Towards an Epistemic Logic of Concepts

Tanja Hötte* and Thomas Müller

Universität Bonn, Institut für Philosophie, Lennéstr. 39, 53113 Bonn, Germany
thoette@uni-bonn.de, thomas.mueller@uni-bonn.de

Abstract. What does it take to possess a concept? Behaviour of various degrees of complexity is based on different levels of cognitive abilities. Concept possession ranges between mere stimulus-response schemes and fully developed propositional representations. Both biological and artificial systems can be described in terms of these levels of cognitive abilities, and thus we can meaningfully ask whether a given system *has concepts*. We regard that question not in terms of behavioural criteria, but from a formal point of view. We focus on the interrelation between a given objective structure of concepts and a subject's representation of that structure. The main question is how much of the structure of the objective side needs to be mirrored subjectively in order to grant possession of concepts. Our approach shows a strong parallel to epistemic logic. There, the objective side can be represented by an algebra of true propositions, and an epistemic subject can represent some of these propositions as what she believes to be true. As in propositional epistemic logic, in an epistemic logic of concepts the main issue is finding adequate closure conditions on the subjective set of representations. We argue that the appropriate closure conditions can be stated formally as closure under witnesses for two types of relationships among concepts: in order for a subject to possess a concept *c* she has to represent both a *sibling* and a *cousin* of *c*. We thus arrive at a first formally perspicuous candidate for a psychologically adequate epistemic logic of concepts.

1 Introduction

Biological systems show behaviour of various degrees of complexity, some of which is based on cognition. Some types of behaviour clearly do not presuppose cognition; among them are reflexes, but also mere stimulus-response behaviour such as that exhibited by Pavlov's dogs who slobber when they hear a bell ring. On the other hand, some types of human behaviour are clearly based on propositional representations and thus, on a high form of cognitive abilities; reading this article is among them. In between there are types of behaviour that presuppose conceptual representations without demanding full propositional cognition. Such behaviour is not confined to human beings: there are also some types of animal behaviour that fall within that range, e.g., the behaviour of parrots that Pepperberg describes in her book *The Alex Studies* [8]. Once these different levels of

* Research project "Wissen und Können", working group Stuhlmann-Laeisz, funded by the VolkswagenStiftung. Project website: www.wuk.uni-bonn.de

cognitive abilities have been recognised, it becomes possible to use them for describing not just biological, but also artificial systems. Thus we can meaningfully ask whether a given biological or artificial system *has concepts*.

When does a system have concepts? We will address this question from a formal point of view: we will presuppose that we know what a system *represents*, and we will try to describe minimal structural conditions that have to be met by the set of those representations in order to call them *conceptual*. Thus we leave the empirical question of how to derive representations from behaviour aside. The main link between behaviour and representation ascriptions will clearly have to be via various types of classificatory tasks. — We will also not consider the problem of concept acquisition, but aim at describing fully competent epistemic subjects. The concepts we focus on are perception-based ones like colour concepts or concepts of natural kinds, not theoretical concepts.

Our approach shows a strong parallel to *epistemic logic*. In that well-established branch of modal logic, one assumes that there is an objective side (that which can be represented) in the form of an algebra of *true propositions*, and an epistemic subject can represent some of these propositions as what she *believes to be true*. Whether the subject is granted knowledge normally depends on a number of side conditions, most of which are *closure conditions* for the set of propositions believed to be true. E.g., many systems of epistemic logic presuppose closure under logical consequence, so that a subject will only be granted knowledge of a proposition *p* if she also represents all logical consequences of *p* as believed-true. There is a long-standing debate about the adequacy of such closure conditions under the heading of “the problem of logical omniscience”: empirical subjects just cannot represent all logical consequences of any proposition as believed-true, since empirical subjects are finite, but there are infinitely many logical consequences of any given proposition (cf. [2], [3]).

