First thoughts

Daniel A. Weiskopf

Abstract: Jean Mandler proposes an original and richly detailed theory of how concepts relate to sensory and motor capacities. I focus on her claims about conceptual representations and the processes that produce them. On her view, concepts are declarative representations of object kind information. First, I argue that since sensorimotor representations may be declarative, there is no bar to percepts being constituents of concepts. Second, I suggest that concepts track kinds and other categories not by representing kind information per se, but rather by being subject to the appropriate sort of inferential dispositions. These dispositions themselves may apply equally to perceptual and non-perceptual representations. Third, I argue that Mandler's proposed redescriptive mechanism for producing conceptual primitives can be viewed as a kind of Fodorian triggering device. Hence there may be less distance between her view and Fodor's than either one has supposed. I suggest that redescription needs to be supplemented with several other kinds of more flexible and open-ended concept learning mechanisms. Finally, I briefly sketch the view of conceptual development that results from adopting these proposals and contrast it with Mandler's.

Having concepts is a special kind of cognitive achievement. Not every creature that represents its environment does so conceptually. On some views, only humans are capable of true concept possession (Davidson, 1982/2001). This extreme position aside, it seems certain that our conceptual capacities far exceed those available to any other terrestrial creatures. How do we come by this achievement? Putting the question in a way that contemporary developmental psychology can get some purchase on, we might ask: when do human infants and children first begin to display *distinctively conceptual* capacities?

Much of Jean Mandler's research has been focused on precisely this question. The fruits of that research, summarized in her 2004 book, *The Foundations of Mind*, constitute an enormously impressive achievement: a view of the overall place of concepts in our mental lives as well as an empirically detailed picture of how they emerge from non-conceptual processes and

mechanisms and how they relate to faculties such as perception, memory, and language. The target article extends and enriches one thread of this research, having to do with the nature of our primitive concepts. Specifically, she argues that the first concepts we come to possess may be both (1) relatively late in developing, compared to our skills at perceptually discriminating categories; and (2) relatively global, or high-level, as opposed to belonging to the so-called 'basic level' (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976).

I can't hope to do justice to the richness of Mandler's position, or the evidence she assembles for it, in this brief space. So I will focus here on several related aspects of her overall theory of conceptual development. In Section 1 of the paper I lay out the conditions that Mandler proposes on what makes something a concept. I then criticize the representational distinction she draws between perceptual and motor states (which I will generally label *percepts*) and conceptual states. I suggest that the one condition Mandler uses to distinguish percepts and concepts—namely, being declarative as opposed to procedural—doesn't in fact do so. In particular, I'll argue that some percepts can be non-procedural, and hence may (by her criterion) be conceptual.

In Section 2 I turn to the claim that the main distinction between percept and concept lies in their content. Again, I argue that the condition as stated doesn't do the job of drawing the appropriate distinction. But something like Mandler's proposal seems to be on the right track. I suggest a way of modifying her criterion in a way that draws the percept-concept boundary not in terms of content, but rather in terms of functional role. Once we see that this functional distinction is key, we open the door to a modest empiricist position on which our first thoughts and concepts, if not those that come later in development, might be encoded in a perceptual format.

Finally, in Section 3, I examine the mechanism by which Mandler claims that perceptual information is recoded or redescribed in a more abstract conceptual format. I argue, first, that this mechanism is one that concept nativists like Fodor could be quite happy with, and second, that we will need to supplement Mandler's mechanism with other concept acquisition mechanisms in order to do justice to the richness of our conceptual repertoire. The overall picture that emerges from this discussion is one that is compatible with much of what Mandler says about conceptual development, but which takes a different view both of the first conceptual thoughts that we entertain, and of the ways that our conceptual repertoire can develop once it gets an initial foothold.

1. On sensorimotor representations

No discussion over when concepts emerge in development proceeds without some characterization of what it takes for something to be a concept. In the target article, Mandler proposes that concepts are constituted by "declarative knowledge about object kinds and events that is potentially accessible to conscious thought" (Mandler, 2007, p. 3). That they are declarative is a fact about the *format* in which they are represented. That they refer to kinds and events is a fact about their *content*. That they are potentially conscious is a fact about how they are *processed*. I'd also note that, for Mandler, the third condition is strictly speaking redundant, since conscious accessibility is purportedly entailed by declarative encoding (Mandler, 2004, p. 55). I'll first discuss the necessity of the format and processing conditions before turning, in the next section, to the conditions on conceptual content.

We can begin by considering what something's being *non-conceptual* consists in. On Mandler's taxonomy, sensorimotor knowledge is procedural, and procedural knowledge is non-conscious. One immediate question that arises is whether it's really true that sensorimotor

representations are invariably non-conscious. The canonical evidence for this claim, on the sensory side, involves our inability to produce accurate drawings of extremely familiar common objects like coins and human faces (in the absence of actual sample items, anyway), despite being able to access stored perceptual information about these objects in other contexts, e.g., face recognition. On the motor side, we are typically unaware of the stages that go into our fluid performance of acts like tying our shoes or shifting a manual transmission. The fact that we can't access the information the underlies performance on these tasks suggests that it is represented in a distinctive format that prohibits its entering consciousness.

