Berthouze, L., Prince, C. G., Littman, M., Kozima, H., and Balkenius, C. (2007).
Proceedings of the Seventh International Conference on Epigenetic Robotics: Modeling
Cognitive Development in Robotic Systems. *Lund University Cognitive Studies, 135*.

Modelling the thing-in-itself – a philosophically motivated approach to cognitive robotics

Matthias J. Schlemmer **Markus Vincze** **Bernard Favre-Bulle**
Vienna University of Technology
Automation and Control Institute
Gusshausstrasse 27-29/E376
1040 Vienna, Austria

Abstract

In this paper, we will try to outline a few major questions that the design of artificial cognitive systems should tackle. The argumentation will be led from a philosophical point of view, reviewing some of the latest epistemological positions, namely from the 20^{th} century. The notion of the thing-in-itself, i.e. the substantial core of an object, will be central in the discussion of abstraction in cognitive systems. A second important point is intentionality (in the philosophical sense): Doing something is always guided or biased by a task that is currently pursued. Last but not least the old question of symbol anchoring is to be posed. A preliminary system design for abstracting knowledge is shown that forms a possible implementation for tackling these aspects for the domain of robotic vision.

1. Why start from philosophy?

It has become obvious in the past years that the study of cognitive systems, in which epigenetic robotics takes an important part, not only profits from but even requires an interdisciplinary approach. Philosophy has an outstanding role in the history of sciences, often termed as „mother of all sciences“. Its unique approach can provide researchers with new perspectives and simultaneously retaining the big picture. The usual suspects of disciplines involved in the study of cognitive systems, namely biology/neuroscience, developmental sciences, psychology and artificial intelligence, have evolved from philosophy when sciences started to split into more detailed directions. Nevertheless, their roots are philosophical.

Second, at least since the ground-breaking article „Epistemology Naturalized“ by Willard van Orman Quine (van Orman Quine, 1969), natural sciences and philosophy started to move together. Quine rejected the notion of a „prima philosophia“ (first philosophy) in the sense of a philosophy being in the elitist atrium of sciences, but rather stood up for philosophy being seen as a science like any other, simply arguing in a more abstract sense (Schneider, 1998).

Last and maybe the most exciting point is, that philosophy has the tendency to pose the hard and uncomfortable questions. There are some issues in cognition that are spooking around but are tackled explicitly only in philosophy (e.g. the discussion of free will). As we will see later, exactly the problem of what makes a cognitive system *cognitive* is – sometimes with other terms – one of the oldest philosophical topics. An approach to cognitive systems from a philosophical epistemology seems to make an important yet arguable implicit assumption: cognitive systems are intended to be human-like or human-inspired. As we will show later on, this paper argues for a very restricted human-relation, however.

Summarised, there are issues that philosophy might help to address. With this paper we try to start with the philosophical notion of the thing-in-itself in the context of abstraction in cognitive systems. Second, we tackle intentionality as the philosophical notion for the system's bias due to the task currently pursued. Finally, semantics are to be looked at – a topic well-known to artificial intelligence. To put it in a nutshell, it is our firm belief that the look at epistemological positions has the ability to bring up new approaches and perspectives to an understanding of cognition in engineering as well.