In developing an epistemic logic of concepts, we have to face a similar problem. In parallel to propositional epistemic logic, we distinguish between an objective side of *concepts* and a subjective side of *representations*. Our main question can be phrased in terms of closure conditions: which conditions does a subjective set of representations have to meet in order to be truly conceptual? Just like for propositions, a trivial condition which would allow for representing only a single concept won’t do. On the other hand, imposing strict closure conditions runs the risk of ruling out too many empirical subjects for concept possession at all. Our task will be to find a good balance between these two extremes.

In Section 2, we will first describe the objective side of how concepts are structured. In Section 3, we then address the question of how much of that structure must be mirrored on the subjective side in order to grant subjective possession of concepts.

2 The Structure of Sets of Concepts

In this section we will regard relations between the elements of a set of concepts. But let us first say a word about the kind of such elements: we focus on perception

based concepts and we do not assume that the epistemic subject possesses a language. This restricts the concepts that we will be dealing with to rather simple ones. In particular, we do not consider second-order concepts, i.e., concepts that are themselves sets of concepts. Also, as we are interested in minimal conditions of concept possession, we presuppose only a thin notion of concept, not one that has rich internal structure. In this respect, our approach differs from, e.g., the detailed account of human concepts proposed by Kangassalo [6].

In line with this approach, we do not wish to presuppose too fine-grained a structure among the concepts that we are dealing with. We will focus exclusively on the *subconcept* relation and on relations definable in terms of it. The subconcept relation is the most basic relation on any set of concepts, and it will have to be considered in any case. There are further relations on a set of concepts that might be interesting. E.g., each concept could have a *complement*, or we could consider an *incompatibility* relation among concepts. However, in this paper we stick to positively defined, perception based concepts. With respect to these, negation appears to presuppose an additional layer of theory. Finally, having opted for first-order concepts only, the *element* relation among such a set of concepts is empty.

The fact that we do not require language capability as an essential ingredient of concept possession means that we must not assume that the epistemic subject can individuate concepts by name. Thus, the fact that languages can have synonymous terms referring to one and the same concept, does not pose a problem for our approach.

We will only consider finite structures in this paper. This will allow us to sidestep a number of technical issues, but there is also a philosophical basis for this restriction: we are interested in ascribing concept possession to actual (biological or artificial) epistemic subjects, i.e., to finite beings.

We will denote a set of concepts by $\mathcal{C}$. Single concepts will be denoted by lower case sans serif expressions, e.g., *a*, *b*, *cornflower* etc.

2.1 The Subconcept Relation

$\mathcal{C}$ is structured by the (reflexive) subconcept relation, usually dubbed *ISA* — e.g., as in *bird ISA animal*.

Now let us regard how the subconcept relation orders a set $\mathcal{C}$ of concepts. If $\mathcal{C}$ is a singleton there is not much to order, except that *ISA* is reflexive on the element. If $\mathcal{C}$ contains two concepts, we have the following possibilities: one concept can be a subconcept of the other or they can have nothing to do with each other. In case $\mathcal{C}$ contains three or more elements, two concepts can have a common sub- or superconcept, and they can be otherwise connected by the subconcept relation via a chain of intermediate concepts. It never happens that a concept is both a proper sub- and a proper superconcept of another concept. This suggests that, formally, the set of concepts $\mathcal{C}$ is at least partially ordered: *ISA* is reflexive, antisymmetric and transitive on $\mathcal{C}$.

Does the notion of a partial order really characterise *ISA*? The alternatives would be either to drop or to add formal requirements. Dropping requirements

appears implausible: ISA is certainly reflexive. It is also transitive: a subconcept of a subconcept of c is itself a subconcept of c . How about antisymmetry? Could there be concepts a and b such that a is a subconcept of b and b a subconcept of a while $a \neq b$? We frankly do not understand what that would mean. Thus, ISA is at least a partial order.