Note that something's being non-conscious doesn't by itself entail that it's procedural. There might be many cognitive states that aren't conscious (e.g., states of the language faculty or our intuitive numerical system for calculating quantity; on the latter, see Dehaene, 1997) but also aren't procedural (since these systems might all employ symbol systems that encode propositional information). So even if it were true that sensorimotor processing was non-conscious, it wouldn't follow that it's invariably procedural. However, given the inclusion hierarchy sketched in the last paragraph, Mandler can agree with this and still claim that, nevertheless, sensorimotor representations are non-conscious, since they *are* procedural.

But I think we can argue directly that merely being sensorimotor doesn't entail that a representation is non-conscious. Even granted that motor skills and certain forms of perceptual skill (e.g., reading upside-down text, identifying fighter jets by their silhouettes) are usually assumed to be acquired and deployed by procedural memory systems, it still might be the case that some sensorimotor representations are consciously accessible.

Sensorimotor imagery provides one example of this phenomenon. While there may be some individual variation in how frequently people use imagery or in how vivid and detailed

their imagery is, most people use mental images to help solve problems on a regular basis. It might be that this imagery is not always very accurate—there are, for example, systematic distortions in how we imagine even familiar geographic features being located relative to each other (Tversky, 1981). Mandler's own example of failing to accurately imagine how ears are located relative to eyes on the human head provides another nice case (Mandler, 2004, p. 48). But the fact that this imagery is inaccurate doesn't show that there is no conscious access to perceptual representations. Images can be unspecific on many details. Analogous points could be made about motor imagery: my motor images of what is involved in making a free throw in basketball might be fairly inaccurate (something to which my performance testifies), but these images are still consciously manipulable.

This suggests that sensorimotor representations can be conscious. For Mandler, this requires showing that they can be encoded non-procedurally. And, in fact, it seems generally unlikely that all stored perceptual representations of objects are encoded in a procedural format that is unavailable for conscious use. If we accept the drawing test as a criterion, for instance, it seems clear that we can sometimes make accurate drawings of objects that get their gross perceptual features right (e.g., the general bioplan of an organism, such as the layout of its torso, limbs, tail, eyes, ears, mouth, etc.). A sketch of a cat and a dog, even if neither is incredibly detailed or accurate, doesn't depict them as looking similar. A sketch of a toaster and a laser printer also might get their gross featural structure right, including their similarities and differences (e.g., the toaster doesn't have a paper tray). So it doesn't seem likely that all perceptual representation is locked up in a procedural format that can't be accessed by these tasks.

Roughly, on the view I am proposing here, sensorimotor information may either be encoded procedurally or non-procedurally. When we train ourselves to perform rapid perceptual identifications or practice stereotyped motor skills, we generate encapsulated procedural representations that incorporate the needed sensorimotor representations. But the same sensorimotor representations that serve as input to this process of procedure-making remain available for use in other tasks. Packaging them in a procedural format does not eliminate them; rather, it involves a kind of representational re-description. Once embedded in a procedure, these representations become context-bound and inflexible. But the representations on which these new procedures are based can themselves remain available for other processes to access. So the link between the sensorimotor and the procedural can sometimes be severed. This opens space for sensorimotor representations to be both declarative and conscious, and hence for them to satisfy some of Mandler's conditions on being concepts.

A possible rejoinder to this suggestion is that the information tapped in tasks that are allegedly criterial for non-procedural encoding (e.g., drawing) is in fact encoded *conceptually*, not perceptually. This might explain the lack of access to detailed perceptual information in these tasks. Mandler notes, for instance, that it might be part of our FACE concept that faces are OVAL-SHAPED. The concept OVAL-SHAPED encodes information about how faces look, but not in the same way that an image of a face would. Because these conceptualized descriptions typically omit many specific details that are present in percepts themselves, we could explain how people might generate drawings that are roughly accurate at the level of large structural features but poor in fine details. The same sort of explanation would go for our image-generation capacities.

The trouble with this rejoinder, though, is that both imagery and drawing can occur under the guidance of conceptualized thought and still involve non-conceptualized representational

elements, e.g., stored percepts. Concepts might serve as retrieval cues for getting perceptual representations out of long-term memory. So it isn't generally true that having a concept that carries certain perceptual information about an object rules out also having a *separate* perceptual representation of that information. In fact, I will argue in the next section that Mandler needs for there to be separately stored, relatively fine-grained perceptual representations not only to explain facts about object recognition, but also in order to account for important facts about concept individuation. For present purposes, all we need to note is that we don't yet have a way to distinguish between (1) people using concepts to retrieve separate, but somewhat inaccurate or degraded, percepts from memory; and (2) people using concepts to guide behavior without relying on stored percepts, hence producing some inaccuracies and distortions.

Let's take stock. I concede that procedural representations are non-conscious. But I've argued that sensorimotor representations need not be procedural, and may be conscious. By Mandler's conditions, this pushes them (potentially) closer to the domain of the conceptual. To show that, *contra* her account, perceptual and motor representations can be part of our concepts, we would need to show that they can also be encoded declaratively and can represent kinds and event types.

