WHAT DO WE NEED CONCEPTS FOR?

RADU J. BOGDAN

If we are serious about concepts, we must begin by addressing two questions: What are concepts for, what is their job? And what means are available in an organism for concepts to do their job? One is a question of raison d'ê tre, the other of implementation. Here are a few reflections on these questions.

What is the business of concepts? To pick up relevant and useful properties of the environment. Why should they do that? To identify goal satisfying conditions and guide behavior toward them. And why should concepts identify and guide? Because organisms, whatever their complexity, have basic goals (replicate, survive, maintain appropriate energy levels by eating and resting, and so on) which they must satisfy. Categorizations help them satisfy their goals. At their simplest, categorizations are discrimination capabilities functionally aligned to the basic goals of the organisms: they distinguish aspects of events and objects which enable behaviors to track and bring about the goal satisfying conditions. 1 So far so trivial. How these discrimination capabilities are organized and what semantic and epistemic obligations they have is quite another and less trivial matter.

We need first some working notions and distinctions. Assuming that categorizations of all sorts guide action, let us have the notion (familiar in cognitive psychology and AI) of condition-production rule, or CP rule, characterize the simplest information-to-action sequence. A CP rule can be described in an 'IF (condition), THEN (production)' form. The condition is typically a data structure about some state of the world, the production typically an action. In simple and prelinguistic organisms the CP rules take the form of behavioral categories. An animal may, for example, have the following CP rule as its behavioral category: 'IF small, dark and smells good, THEN chase it'. What is important about behavioral categories is that the production part is inseparable from the condition part. Even when the same animal recategorizes the same prey in terms of 'IF only head is visible and has this shape, THEN chase it', the new condition cannot be added to the old and detached into an independent rule, cognitively closed and free of behavioral productions. To do that is to form a concept. Our animal has neither the resources nor the need to form concepts; it manages well (thank you) with behavioral categories.

A behavioral category can be primitive and yet package in its condition slot a number of properties. The category can be primitive in that it is not built out of other categories: its structure is either wired in or learned as a whole. Yet, as in our earlier illustration, the data structure in the condition slot may contain a good number of what we, but not the animal in question, regard as simple properties (small, dark, etc.). The animal categorizes the whole package. It does so because, presumably, the entire set of properties (and nothing less) is needed to guide behavior to its goals. Visual prototypes are primitive in this sense. In the long, evolutionary run, it is the success of behavior in satisfying basic goals which selects (by feedback) what sort of data structures end up in the condition of a category.

Not all behavioral categories of animal cognition need be primitive. Animals may learn to combine categories by either disjunctive listing of conditions (IF either small and dark, or IF head is thusly shaped, THEN do it) or by linking conditions in some order (IF such and such sounds, or such and such movements, THEN gather some more information, THEN (IF small and dark, THEN you know what next)). The latter is not an inference, just a conditional expansion of the data structure needed for the behavioral production. In either form, such categories remain essentially behavioral.

To put a bit more flesh on the skeletic notion of concept introduced earlier, think informally of concepts as CP rules which pick up invariant aggregations of properties of objects and events in the environment. Concepts are likely to emerge when the production in the CP rules must itself be cognitive. (The reason for that in a short while.) A concept can be represented as having the form 'IF object has property X, THEN it has property Y' (e.g., 'IF object barks, THEN it has four legs'). To have a concept requires both a role for it in the cognitive economy of the organism and the resources to form and efficiently access it. There is no point in having concepts if behavioral categories will do; and there is no way to have a concept if the resources are not there.

With these notions and distinctions in mind, let us ask ourselves: Do animals have concepts, in addition to behavioral categories? From a design standpoint, this amounts to asking whether behavioral categories suffice to guide behavior to goals. (I assume that, for thermodynamical reasons, nature is economical and therefore inclined to frown upon those which try to do with more expenditure of material and energy what they can do with less.) If behavioral categories are not enough, and if the behavior itself is observed to evidence cognitive versatility and sophistication, we must look for the cognitive resources (language, memory, whatever) needed to form, access and utilize the data structures operating as concepts.

Take the much discussed case of pigeon concept formation. We are told by animal psychologists that hungry pigeons can be trained to form the ("abstract") concept of redness through exposure to various color slides, with only red being rewarded with food. Pigeons are said to learn many other concepts in such conditions, most of them (about photographed fish, women dresses, and the like) being quite novel and unnatural for the species. My philosophical naivete recommends questioning the notion that we are witnessing genuine concept formation. All we seem to be shown is that pigeons are capable of new, systematic discriminations which have behavioral consequences. The question is how these discriminations are encoded and in what forms they are utilized. Both data and theory suggest behavioral categories of the form 'IF ....red, ..., THEN [IF peck, THEN food]'.