Maybe ISA satisfies additional formal requirements? Two intuitions pull in different directions at this point. On the one hand, in the tradition of formal concept analysis [5], it is customary to suppose that a set of concepts forms a lattice. This means that any two concepts have a common subconcept and a common superconcept, which is a lot of structure. On the other hand, in a philosophical tradition stemming from Aristotle, concepts can be thought of as ordered in a taxonomic tree. Formally, ISA on a taxonomic tree is a partial ordering that fulfills the additional formal requirement of backward linearity. In a taxonomic tree, concepts are distinguished by their genus proximum and their differentia specifica. Ordering concepts in such trees is intuitively appealing — think of zoology. However, such an order is unable to account for many of our intuitive uses of concepts, as taxonomy is always just with respect to *one* hierarchy of classification. We do however mix different classifying hierarchies. For example, peas are both vegetables and green. In a taxonomic tree, it follows that **vegetable** must be a subconcept of **green** or vice versa. Apparently, this is not the case. One can force concepts into a tree by adding multiple copies of nodes, as in Fig. 1. This is, however, both inelegant and yields an exponential overhead of new nodes.

Faced with the task of deciding between general partial orders, lattices or trees, we do not see overwhelming, univocal arguments in favour of one of the more specific

Fig. 1. From a partial order to a tree

structures. Thus we will only presuppose that $\langle \mathcal{C}, \text{ISA} \rangle$ is a partial order. This does not exclude the more specific cases, while keeping our theory general.

Note that we do not require **ISA** to connect all of $\mathcal{C}$. Let $\text{con}(\mathbf{a})$ be the set of concepts connected to a concept $\mathbf{a}$ via the **ISA** relation (i.e., the reflexive transitive closure of **ISA** and its converse). There could be some $\mathbf{b} \in \mathcal{C}$ such that $\mathbf{b} \notin \text{con}(\mathbf{a})$, in which case $\text{con}(\mathbf{a}) \cap \text{con}(\mathbf{b}) = \emptyset$. We do not exclude the possibility of $\mathcal{C}$ consisting of several of such **ISA**-clusters, isolated from each other. In fact, we do not think it likely that this is the case for an objective ordering of concepts, but in this paper we will maintain full generality. In Sect. 3.3 we will argue that it is useful to allow a *subjective* ordering of concepts to contain disconnected elements.

In the finite case (as well as in ‘friendly’ infinite cases), the **ISA** relation gives rise to a *covering relation* that relates just closest subconcepts in the partial order. Let $\mathbf{a}, \mathbf{b} \in \mathcal{C}$. Then $\mathbf{a}$ is covered by $\mathbf{b}$, written $\mathbf{a} \prec \mathbf{b}$, iff (i) $\mathbf{a} \text{ ISA } \mathbf{b}$, (ii) $\mathbf{a} \neq \mathbf{b}$, and (iii) on the assumption that $\mathbf{a} \text{ ISA } \mathbf{c}$ and $\mathbf{c} \text{ ISA } \mathbf{b}$ we have $\mathbf{c} = \mathbf{a}$ or $\mathbf{c} = \mathbf{b}$.

2.2 Siblings and Cousins

Let us introduce two useful notions with the help of $\prec$: the sibling relation and the cousin relation.

The intuition for the *sibling* relation comes from trees: in a tree, a sibling of a node is a node with the same mother. Here, we generalise this notion to partial orders, excluding the reflexive case.

$$\text{sib}(\mathbf{a}, \mathbf{b}) \Leftrightarrow_{df} \mathbf{a} \neq \mathbf{b} \wedge \exists \mathbf{c} (\mathbf{a} \prec \mathbf{c} \wedge \mathbf{b} \prec \mathbf{c}).$$

This relation is irreflexive and symmetric, but not transitive (in view of $\mathbf{a} \neq \mathbf{b}$). Sharing a covering node ($\exists \mathbf{c} (\mathbf{a} \prec \mathbf{c} \wedge \mathbf{b} \prec \mathbf{c})$) is not necessarily transitive (though on trees it is). Staying in the family metaphor, one could say that partial orders allow for half-siblings. And the half-sister of Eve’s half-brother need not be her sister at all (see Fig. 2 (a)).