Declarative memory can be defined in three ways: by its verbalizability, content, or function. First, declarative memories are supposed to be those that can be verbally articulated. We can linguistically express what it is that we learn or remember declaratively. The content of declarative memories, meanwhile, is supposed to be propositional. These two claims are mutually reinforcing, since the aim of verbalizing is to express a proposition, and presumably most propositions can be so expressed. If, however, there can be propositions that are not readily articulated verbally, then the content condition may allow more memories to count as declarative

than the function condition would. Arguably this is the right approach: we should not assume that every content that can be grasped in thought has a natural form of linguistic expression. Finally, declarative memories are supposed to be multifunctional. They can be acquired rapidly, even on single trials, and the information acquired can be used in a variety of tasks and contexts. This is in contrast with the context-bound nature of procedural memories.

Sensory experience encodes highly specific information about how things look–their color, shape, texture, etc. Motor plans encode specific information about how one's body is engaging with the world—how precisely to tilt one's hand to put a letter in a mail slot, for instance. This information is frequently difficult to verbalize. We lack precise descriptive terms for many perceivable qualities and ways of enacting movements. By the verbalizability criterion, then, these representations cannot be declarative.

However, the verbalizability criterion may be too demanding. First, Mandler claims that concepts are declarative, and hence their content can be verbalized. But it seems plausible that there may be aspects of conceptual content that escape ready verbalization. For example, the use of multidimensional scaling techniques can reveal underlying dimensions along which concepts can be located and ordered (Nosofsky, 1992). The dimensions that these techniques reveal may not be ones that are immediately available for verbalization. To take a well-known example, our bird concepts may be organized according to such features as size and predacity (Smith, Shoben, & Rips, 1974). These may not be the features that we most immediately *verbalize* when asked to characterize particular kinds of birds, though.

Second, it isn't even clear that Mandler's proposed contents for global concepts are readily verbalizable, either. She suggests that concepts like ANIMAL carry the content *self-moving interactor*. But this doesn't seem like the sort of content that we would verbalize if we were to

describe how we conceive of animals. (This despite Mandler's claim that these general concepts prove easier to define than 'basic-level' concepts.) It might, however, be the sort of content that would be revealed by analysis of classification and inference behavior directed at animals. Suffice it to say that if these concepts fail to be verbalizable, then by Mandler's lights we should give verbalizability correspondingly little weight in determining whether a representation is declarative or not.

It is less clear whether the content criterion rules out sensorimotor representations as being declarative. The issue here turns on whether those states encode propositions or not. This in turn depends on what one counts as a proposition, a question that has attracted considerable attention by philosophers. On some accounts (e.g., Peacocke, Tye), the content of experience is non-propositional; rather, it is spatial, analogue, and unit-free. On others (e.g., McDowell, Brewer) experiences may carry content that is extremely fine-grained but still propositional. Resolving this complex debate is beyond the scope of the present discussion (see the papers in Gunther, 2003 for discussion). As we will see in the next section, however, Mandler's account of the format of primitive concepts commits her to thinking that they are encoded in a spatial format. These spatial representations are nevertheless supposed to be able to combine compositionally to form a range of fully propositional thoughts. If abstract spatial representations can be declarative, though, it is hard to see how we might rule out the prospect of other quasi-perceptual representations from also being declarative.

Finally, are sensorimotor representations multifunctional? The evidence already cited suggests that at least some of them are. In executing many tasks, we may find it extremely useful to carry out simulations of what we might see and how we might act in certain situations. This might happen spontaneously, as in on-line text comprehension, or intentionally, as in the use of

imagery in problem-solving. This suggests that this information can be drawn on in a wide range of contexts. By the functional criterion, then, this puts it in the camp of the declarative. And it is arguably the ability of representations or memory systems to function in a relatively context-free way that is most central to their being declarative. If we give this criterion the greatest weight, there is nothing to prohibit some sensorimotor representations from being declarative. This brings them at least part of the way towards potentially being conceptual. For the moment, then, let us set aside considerations of representational format and turn to questions of content.

2. From content to function

Mandler holds that concepts represent different sorts of information than do percepts. I've already touched on one such difference in content: percepts are characteristically fine-grained, while concepts need not be. As she puts it, "perceptual categories contain more detailed information (at least in infancy) than do conceptual categories" (Mandler, 2004, p. 199). Another purported difference in content is that "perceptual categories work on different kinds of information than do conceptual categories. […] There is an important distinction to be made between people's summary representations of what things look like and their summary representations of what things are" (Mandler, 2004, p. 197). Percepts only contain information about looks, while concepts contain information about 'what things are'. As a consequence of this distinction in content, percepts and concepts participate in different sorts of categorization processes: "Perceptual categorization is used for recognition and object identification. Conceptual categories, on the other hand, are used to control inductive generalization and … for recall of the past as well" (Mandler, 2004, p. 199). Concepts are for determining the kind to which an object belongs and drawing inferences from its kind membership, percepts are for recognizing or identifying objects.

The distinction between having a percept and having a concept, then, corresponds roughly to the distinction between knowing what something looks like and knowing what kind of thing it is. One immediate question is whether this distinction can bear much significant theoretical weight. It is possible to interpret the demand that concepts contain kind information in an extremely stringent way. Kind information might be information that captures the essential or defining properties of a kind. But it's wildly unlikely that we have such information for most of our concepts. For many objects, all the information that we have about them may be that they are things that appear thus-and-so. My knowledge of precious gemstones, for example, is more or less limited to facts like this. Yet this information may be enough, in certain contexts, for me to determine what kind of thing I'm confronted with; that I'm looking at a piece of green agate, not a piece of jade, for example. Looks information can, and often does, serve as kind information, at least insofar as it underpins our judgments about kind membership.

