOWLEDGMENTS

I would like to thank Mark Sprevak and Tillmann Vierkant for their help and their extensive comments on several drafts of this paper. Also, I would very much like to thank the referees for their constructive criticism, and for pointing me toward a wealth of additional literature. Finally, I want to thank the editors of this special issue for their support.

⁷It was Robert Goldstone who pointed this out to me.

REFERENCES

- Barlow, H. (2001). "Feature detectors," in *The MIT Encyclopedia of the Cognitive Sciences*, eds R. A. Wilson and F. C. Keil (Cambridge: MIT Press), 311–314.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behav. Brain Sci.* 22, 577–660.
- Bhatt, R. S., and Quinn, P. C. (2011). How does learning impact development in infancy? The case of perceptual organization. *Infancy* 16, 2–38.
- Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* 94, 115–117.
- Burns, B., and Shepp, B. E. (1988). Dimensional interactions and the structure of psychological space: the representation of hue, saturation, and brightness. *Percept. Psychophys.* 43, 494–507.
- Cowie, F. (1998). Mad dog nativism. *Br. J. Philos. Sci.* 49, 227–252.
- Fodor, J. A. (1975). *The Language of Thought*. New York: Thomas Crowell Publishing.
- Fodor, J. A. (1981). "The present status of the innateness controversy," in *Representations*, ed. J. Fodor (Cambridge: MIT Press), 257–316.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. A. (2008). *LOT2: The Language of Thought Revisited*. Oxford: Oxford University Press.
- Fodor, J. A., and Pylyshyn, Z. W. (1981). How direct is visual perception? Some reflections on Gibson's "ecological approach." *Cognition* 9, 139–196.
- Fodor, J. A., and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition* 28, 3–71.
- Gauker, C. (1998). Building block dilemmas. *Behav. Brain Sci.* 21, 26–27.
- Gerganov, A., Grinberg, M., Quinn, P. C., and Goldstone, R. L. (2007). "Simulating conceptually-guided perceptual learning," in *Proceedings of the Twenty-Ninth Annual Conference of the Cognitive Science Society* (Nashville: Cognitive Science Society), 287–292.
- Gibson, E. J. (1963). Perceptual learning. *Annu. Rev. Psychol.* 14, 29–56.
- Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *J. Exp. Psychol. Gen.* 123, 178–200.
- Goldstone, R. L. (1998). Perceptual learning. *Annu. Rev. Psychol.* 49, 585–612.
- Goldstone, R. L. (2003). "Learning to perceive while perceiving to learn," in *Perceptual Organization in Vision: Behavioral and Neural Perspectives*, eds R. Kimchi, M. Behrmann, and C. Olson (Mahwah: Lawrence Erlbaum Associates), 233–278.
- Goldstone, R. L., and Barsalou, L. W. (1998). Reuniting perception and conception. *Cognition* 65, 231–262.
- Goldstone, R. L., Braithwaite, D. W., and Byrge, L. A. (2012). "Perceptual learning," in *Encyclopedia of the Sciences of Learning*, ed. N. M. Seel (Springer), 2580–2583.
- Goldstone, R. L., and Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdiscip. Rev. Cogn. Sci.* 1, 69–78.
- Goldstone, R. L., and Landy, D. (2010). Domain-creating constraints. *Cogn. Sci.* 34, 1357–1377.
- Goldstone, R. L., Son, J. Y., and Byrge, L. (2011). Early perceptual learning. *Infancy* 16, 45–51.
- Landy, D., and Goldstone, R. L. (2005). How we learn about things we don't already understand. *J. Exp. Theor. Artif. Intell.* 17, 343–369.
- Newell, A., and Simon, H. A. (1976). Computer science as empirical inquiry: symbols and search. *Commun. ACM* 19, 113–126.
- Prinz, J., and Barsalou, L. W. (2000). "Steering a course for embodied representation," in *Cognitive Dynamics: Conceptual Change in Humans and Machines*, eds E. Dietrich and A. Markman (Cambridge: MIT Press), 51–77.
- Quinn, P. C., Schyns, P. G., and Goldstone, R. L. (2006). The interplay between perceptual organization and categorization in the representation of complex visual patterns by young infants. *J. Exp. Child. Psychol.* 95, 117–127.
- Rumelhart, D. E., and Zipser, D. (1985). Feature discovery by competitive learning. *Cogn. Sci.* 9, 75–112.
- Schyns, P. G., Goldstone, R. L., and Thibaut, J.-P. (1998). The development of features in object concepts. *Behav. Brain Sci.* 21, 1–17.
- Schyns, P. G., and Murphy, G. L. (1991). "The ontogeny of units in object categories," in *Proceedings of the XIII Meeting of the Cognitive Science Society* (Hillsdale: Erlbaum), 197–202.
- Schyns, P. G., and Murphy, G. L. (1994). "The ontogeny of part representation in object concepts," in *The Psychology of Learning and Motivation*, ed. D. L. Medin (Waltham: Academic Press), 305–349.
- Schyns, P. G., and Rodet, L. (1997). Categorization creates functional features. *J. Exp. Psychol. Learn. Mem. Cogn.* 23, 681–696.
- Smith, L. B., and Kemler, D. G. (1978). Levels of experienced dimensionality in children and adults. *Cogn. Psychol.* 10, 502–532.

Conflict of Interest Statement: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 30 October 2011; accepted: 13 March 2012; published online: 30 March 2012.

Citation: Stöckle-Schobel R (2012) Perceptual learning and feature-based approaches to concepts – a critical discussion. *Front. Psychology* 3:93. doi: 10.3389/fpsyg.2012.00093

This article was submitted to *Frontiers in Cognition, a specialty of Frontiers in Psychology*.

Copyright © 2012 Stöckle-Schobel. This is an open-access article distributed under the terms of the Creative Commons Attribution Non Commercial License, which permits non-commercial use, distribution, and reproduction in other forums, provided the original authors and source are credited.