The second relation that we will introduce is the *cousin* relation. Two concepts are cousins if they have immediate superconcepts that are siblings. Cousins are more independent than siblings, but still connected via the **ISA** relation.

$$\text{cousin}(\mathbf{a}, \mathbf{b}) \Leftrightarrow_{df} \mathbf{a} \neq \mathbf{b} \wedge \exists \mathbf{c}, \mathbf{d} (\mathbf{a} \prec \mathbf{c} \wedge \mathbf{b} \prec \mathbf{d} \wedge \text{sib}(\mathbf{c}, \mathbf{d}))$$

Fig. 2. (a) patchwork families, (b) $\mathbf{a}$ is a candidate for being one’s own cousin

Again, reflexivity it excluded. We have to put in this condition explicitly because otherwise in the case of the partial order of Fig. 2 (b), *a* would be its own cousin. Note that *cousin* is a symmetric relation.

3 Possession of a Concept

After we have made some observations on the form of sets of concepts in general and introduced relations on them, we now turn to the subjective side of concept possession. We thus consider the epistemic view that an animal or artificial system has of the objective set of concepts. The question is how much of that set has to be grasped in order to ascribe concept possession sensibly. In the following we will first discuss the extreme cases of requiring no structure at all and of demanding so much that even humans cannot be ascribed concept possession any more. Then we will present an intermediate position, based on results of cognitive science as well as on solutions to the problem of omniscience (see Sect. 1).

To state the question formally, let i be an epistemic subject. Now let us introduce a set $\mathcal{P}_i \subseteq \mathcal{C}$ for each i . $\mathcal{P}_i$ shall denote the set of concepts i possesses. Now we can state the question as follows: given $a \in \mathcal{P}_i$, what other concepts must $\mathcal{P}_i$ contain? And what structure does $\mathcal{P}_i$ need to have?

Relations on $\mathcal{P}_i$

Each $\mathcal{P}_i$ might just be a small subset of $\mathcal{C}$. Still, $\mathcal{P}_i$ can have an internal structure. In the following, we presuppose that the epistemic subject does not err in the concepts nor in their relations (cf. Sect. 1). A similar presupposition is operative in propositional epistemic logic. In general, we would not ascribe possession of some cognitive content to a subject who mixes up the inferential relations in which that content stands. Thus, for $a, b \in \mathcal{P}_i$ we have

$$a \text{ ISA}_i b \text{ iff } a \text{ ISA } b,$$

where ISA_i denotes the ISA relation defined on $\mathcal{P}_i$ instead of on $\mathcal{C}$. From the epistemic point of view this is reasonable, since if an epistemic subject possesses both a concept and one of its superconcepts, he has to treat them as concept and superconcept. If he treated them differently, e.g., as siblings, our intuition would not grant him possession of both concepts. Formally, $\mathcal{P}_i$ is thus a subordering of $\mathcal{C}$.

This perfect matching need not hold for the covering relation. $\mathcal{P}_i$ is likely to be less fine-grained than $\mathcal{C}$. In the examples of Fig. 3, all nodes belong to $\mathcal{C}$, but just the circled ones belong to $\mathcal{P}_i$. Let $\prec_i$ denote the covering relation of ISA_i . We have $\text{plum} \prec_i \text{fruit}$, but not $\text{plum} \prec \text{fruit}$. Similarly with all the other nodes at the bottom level. This does not conflict with our use of concepts. What we use extensively is the ISA relation, but we would always be reluctant to label a concept as a direct superconcept absolutely — there is always the possibility that we do not know enough of the area, or that we have missed some in-between concept.