According to Mandler, the first genuine concepts children form, somewhere between the ages of 7-11 months, are *global* (or *superordinate*) concepts. These include ANIMAL, VEHICLE, FURNITURE, PLANT, and so on. As noted above, a concept like ANIMAL has the content *self-moving interactor*. What a child understands in grasping this content is that animals are prone to certain patterns of motion and responsiveness to their surroundings. Is this *kind* information, though? Insofar as it just characterizes animals by their typical mode of spatial interaction, it is hard to see how it could be. It certainly isn't a specification of any essential properties of animals (nor does Mandler require that it be). It doesn't, at any rate, seem any more like information about what kind of things animals are than any description of their other perceivable attributes would be.

What this suggests is that the relationship between kind information and looks information posited in Mandler's work is complex. For instance, she claims: "The infant looking at dogs and cats does not at first conceptualize them as two different kinds, but … most likely as two different-looking, self-moving interactors" (Mandler, 2007, p. 7). On one reading, this seems reasonable. If ANIMAL is among the first concepts available, then animals like cats and dogs will be represented as ANIMAL THAT LOOKS *F*, where *F* differs for each different-looking creature. Of course, *F* here is not supposed to be a conceptual representation, but something like a pointer to a complex percept. This actually suggests that on Mandler's own view perceptual information *needs* to be part of concepts. The argument for this runs as follows:

(1) The concepts CAT and DOG are distinct for the preverbal child.

(2) The distinction between the two does not come in terms of what they encode conceptually.

(3) Perceptual information, e.g., LOOKING *F*, can be used to individuate these concepts.

(4) Information that individuates a concept is constitutive of that concept.

(5) Hence, some concepts are partially constituted by perceptual information.

Mandler grants (2), and seemingly grants (3) as well. Premise (4) is close to being a matter of definition. One possible way to avoid this conclusion, then, would be to deny premise (1). The trouble, however, is that this entails claiming that preverbal children don't distinguish in thought between cats and dogs. Then, one might wonder, what justifies attributing to them possession of

these two separate concepts at all? I suggest that, absent an answer to this question, we should accept that some concepts can have some perceptual constituents.

This just raises again the question, though, of what more would be needed to conceptualize these creatures as belonging to different *kinds*? This is especially pressing since the information encoded in these hybrid concepts seems to be entirely about looks. One possibility that seems attractive is that conceptualizing something as a kind doesn't have to do with *how* it is represented at all, but with how one's representations of it are treated. It might be, for example, that whether something is a kind representation involves its standing in the right sort of inferential relations. Kinds are typically thought of as having hidden underlying characteristics that may be present even in the absence of their typical perceptual appearances. So something may be a kind member even if it lacks these typical appearances. Moreover, since these underlying properties are hidden, there may be doppelgangers of kind members that are not genuine kind members. Realistic toy animals and species that mimic the appearances of other species are examples. Kinds—at least 'natural' kinds like gold and goats—can have both atypical members and 'twins'. Our representations of kinds, then, should reflect this fact. In representing something we take to be a kind, we might be sensitive to the fact that even something that satisfies the content under which we are representing that object might not genuinely belong to the kind in question.

What this sensitivity requires is something like the following. If my concept of an F represents Fs as being things that have a certain appearance, I must be willing to accept that (i) not all Fs need to have that appearance, and (ii) some things having that appearance may not be Fs. These paired dispositions might be part of the inferential dispositions that characterize kind concepts, as opposed to looks-based concepts. Mandler seems to have something like these

conditions in mind when she says: “Of course, adults *do* make use of perceptual similarity in their inductions, but they use it to help determine kind, not as a basis for induction in its own right. No matter how much something may look like an animal, if we think (for whatever reason) it is *not* an animal, we will not ascribe animal properties to it” (Mandler, 2004, p. 199).

In fact, I’d suggest that we can generalize these inferential dispositions. Such a generalization is desirable, since not all concepts are kind concepts. Still, something like this way of distinguishing kind concepts might provide a way to separate concepts in general from other mental representations. The basic proposal is this: information encoded as part of our concepts is generally taken to be *fallible* with respect to identifying or drawing accurate inferences about the category that they represent. According to most contemporary psychological models, concepts are locally circumscribed packages of information (Weiskopf, forthcoming-a). In such a package, a concept of Fs may represent them as being G_1, …, G_n; that is, as having a certain complex set of properties. But what constitutes the fact that such a package of information is conceptual is, in part, the fact that no single such package of information is taken to be definitive of what it is to be a member of a category. Indeed, even the totality of ways that we have of representing a category might be fallible. This distinction between how we happen to represent categories and how they actually are gets a foothold only at the level of concepts. Being a concept, then, is at least in part a matter of being processed by a system that treats all of its representations as potentially revisable in the face of new evidence.