It is likely that the property of redness is packaged in the condition slot of the pigeon's behavioral category with other properties such as 'shape on the wall screen of a box' or the like. The experiment varies the images projected on the screen, not the screen itself or the wall or the box. We may therefore assume that more than redness remains invariantly encoded in the condition of the behavioral category. If the pigeon were to form a concept, that concept would rather be complex and distal (wall screen, etc.) rather than simple and sensory. Color concepts, per se, do not have great practical value for nonlinguistic beings. By themselves, they are not very likely to guide behavior to too many goals. When they do, it is typically as part of practically motivated packages of properties of middle sized objects and robust events.

Child concept formation seems to fit the same pattern. We are told that children master middle-size objects concepts of average abstractive power, such as those of table or dog, much before and much more easily than they master color concepts. It also appears that it takes words to really nail down the latter concepts. This makes practical sense. Given what children need to know and do, familiar middle sized objects must have priority in being represented, prototypically, in the conditions of their initial and practically motivated behavioral categories. When the time is ripe, the resources are there, and so is the reason for their utilization, the prototypes get their own conceptual rules and even names.

Since various properties, from colors to shapes to tablehood to red screens delivering food, are all out there waiting to be encoded, the fact that categories gravitate toward encoding some properties, and not others, some combinations of properties, and not others, at some levels of abstraction, and not others, must have an explanation in how categorial encodings guide an organism's behaviors toward goal satisfaction. This is why, I suggest, we must take a look at goals and behaviors before deciding on the semantic power and focus of an organism's cognition. When, then, does the guidance of behavior to goal require concepts, not just behavioral categories? When do concepts become useful as well as possible?

Concepts begin to matter, I think, when the organism needs (because has found it possible and beneficial) to rehearse cognitively, on some internal model of its immediate or more distant environment, what it should do next or what may happen next as a result of some event or how it should react as a result of some feedback. Rehearsing cognitively means expecting, anticipating or inferring some properties from others, considering subgoals and subproblems before tackling the big goals and big problems, and so forth. As cognitive CP rules, concepts can conveniently pack properties found to be constantly or invariantly associated, thus facilitating expectations, anticipations and inferences. The encoding and accessing of concepts is a matter of new and powerful resources. Language of course helps enormously by providing, through words, means of tracking concepts and, through grammar and logic, regimented and flexible means of deploying them. Language also regimets the public use of concepts, their semantic anchoring and epistemic values.

It is in the nature of a CP rule to systematically correlate the information tokened in the condition slot with its productive effects. Such a systematic correlation is not possible unless the information in the condition fits and serves the production. Animals do not live (eat, rest, procreate) in order to cognize but rather the other way around. Life is full of goals that the CP rules must serve. Although the order of a CP rule operation is, IF condition is satisfied, THEN production follows, the order of constraining (what selects what) must ultimately be the other way around. It is because concepts and behavioral categories enable an organism to guide its behavior to satisfy its basic and then derivative goals that certain data structures end up organized and encoded the way they are. Otherwise, feedback, reinforcement and learning would have no frame of reference, no standard of success and failure.

This is not a behaviorist thesis. It is a natural selection thesis about cognition. I am not saying that the information encoded by behavioral categories and concepts is not objectively determined by, and hence is about, features of the environment, and is not acquired by causal exposure to stimuli. (I am a proud card carrying naturalist.) I am however saying that whereas the environment and its stimuli are the ones which propose the information, it is the feedbacks from comparisons between goals and