Regarding *sib* and *cousin*, which depend on the covering relation, even more can change if we determine them via $\mathcal{P}_i$. Let sib_i and $cousin_i$ denote the relations defined on $\mathcal{P}_i$ instead of on $\mathcal{C}$, i.e., in terms of $\prec_i$ instead of $\prec$, via the respective definitions from Sect. 2.2. Now regard the example of Fig. 3(a) once more. We have $sib(\text{plum}, \text{apricot})$ and $cousin(\text{plum}, \text{grape})$, but no $cousin_i$ relation holds in $\mathcal{P}_i$, and we have derived relations like $sib_i(\text{plum}, \text{grape})$ etc. For this special example, sib_i is even larger than *sib*, while $cousin_i$ gets trivialised on $\mathcal{P}_i$. Using a different example, for $\mathcal{P}_j = \mathcal{P}_i \setminus \{\text{fruit}\}$, $sib_j = \emptyset$, we can enlarge $cousin_i$ as well (see Fig. 3(b)). These examples illustrate the fact that sib_i and $cousin_i$ generally are only weakly related to *sib* and *cousin*. If we use sibling or cousin relationships in specifying closure conditions, it therefore seems advisable to employ the objective relation, i.e., the restrictions of *sib* and *cousin* to $\mathcal{P}_i$. This means that these relations cannot in general be defined internally (in terms of $\mathcal{P}_i$ and $\prec_i$): $sib(a, b)$ can hold in $\mathcal{P}_i$ even if the superconcept *c* witnessing the sibling relation is not in $\mathcal{P}_i$. The same holds for *cousin*. We therefore adopt *sib* and *cousin* as primitive relations on $\mathcal{P}_i$: $sib(a, b)$ holds in $\mathcal{P}_i$ if and only if $a, b \in \mathcal{P}_i$ and $sib(a, b)$ holds in $\mathcal{C}$.

This way, concepts not connected via ISA_i can still stand in the sibling or cousin relation in $\mathcal{P}_i$.

Fig. 3. Examples of two $\mathcal{P}_i$ s smaller than $\mathcal{C}$: (a) more siblings (b) more cousins

3.1 Atomism

The absolutely minimal requirement for possession of a concept is just possession of that very concept. The approach to concept possession that specifies no additional requirement is called atomism. That view appears to be endorsed by Fodor [4, p. 121]. Atomism yields maximal independence of concepts. In particular, there are no conditions on the complexity of the structure of the set of concepts an epistemic subject can possess. Since every concept can be possessed on its own, the test for concept possession cannot go far beyond testing the capability to classify objects with respect to the concepts. This opens up the possibility of ascribing concepts to small children and animals. However, it turns out that atomism is untenable.

Seen from the point of cognitive science, the problem with atomism lies exactly in its generous openness. Since there is no complexity of the structure

of concepts required, there is no distinction between concept possession and stimulus-response behaviour. As discussed in Sect. 1, this is not adequate.

3.2 Holism and Other Closure Conditions

Now we address the issues on the other side of the scale: demanding as much structure as there is.

Holism. According to a stricly holistic approach, an epistemic subject needs to possess the whole set of concepts in order to be granted possession of even a single concept. This is too harsh a requirement even for humans. Even if the objective set should be finite, it is probably very large; too large for the capacity of a single mortal being. Recall from Sect. 2.1 that $con(a)$ is the set of concepts connected to a concept a via the ISA relation. Then the holism condition reads as follows:

$$\forall a, b[(a \in \mathcal{P}_i \wedge b \in con(a)) \rightarrow b \in \mathcal{P}_i].$$

Note that this condition does not necessarily require possession of all concepts from $\mathcal{C}$. It just requires possession of all concepts that are connected to the concept in question via ISA. Should there be two or more isolated structures of concepts (cf. Sect. 2.1) all but one could be left out completely.