An example may help to clarify this point. Suppose one’s CAT concept contains features like HAS A TAIL, IS FURRY, MEOWS, PURRS, IS KEPT AS A PET, etc. This is probably the housecat stereotype most people possess. But we can perfectly well recognize that there can be cats that lack tails or fur, or both, whether due to the particular breed they belong to or some sort of

accident. Some cats may not meow, or may purr little. But they are still all cats in good standing, and we can recognize this without feeling any sense of conceptual incoherence. How can this be, if the very concept CAT encodes these features, however? I suggest that what *makes* this cat representation a concept is in part the fact that it is treated as a fallible guide to the properties of the things it refers to. Hence the fact that a concept F contains certain information G does not rule out our thinking that something might be an F but not G. We can revise our concepts of entities in order to accommodate these new facts. We do so all the time in, for example, certain sorts of counterfactual reasoning where we can imagine all sorts of highly untypical creatures (e.g., talking cats made of glass).

If this proposal is on the right track, then the main difference between percepts and concepts is not fundamentally a distinction of content or format. Rather, it is a distinction in the systematic way that these representations are treated. The perceptual system can represent objects accurately or inaccurately; but it is not in the business of revising its representations of those objects on the basis of new evidence. If the perceptual system has (say) a stored template for how human faces look, those templates are applied on the basis of whether a new entity looks relevantly similar to any particular stored template. The application of these stored templates is what explains the fact that certain things can look like faces to us. Even if the region of space that triggers the activation of a face template doesn't correspond to a human being (it might be part of a cloud or the pattern of bark on a tree), it may still *look* face-like; indeed, this can be the case even if we *know* that it isn't really a face that is being perceived. The question of whether there could be something that *looks F* (or *is F*) but does not fall under this perceptual template doesn't arise for the perceptual system.

However, notice that this proposal is consistent with some perceptual representations being subject to the right sort of concept-defining inferential dispositions. We might encounter a new kind of object and encode a representation of its perceptual qualities without knowing any more information about it (including whether it was, say, inanimate, an artifact, or an animal). This stored percept could serve as our concept of that kind of thing so long as we were able to use it in identifying and reasoning about objects like that, but also inclined to think that there might be objects belonging to the same kind that failed to display that appearance, or similar-appearing things that failed to belong to that kind. This would, on my view, constitute a way in which percepts might be recruited to serve as concepts by being freed from the restricted range of processes that perceptual systems employ and made available in some way to the conceptual system. Just what is involved in this process of 'freeing up' percepts for use as concepts is the next topic to which we turn.

3. Redescription and beyond

So far I've argued that the functional distinction between concepts and percepts may be, but need not be, tracked by differences in content or format. There is no principled bar to percepts themselves, not just representations that encode perceptual information, being part of our concepts. A further question, then, is why some concepts can't be *entirely* composed of perceptual representations. Suppose we agree with Mandler that "[c]oncepts can consist in part of unconceptualized sensorimotor information. This seeming oxymoron expresses that some concepts consist of a spatial description pointing to a sensorimotor experience" (Mandler, 2007). This establishes that concepts can be *partially* constituted by such representations. In fact, if Mandler is right, some concepts *must* be so constituted. Recall that on her account, we begin with global concepts and then differentiate them into concepts of more narrow domains. But we

often lack concepts for the qualities that distinguish these subdomains. These are often *perceived* qualities—having certain specific looks, making certain noises, smelling a certain way, etc. These qualities can be 'hooked onto' a concept to distinguish it from other concepts in the same general family.

But here is the question: if we have a mechanism for hooking sensorimotor contents into the conceptual system, why should we assume that it can only operate to hook these contents onto pre-existing global concepts? There are two alternative possibilities here: either these contents can just be straightforwardly co-opted by the conceptual system, or there is a mechanism for spontaneously generating new 'hooks' to hang them on, without these hooks themselves being general concepts. Either one of these possibilities avoids the need for there to be a complex pre-existing conceptual structure, and each permits there to be concepts that are not just partially, but wholly constituted by perceptual symbols.

On such a view, it might be that the infants in the studies by Quinn, Eimas, and Rosenkrantz (Eimas & Quinn, 1994; Quinn, Eimas, & Rosenkrantz, 1993) are actually acquiring and deploying the concepts CAT and DOG when they learn to distinguish instances of those categories by their looks. Mandler proposes that *attention* is required to trigger the redescriptive process that produces the first global concepts from such perceptual experiences. But it seems plausible that the infants in these studies *are* attending to the stimuli. If attention is required to group together information for use in thought, then, they might well satisfy one key precondition of concept attainment. The other condition is that the information attended to must be re-encoded into a more abstract format. In the present context, though, this requirement itself is what is up for debate, so it would be question-begging to impose it. Whether there can be concepts that don't arise from prior redescriptions is just the issue. Again, we are left with the question: why

can't at least some concepts—perhaps the *first* concepts—be constituted by such attended percepts?

The only remaining objection to the idea that first concepts might be entirely perceptual is that these representations might fail to satisfy some further condition on concepts. One such condition is *compositionality*: concepts comprise a system in which larger structures can be built up out of smaller ones. Hence possessing a concept for a certain category entails that one also possesses the ability to combine that concept with others to form complex thoughts. For instance, possessing CAT entails that one is able to entertain complex concepts such as BLACK CAT, ANGRY CAT, CAT WITH ONE WHITE PAW, CAT UNDER THE TABLE, and so on (as long as one possesses the other concepts like ANGRY, UNDER, etc.). And it might be that infants are able to attend to objects and represent them perceptually—perhaps even to label them as CATS—without yet having this ability.