behaviors which ultimately dispose by selecting the appropriate encodings for the information proposed. (A naturalist must bring natural selection inside cognition.) Although it is the manipulated environment in a Skinner box which determines what behavioral categories a pigeon must form to eat, it is the feedback from the successes and failures of the pigeon's behavior which retains, reinforces and adjusts the scope, organization and level of abstraction of the information structures needed to activate its productions. And, likewise, although it is the home environment which suggests to a child various types of objects in the form of cats, tables or redness, it is her cognitive and behavioral productions which stabilize the information structures in her categories and concepts around the useful and frequent prototypes, and not around data at a higher or lower level of abstraction.
This means that a lot of organisms discriminate and categorize only what they need to, in order to do their internal modeling and guide their behavior toward their goals. Behavioral categories and concepts are bound to reflect this functional subordination. An organism can therefore have behavioral categories or concepts about some property or object without knowing either the essentials of that property or object or indeed the standard conditions in which they exist. The pigeon need not know either the nature of redness or the normal circumstances in which redness is displayed in nature to encode redness (together with other properties) in the condition of its behavioral category.
This is by now a familiar result in philosophy, backed by psychological evidence. Yet, ironically, it is challenged by the very philosophical reflections on animal and human cognition which in recent years have contributed considerably to a deeper understanding of semantic cognition and conceptualization. I am thinking of the causal, information based and socially sensitive naturalizations of cognition and language. What I find ironic about the challenge is that it obscures the biological function of concepts by underrating the internal nature of organisms. The naturalism I am criticizing is a naturalism of externality, almost neobehaviorist.
It all began with Hilary Putnam imagining two worlds, W1 and W2, which differ in some deep compositional (say, chemical) respects yet send similar stimuli to two organisms, O1 and O2, inhabiting each of the worlds. The organisms are similarly constituted to the last cellular detail. O1 in W1 gets information from a structure whose chemical composition is XYZ while O2 in W2 from a structure whose composition is ABC. Since both XYZ and ABC send the same sensory input (both look, feel and taste like our W3, i.e. H2O, water), and O1 and O2 encode and use the information from XYZ and ABC in quite similar ways, they form the same concepts.

But they don't, we are told. There is a well known story about words, meanings and language based attributions, which I cannot go into here, according to which O1 and O2 end up with different meanings and beliefs. If O1 and O2 were social and linguistic beings, that could well be their story. It is not ours here. Our O1 and O2 are set up as simple animals which form behavioral categories or concepts but have no language, no meanings, and no fancy cognitive attitudes to attribute to each other. Why should O1 and O2 form different concepts? Because the concepts organisms form are determined by the kinds of information to which they are exposed. (Thus Fred Dretske.) As the laws of XYZ and ABC in W1 and W2 are different and determine different kinds of information being transmitted to and then tokened inside O1 and O2 by the physically indistinguishable inputs and internal states, the concepts O1 and O2 form about XYZ and ABC must be different. They are said to be about different (chemical) substances. But are they, really?

Let us assume that the concepts of XYZ and ABC are primitive (in the sense introduced earlier). "[T]he (primitive) concepts one acquires are limited by the kind of information available in the signals to which one develops a selective response, writes Dretske, and the concepts so acquired have their identity determined by this information." 2 True, one's primitive categories or concepts must reflect the kinds of information one has sensory access to, but they need not and had better not reflect all of that information. As Dretske has so well argued, concepts are filters: they select certain aspects conveyed by the input and filter out others. It is just possible that the primitive concepts O1 and O2 form in their worlds about the XYZ and ABC driven substances do (for practical reason) filter or leave out precisely the properties of their deep compositional structures, i.e. their XYZness and ABCness. But then O1's and O2's concepts cannot be about XYZness and ABCness; they must be about what is filtered in, the superficial WATERness (as we call it). It is about WATER that O1 and O2 have the same concepts. Differently said, their concepts do not reach as deep as XYZ and ABC, although the physical information from the latter is available in the respective inputs. But concepts are filters, not copies, of the input.

The possibility just contemplated turns into psychological reality if we recall that in the long run basic behavioral categories and concepts encode information about those properties of the world, as reflected in the input, which usefully guide their cognitive and ultimately behavioral productions to goals. If the set of superficial properties grouped around (what we call) WATER does this guidance job for O1 and O2 in their worlds, and neither O1 nor O2 has the means or needs to conceptually encode (i.e. filter in) information about XYZness or ABCness, respectively, then their concepts must be about WATERness only.

To sum up. Concepts and behavioral categories begin by being biological functions. It matters how they arise or mutate, but not that much. What matters more is that in the long run they are naturally selected by their effects, i.e. by the ability of behavior to produce goal satisfaction. Concepts and categories must be about the world to do their job. Yet the scope, texture and depth of their aboutness is not exclusively a matter of what the world has to offer; it is also a matter of how much of the world the organism can possibly as well as profitably accommodate. If this plain metapsychological truth meets so much naturalist resistance, it must be in part because many naturalists jump too easily to natural languages with their meanings and anchored references, to human societies with their semantic norms of word usage, expert division of cognitive labor and epistemic expectations of truth and knowledge. Behavioral categories and concepts are busy doing their job much before nature can afford these fancy developments. It helps to begin from ur and advance tout doucement. The semantic and epistemic normativity required by language and society should not be allowed to obscure what behavioral categories and concepts are initially and essentially for.

NOTES

1 I develop this teleosemantic position in Information and Semantic Cognition, MIND & LANGUAGE, 3, 2, 1988; and Guidance to Goal: The Roots of Teleosemantics, forthcoming.

2 Fred Dretske, Knowledge and the Flow of Information, The MIT Press, 1981, p. 227.