Closure under Superconcepts. Closure under superconcepts is a bit less demanding.

$$\forall a, b[(a \in \mathcal{P}_i \wedge a \text{ ISA } b) \rightarrow b \in \mathcal{P}_i].$$

In this case, one does not need to possess the whole set of concepts, but just everything along the order relation starting with the concept in question. This recursive abstraction and generalisation is very natural for humans (a dog is a mammal, a mammal is an animal, an animal is a living being, ...), but it is not clear why it should be essential for concept possession. In particular, it will be next to impossible to find an animal capable of this level of abstraction. On the other hand, there are clever animals to whom we want to ascribe concept possession, e.g., the parrot Alex (see Sect. 1). So this condition demands too much.

Closure under Subconcepts. Closure under subconcepts is formally similar, but has a different motivation.

$$\forall a, b[(a \in \mathcal{P}_i \wedge b \text{ ISA } a) \rightarrow b \in \mathcal{P}_i].$$

Again, the order relation is followed, but this time downwards in the set of concepts. It is not a process of abstraction but of specialisation that is carried out. Once more, this is natural for humans (some dogs are large dogs, some large dogs are St. Bernhards, some St. Bernhards are trained for rescue in the mountains, ...), but might rule out concept possession for animals.

Closure under Sub- and Superconcepts. For the sake of completeness, let us regard the conjunction of the closures of the above, closure under both superconcepts and subconcepts.

$$\forall a, b [(a \in \mathcal{P}_i \wedge (a \text{ ISA } b \vee b \text{ ISA } a)) \rightarrow b \in \mathcal{P}_i].$$

This closure is subject to both objections raised above. In addition, as this condition has to hold for every concept in $\mathcal{P}_i$, this type of closure yields the whole set of concepts connected to a : we are back at holism.

Closure under Siblings and Cousins. A closure condition that requires less abstraction is closure under sibling and cousin concepts.

$$\forall a, b [(a \in \mathcal{P}_i \wedge (sib(a, b) \vee cousin(a, b))) \rightarrow b \in \mathcal{P}_i].$$

Note that we talk here about the *sib* and *cousin* relations on the objective set of concepts (cf. the discussion before Sect. 3.1). But again, do we have to possess all sibling concepts in order to possess one single concept? In the Indo-European language family, we have eleven basic colour concepts. There are other languages that have considerably less basic colour concepts (down to just two concepts, for black and white). If closure under siblings were a necessary condition, people with just the concepts of black and white could not even have the same colour concepts for these two colours as we have. But they do, as Berlin and Kay have shown in [1]. So this condition is not adequate either.

Closure under Witnesses. Another approach towards more realistic closure conditions consists in weakening the universal quantification to an existence quantification. Instead of requiring all concepts of a certain kind (e.g., all superconcepts) one can just require (at least) one concept of that kind. This one concept witnesses that the condition does not run empty (which might be another danger for all-quantified conditions). We will call this kind of closure *closure under witnesses*. For example, one can demand closure under the following condition:

$$\forall a [(a \in \mathcal{P}_i \wedge \exists b (a \text{ ISA } b)) \rightarrow \exists b (a \text{ ISA } b \wedge b \in \mathcal{P}_i)],$$

which means that if i possesses a concept a and there is a superconcept of a at all, i possesses at least one superconcept of a as well. This condition is met if $\mathcal{P}_i$ contains just one chain of concepts connected through ISA. Analogously, one could formulate a definition using subconcepts, going down in the order instead of up.

The problem of this approach lies again in our goal to leave a possibility for animals to possess concepts. As we discussed already in the case of closure under all superconcepts, humans do not appear to have a problem with iterated generalisation. Neither do they have a problem with iterated specialisation. But there is no reason why these faculties should be essential for concept possession. There is even a reason against that. From psychology we know that there is

a level of perceptually based prototype concepts that are learned easiest by children. Learning how to generalise or to specialise such concepts comes much later in the development. Still our intuition tells us that children possess these concepts as soon as they show the appropriate classification behaviour.