If this is correct, then the emergence of perceptually based concepts might not coincide with the first emergence of discriminative abilities for categories. The representations underlying these discriminative abilities would need to undergo at least two sorts of functional transformations. First, as noted, they would need to be made subject to the revisability dispositions that were described at the end of the previous section. Being subject to such dispositions is part of what is involved in treating a representation as a device that tracks a category whose nature is not exhausted by its perceptual appearances. And treating categories in this way is, as Mandler and I agree, central to conceptualizing them. Second, as was just proposed, these representations need to be made available for potentially unbounded combination with other representations. Concepts are part of a system of states that has just such

an open-ended combinatorial potential, and not every way of tracking a category is automatically part of such a system.

I don't aim to take a stand here on precisely when a child's representations cross the functional threshold that divides conceptual thought from the variety of nonconceptual ways of representing the world. While this has been a major goal of Mandler's research program, my goal is instead to clarify just what crossing that line entails. It's a consequence of the view that I have been laying out that some perceptual representations might be able to serve as concepts, if they can be co-opted so as to function in the right way in the system as a whole. Some neo-empiricist philosophers and psychologists have gone further than this, however, and claimed that *all* conceptual thought whatsoever is carried out using perceptual vehicles. Thoroughgoing empiricism of this sort seems to me misguided, although I can't argue the case in detail here (see de Rosa, 2005; Machery, 2007; Markman & Stillwell, 2004; Weiskopf, 2007 for further discussion). Instead, I will focus on developing the more modest sort of empiricism sketched so far.

The primary issue between empiricists and their opponents is whether there are any amodal representations in cognition; that is, any representations that are not proprietary to some dedicated sensory system. While Mandler sometimes refers to herself as an empiricist, then (see, e.g., Mandler, 2004, p. 61), it isn't entirely clear whether this is apt. On her view, the main currency of the conceptual system proper is a kind of abstract spatial representation. These can be supplemented, as we've seen, with sensorimotor representations, but on the strictly conceptual side of things we operate with combinations of simplified spatial image schemata. Whether this counts as a form of empiricism depends on whether the peripheral systems also operate with such schemata, or whether they are generated exclusively for use by central systems.

Given that Mandler characterizes the conceptual primitives as 'redescriptions' of perceptual input, and notes that they involve putting information into a different vocabulary from that used by the perceptual systems, it seems clear that she intends these to constitute a format distinct from that used by the senses. The redescription process can in one sense be seen as involving abstraction of information from a perceptual array. Particular objects—living animals, say—move and interact with their environment in a certain way. But (a) not all of the details of their movement patterns are encoded in the general ANIMAL concept, and (b) none of the fine-grained perceptual details of their intrinsic structure are included in that concept, either. To say that redescription is a kind of abstraction, then, is to say that it *discards* information that is available in perception in the process of creating new symbols.

If this is right, though, note that redescription cannot produce any content that is *new* to the cognitive system. Abstraction discards information (see Mandler, 2007, p. 11), but the information must have been carried by some earlier state of the system in order for it to be discarded at a later stage (Dretske, 1981). Moreover, there is even a sense in which the process of redescription is a rather mechanical one. While attention is needed to activate a new conceptual primitive and make it available to thought, the process involved in concept activation seems to be a fairly inflexible, automatic one.

To see this, consider that the primitives can seemingly be 'matched' with their releasing inputs by a simple mechanical process. Perhaps this process might involve comparing perceptual inputs with the stored image schema template of, e.g., CONTAINER, and releasing the template when a sufficient match is found. This renders the problem of 'deciding' which primitive to release given a certain input easy to solve. In fact, the system would seem to have relatively little freedom about what concept it should acquire given the input. Such an account would have the

advantage of making sense of why certain templates are released under the influence of certain kinds of perceptual experiences: those are the ones that best match them in their gross spatial features. Moreover, this process of releasing primitives that match percepts also makes sense of the fact that templates carry a subset of the content encoded by the initial percepts. So this seems to be a fairly natural way (although perhaps not the only way) to understand what happens in Mandler's redescription process.

There are, then, two questions about this process. First, is it really a story about concept learning, as Mandler (2007, p. 16) suggests? This is important if it is to constitute an alternative to Fodor's triggering-based account. Second, is redescription so understood enough by itself to account for the range of concepts that we can entertain?

Take first the point about learning. As I've characterized it, redescription itself is a relatively inflexible process. When a stimulus having the right sort of characteristics is attended to, it is fed to the redescriber, which releases a more abstract representation that fits it (e.g., CLOSED-PATH). The redescriber has a certain number of symbols that it can release. These constitute its total available vocabulary. Once they are made available, all other concepts are formed by combining them in various ways (e.g., VEHICLE = MOVING + OUTDOOR THING + INANIMATE THING + CONTAINER). The process of combining pre-existing concepts is not itself a learning process. But a case can be made that redescription itself is not a learning process, either. Given the restricted range of stimuli that can release a conceptual primitive and the mechanical, template-fitting process that is involved in their release, it bears a closer resemblance to what Fodor calls, echoing the ethological literature, *triggering*. A triggering process is characterized by sensitivity to a small range of releasing inputs and relative insensitivity to other cognitive

factors (e.g., what else the creature thinks, what evidence it possesses, etc.). Redescription as I have understood it here seems to fit the bill.

If this is correct, there is no essential tension between Fodor's (1981) view of concept acquisition and Mandler's. While they disagree over the status of everyday lexicalized concepts like CAT and PLANE—with Fodor holding that these are unstructured atoms and Mandler holding that they are complex—on the issue of how we come by the conceptual primitives they don't really disagree at all. Fodor doesn't, of course, trouble to tell a psychological story about how triggering works psychologically, but Mandler's account can be seen as filling in this gap, at least for some concepts. Although it is substantially harder to see how it might be extended to account for the triggering of concepts like MORTGAGE and FATHER, this may just speak against the implausibility of thinking that these concepts are among the set of triggered concepts. Where Fodor and Mandler turn out to agree is that all concepts are either triggered or constructed from triggered concepts (plus or minus some sensorimotor information).

The final question to which I'll turn is whether a redescription mechanism such as this is really adequate to capture the richness of our conceptual repertoire. Here is one example. Birds and airplanes have wings, and it's likely that this fact is reflected in our BIRD and PLANE concepts. So HAS WINGS is a feature that belongs to both concepts. Now, is this feature *perceptual*? That is, is it encoded as a kind of image that is indexed by a pointer in the concept's structure? Or do we represent the fact of wingedness in a more abstract way than this when we *conceive* of birds and planes, as opposed to perceiving or imagining them? I'd suggest that the latter is the case. We can say—and think—perfectly univocally that both birds and planes have wings. No sense of zeugma attaches to such claims, although bird and plane wings are perceptually distinguishable. This suggests (although hardly proves) that at some level of

representation, the differences between bird wings and plane wings are ignored. This is what one would expect of a conceptual system that eliminates richly detailed perceptual information in order to provide an economical format for representing important facts about and properties of objects.

If this is what the conceptual system does, though, where do these conceptual features come from? Mandler's redescription mechanism can't be their source, since it is only sensitive to gross spatial and temporal properties of objects and events. And in any case it very plausibly has a restricted vocabulary; it can't be poised to coin new concepts for each new object part that needs to be generalized over (consider features like HAS EYES or HAS LEGS, or any others like them). Some *other* sort of concept learning device must be implicated in acquiring these features.

Here we can register a note of agreement with something Mandler says in the target article. She comments that features like BARKS, ascribed to dogs, are attached to concepts as part of a process of refining or honing a pre-existing concept (Mandler, 2007, p. 32). This seems exactly right. One needs some initial concept DOG in order to formulate the hypothesis DOGS BARK, and then to incorporate the feature BARKS into one's default way of thinking of dogs. But there is no reason to think that the feature BARKS is always going to be a perceptual feature. Rather than recovering the auditory image of barking whenever we think of dogs, we might just represent the property of being barkers as such. It's this ability that motivates the search for a different sort of concept acquisition mechanism than mere redescription.

Other sorts of example that pose a challenge for Mandler's account are those involving concepts that are learned through language. For an example of language guided concept learning, consider the sorts of experiments done by Waxman and colleagues (Waxman, 1990, 1999; Waxman & Gelman, 1986), in which children are presented with three perceptually dissimilar

objects and given a novel label that applies only to two of them (e.g., “These are blickets”). Despite the manifest perceptual dissimilarity of the named objects, children will go on to group them together just in virtue of having heard the new word used to label them. This suggests that they have changed the way in which they conceive of the objects. Now they conceive of at least two of them as being BLICKETS. While the children might have no notion of what a blicket is, hearing the word used might induce them to form a new concept to pick out that class of things that the adult authority is trying to demonstrate. Having formed such a new concept, they can go on to project hypotheses about what blickethood consists in, and to elaborate their BLICKET concept in ways that parallel the case of the DOG concept discussed above.

A general concept learning mechanism that explains how such language guided learning of new concepts might take place is described in Weiskopf (forthcoming-b). There it is argued that new concepts can be introduced when a learner has reason to think, given a certain class of evidence, that there exists in the world some class of objects, kind, or property that is of sufficient importance or recurrent interest that it will likely be a repeated target of thought and action. If such an assessment is made, the learner can coin a new concept that has the function of picking out that category, kind, or property. Importantly for our purposes here, the kinds of evidence that can lead to coining new concepts in this way can be comparatively indirect. The learner need not have had a direct perceptual confrontation with the object; it might be enough to have heard its name used in the right context, or to have observed some of its characteristic effects. This latter fact is especially important in explaining how we can acquire concepts of unobservable entities and forces (such as category essences or other causal powers). Given the right sorts of evidence that some such thing exists, we can coin symbols to keep track of and form hypotheses about these entities.

This description of how concept coining might work is necessarily sketchy. For the moment I want to draw attention to just one or two features of it. First, unlike the case of redescription, there is potentially no limit on the sorts of input that might prompt a person to form a new concept. Coining a concept is sensitive to the evidence that the learner possesses for the existence of the category in question. This might be simple perception of category instances; but it might be more complex. Consider, for instance, a case in which we posit the existence of a new kind of entity (e.g., black holes, x-rays) purely on the basis of the observation of its effects. This epistemic sensitivity distinguishes learning from triggering processes, including redescription.