So far we have seen that requiring no structure of $\mathcal{P}_i$ does not give enough complexity for concept possession. On the other hand, stipulating closure conditions that are universally quantified amounts to requiring more cognitive abilities than are necessary for mere concept possession. So the adequate condition must lie somewhere in between. We have discussed weakening the closure under superconcepts and subconcepts to existential quantification. This still runs into the problem of requiring the capability of repeated generalisation or specification for concept possession. In the next section we will regard a weakened condition on closure under siblings and cousins, which appears to be psychologically adequate.

3.3 Closure Under Witnesses for Siblings and Cousins

In their paper [7], Newen and Bartels propose an account of what it takes to possess a concept based on psychology and animal studies. They argue that there are certain levels of complexity of behaviour that can be related to levels of cognitive ability. Concepts are located on an intermediate level. Concept possession requires flexible behaviour more complex than fixed stimulus-response schemata, but less complex than, e.g., planning based on propositional knowledge.

Based on their assessment of the complexity of behaviour typically thought to require conceptual representations, Newen and Bartels propose two structural requirements for concept possession. In order to possess a concept, an epistemic subject should first possess another concept of the same dimension of classification, e.g., not just *red*, but *blue*, too. We formalise this in terms of the (objective) sibling relation

$$\forall a[a \in \mathcal{P}_i \rightarrow \exists b(sib(a, b) \wedge b \in \mathcal{P}_i)].$$

The other requirement is to possess another concept on the same level of abstraction that is not, however, a sibling. The intuition behind this goes back to classifying objects: an epistemic subject shall be capable of sorting objects under different dimensions of classifications, with respect to different aspects. This, of course, need not work for an arbitrary object — just for those that fall under more than one concept that the epistemic subject possesses. Again, this requirement is symmetric. So far, the cousin relation is our best candidate for a formalisation.

$$\forall a[a \in \mathcal{P}_i \rightarrow \exists b(cousin(a, b) \wedge b \in \mathcal{P}_i)].$$

This condition does not exactly match the above intuition. What we have is that in a tree, if the condition is fulfilled, we can ascribe concept possession to the epistemic subject. So the condition is sufficient, but it is not necessary: there can be nodes on the same level that are related more distantly. On general partial orders, the condition is not even sufficient. We encounter two issues: a concept can be both a cousin and a sibling of another concept, with respect to different superconcepts. One would have to exclude that this one concept is

used to fulfill both conditions. The second issue is with the intuition of levels of abstraction: it can happen that two concepts on different levels are cousins. In this case it is not possible to ascribe levels unambiguously to $\mathcal{C}$.

Some partial orders admit the definition of levels, while others do not. Formally, we can capture this distinction as follows: If $\mathbf{a} \text{ ISA } \mathbf{b}$, then $\mathbf{c}_1, \dots, \mathbf{c}_n \in \mathcal{C}$ form a *direct connection* from $\mathbf{a}$ to $\mathbf{b}$ if $\mathbf{a} = \mathbf{c}_1 \prec \dots \prec \mathbf{c}_n = \mathbf{b}$. We call n the *length* of the connection. Let $dc(\mathbf{a}, \mathbf{b})$ be the set of all direct connections from $\mathbf{a}$ to $\mathbf{b}$. Then we can say that $\mathcal{C}$ *respects levels* if for any elements $\mathbf{a}$ and $\mathbf{b}$, all direct connections are of the same length, i.e., if the following holds:

$$\forall \mathbf{a}, \mathbf{b} \in \mathcal{C} \forall \langle \mathbf{c}_1, \dots, \mathbf{c}_n \rangle, \langle \mathbf{d}_1, \dots, \mathbf{d}_m \rangle \in dc(\mathbf{a}, \mathbf{b}) [\mathbf{a} \text{ ISA } \mathbf{b} \rightarrow n = m].$$

That is, in a partial order that respects levels, any two nodes have a fixed distance. In trees, this is obvious: if $\mathbf{a} \text{ ISA } \mathbf{b}$, then $dc(\mathbf{a}, \mathbf{b})$ has exactly one element.