Second, the content of concepts produced in this way is not always contained in the content of the input states themselves. Unlike the case of abstracting a concept for (say) a kind of path from a kind of detailed perception of an event, the concepts produced by coining are not necessarily abstractions from the content of the states that occasion them. Hearing the word ‘blicket’ leads a child to coin a concept for *whatever property the adult is picking out with that word*. But an episode of hearing the word ‘blicket’ does not count as a perceptual presentation of the property *blicket* to the child for abstraction. The word serves as a conventional indicator for the existence of a category. The child’s task is to infer that the content in question exists on the basis of that evidence. Again, this goes beyond what any comparatively simple abstraction mechanism can do. It also distinguishes this sort of mechanism from one on which the role of language is simply to *focus attention* on some property or aspect of the perceived scene. Here the property for which one is learning a concept need not itself be perceived or instantiated nearby for learning to take place.

While it seems likely that further mechanisms such as the ones I've sketched here will be necessary in order to fully explain the range of conceptualized thoughts that we can ultimately entertain, I should stress that I view them as adjuncts to the basic processes and mechanisms proposed by Mandler. Indeed, Mandler too acknowledges that language and analogical extension of spatial relations play a crucial role in concept acquisition beyond its earliest stages. Many concepts might be acquired through perceptual redescription, and it is worth noting that here we have the best empirical evidence for the actual mechanisms that take children from the realm of percepts to the realm of concepts (for a seeming instance of redescription in perception, see Schyns, Goldstone, & Thibaut, 1998). The mechanisms that I have proposed are more speculative, although gathering further evidence for their existence would obviously be desirable.

Putting all of the elements laid out here together results in a picture of conceptual development that is distinct from but related to Mandler's. Moreover, it is compatible with many of her proposals about how to delimit the domain of the conceptual (i.e., in primarily functional terms). On the view I've advanced, our first concepts and thoughts might be entirely composed of percepts. This is the greatest point of disagreement between Mandler's view and mine. These perceptually grounded first thoughts are then taken up by an array of complex concept triggering and learning mechanisms. Some of these, such as Perceptual Meaning Analysis, produce simplified image schemata that serve as our first global concepts. Others produce representations of the perceivable properties of objects, e.g., their parts and functions. Yet others produce concepts for categories that escape direct perceptual detection, e.g., those that are learned of mainly through language and causal reasoning about the world.

The result is that there is a gradual transition from the perceptual into the conceptual, but one that does not correspond to any simple stage-like transition and need not be uniform across ontological domains. And even adult concepts may retain some of their constitutive links with percepts, if it is more efficient for certain tasks (e.g., recognition, planning actions) to deploy fine-grained perceptual information rather than more abstract conceptual representations. While our first thoughts may co-opt percepts as their vehicles, the purposes of conceptual thought—making inferences that go beyond the readily perceivable, reasoning in ways that are insensitive to appearances, efficiently organizing and retrieving remembered information, and so on—are often best served by constructing non-perceptual representations that will most efficiently fit those purposes; and on this latter point, Mandler and I are in agreement.

References

Davidson, D. (1982/2001). Rational animals. In *Subjective, Intersubjective, Objective* (pp. 95-105). Oxford: Oxford University Press.

de Rosa, R. (2005). Prinz's problematic proxytypes. *Philosophical Quarterly, 55*, 594-606.

Dehaene, S. (1997). *The Number Sense*. Oxford: Oxford University Press.

Dretske, F. I. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Eimas, P. D., & Quinn, P. C. (1994). Studies on the formation of perceptually based basic-level categories in young infants. *Child Development, 65*, 903-917.

Fodor, J. (1981). The present status of the innateness controversy. In *Representations* (pp. 257-316). Cambridge, MA: MIT Press.

Gunther, Y. H. (Ed.). (2003). *Essays on Nonconceptual Content*. Cambridge: MIT Press.

Karmiloff-Smith, A. (1992). *Beyond Modularity*. Cambridge: MIT Press.

Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition, 104*, 19-46.

Mandler, J. M. (2004). *The Foundations of Mind*. Oxford: Oxford University Press.

Mandler, J. M. (2007). On the birth and growth of concepts. *Philosophical Psychology*.

Markman, A. B., & Stillwell, C. H. (2004). Concepts á la modal: An extended review of Prinz's *Furnishing the Mind*. *Philosophical Psychology, 17*, 391-401.

Nosofsky, R. M. (1992). Similarity scaling and cognitive process models. *Annual Review of Psychology, 43*, 25-53.

Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month old and 4-month old infants. *Perception, 22*, 463-475.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382-439.

Schyns, P. G., Goldstone, R. L., & Thibaut, J.-P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences, 21*, 1-54.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81*, 214-241.

Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology, 13*, 407-433.

Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies. *Cognitive Development, 5*, 123-150.

Waxman, S. R. (1999). The dubbing ceremony revisited: Object naming and categorization in infancy and early childhood. In D. L. Medin & S. Atran (Eds.), *Folkbiology* (pp. 233-284). Cambridge, MA: MIT Press.

Waxman, S. R., & Gelman, R. (1986). Preschoolers' use of superordinate relations in classification and language. *Cognitive Development, 1*, 139-156.

Weiskopf, D. A. (2007). Concept empiricism and the vehicles of thought. *Journal of Consciousness Studies, 14*, 156-183.

Weiskopf, D. A. (forthcoming-a). Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research*.

Weiskopf, D. A. (forthcoming-b). The origins of concepts. *Philosophical Studies*.