If $\mathcal{C}$ respects levels, we can define a level function: Let $\mathcal{L} : \mathcal{C} \rightarrow \mathbb{Z}$ be a function, assigning every concept in $\mathcal{C}$ an integer indicating its level of abstraction. If $\mathbf{a}$ is a perceptually based prototype concept, $\mathcal{L}(\mathbf{a}) = 0$ shall hold. In addition, we require $\mathbf{a} \prec \mathbf{b} \Leftrightarrow \mathcal{L}(\mathbf{a}) + 1 = \mathcal{L}(\mathbf{b})$.

With this terminology, we can state the following condition for possessing concepts on the same level of abstraction.

$$\forall \mathbf{a} [\mathbf{a} \in \mathcal{P}_i \rightarrow \exists \mathbf{b} (\mathcal{L}(\mathbf{a}) = \mathcal{L}(\mathbf{b}) \wedge \mathbf{a} \neq \mathbf{b} \wedge \neg sib(\mathbf{a}, \mathbf{b}) \wedge \mathbf{b} \in \mathcal{P}_i)].$$

That is, to every concept possessed by the epistemic subject i there is a different concept possessed by i that is on the same level of abstraction, but not in the same dimension of classification (i.e., not a sibling).

Further work will be required to weigh carefully the pros and cons of demanding the additional structure embodied in $\mathcal{L}$. So far, we have the condition on the existence of siblings and we search for another symmetric condition requiring the existence of one more element of $\mathcal{P}_i$. By closure under these conditions, we get a set of at least four concepts.

4 Conclusion

When does a biological or artificial system have concepts? In this paper, we addressed this question from a formal point of view. We focused on the interrelation between a given objective structure of concepts, which we argued is a partial ordering, and a subject's representation of that structure, which we argued should be a subordering. The main question was how much of the structure of the objective side needs to be mirrored subjectively in order to grant possession of concepts. We thus discussed a number of closure conditions for subjective sets of representations. Based on results from cognitive science, we argued that the appropriate closure condition that strikes a balance between atomism and holism is closure under witnesses for siblings and cousins.

As we argued in Sect. 3.3, we believe that closure under witnesses for siblings is both formally and psychologically adequate. Closure under witnesses for cousins

is a formally precise condition that psychologically at least points in the right direction. In order to further develop the epistemic logic of concepts we hope to benefit from discussions with both cognitive scientists and scientists working in the field of AI.

Acknowledgements

We would like to thank Albert Newen and our anonymous referees for helpful comments. T.H. acknowledges support by the VolkswagenStiftung.

References

1. Berlin, B. and Kay, P.: *Basic Colour Terms: Their Universality and Evolution*, University of California Press, Berkeley 1969
2. Cursiefen, S.: *Formen des logischen Allwissens. Eine problemorientierte Darstellung modallogischer Systeme*, PhD thesis, Bonn 2006
3. Fagin, R., Halpern, J., Moses, Y., and Vardi, M.: *Reasoning about Knowledge*, MIT Press, 1995
4. Fodor, J. A.: *Concepts: Where Cognitive Science Went Wrong*, Oxford University Press, Oxford 1998
5. Ganter, B. and Wille, R.: *Formal Concept Analysis: Mathematical Foundations*, Springer, Heidelberg 1999
6. Kangassalo, H., The concept of concept. In Ohsuga, S., Kangassalo, H., Jaakkola, H., Hori, K., and Yonezaki, N. (eds.), *Information Modelling and Knowledge Bases III: Foundations, Theory and Applications*, IOS Press, Amsterdam 1992, pp. 17–58
7. Newen, A. and Bartels, A.: Conditions of Possessing Concepts, submitted
8. Pepperberg, I.: *The Alex Studies: Cognition and Communicative Abilities of Grey Parrots*, Harvard, Cambridge (Mass.) ⁴2002