2. On modelling cognitive systems

The goal of this paper is to tackle some of the hard questions in the interesting research field of building artificial cognitive system. In (Ziemke, 2001), the main approach to cognitive systems is from *radical constructivism*, whose basic principles are: knowledge is not passively received but actively built up, the tendency of the system is guided towards viability and – probably most important – cognition serves the subject's organization of the experiential world, not the discovery of an ontological reality. We are approaching from the domain of home robots, in which the services that such a system should offer to the user take a prominent role. The implicit assumption mentioned in the previous section of judging artificial cognitive systems by the best known – though not fully understood – cognitive system *human* calls out for another tricky question: Shouldn't we model cognitive systems – i.e. robots – like humans because we know they do in fact work? Probably this is the wrong question as we should rather first define what we want such a system to be capable of. The bottom line is that we want a robot to perform a task we want it to. In our belief, this is nevertheless compatible with the main demands of radical constructivism, but possibly gives them slightly an other turn – at least for the „subsystem" vision that does by no means cover the huge area of knowledge representation on the whole. We think that the robot's process of actively building up knowledge about the environmental world, should be driven by the human's needs. This means, we should provide the robot with some basal knowledge in an „innate" manner so that it is able to act in a way that is *viable* – and it is us who judge the viability of the robot by its ability to fulfil the tasks we want it to fulfil. This automatically implies, of course, that there is no need for some absolute ontological reality.

Without willing to continue the never-ending story of trying to define the notion *cognitive* – as this term has been used inflationary in the last few years to avoid difficult words such as *intelligent* or *conscious* – we have to point out that our position is that it has at least something to do with knowing what and why we are doing something. As we will see in the following, this implies important aspects of which we want to tackle three: *intentionality*, *detection of an instance* from known object abstractions and *symbol anchoring*. The latter is a necessary precondition for an artificial cognitive system as this provides the linkage between the perceived objects and their semantics. The considerations on the notion of consciousness that we will give in the following should – as already stated – outline a possible approach to designing artificial cognitive systems. We argue that *consciousness* for a robot shall be defined in the sense we will depict now, and not in the everyday (human's) sense as this is not central for an artificial cognitive system. Hence, we do *not* argue for a robot to necessarily be self-aware or to have a feeling of subjectivity. These aspects are usually used for distinguishing inferior bio-organisms from higher developed mammals. Instead, we are arguing that dealing with our environment in a meaningful way simply needs some sort of „taking part". We are, therefore, only using the term consciousness in order to motivate our approach to the necessity of *intentionality*.

Associating the term *conscious* with the term *cognitive* reminds of the epistemological *theory of objects* (in German: *Gegenstandstheorie*). The theory of objects was an important ancestor of Husserl's phenomenology and for its founder Alexius Meinong, consciousness is *always* levelled at something that could be an external or internal object (Schneider, 1998). External objects can be thought of as real, existing entities (i.e. objects in the everyday sense); with internal objects he refers to inner procedures. Therefore, the real existence of objects is not a necessary criterion for something that the consciousness focuses upon. Of course, the emphasis of conscious as being *conscious towards* something (which already was Franz Brentano's principal assumption), can be found similarly in a second aspect that is frequently investigated in the context of cognitive systems: *attention*. It must be stated that *object* is meant in the broadest – not necessarily materialistic – sense here. The notion of an object from an engineering point of view is thus totally different from the philosophical one. Furthermore, it is to be underlined that *intentionality* is meant exactly in this sense: as the guidance of our consciousness, of which attention is a necessary but not sufficient part. It is at any rate not meant in a teleological sense, i.e. in the sense of having the explicit intention to achieve something.

One of Meinong's central claims was that those objects that we are always conscious towards, i.e. that our intentionality is guided towards, can either be *real* or *ideal*, the latter not being concrete instances of things (i.e. objects in the materialistic sense) but rather anything else – which comprises impossible objects or abstract concepts of objects. This already introduces a hierarchy of objects, from simple real entities of things to higher-level abstracted, non-existent ideal objects (in the philosophical sense).

This leads us to the second big issue that cognitive systems have to tackle: *abstraction*. The central question here is: How are these objects that we are always conscious towards represented in the system? It is unclear, how humans represent concrete items or create and represent generalised concepts or fantastic illusions. Therefore, it is unlikely that humans can serve as a blueprint for synthesising a representational entity, but they might be a good motivation. Psychologically, there are two main approaches to human's object representation, the exemplary model – storing objects in terms of typical examples – and the conceptual one – storing *concepts of objects* rather than examples. In (Feldman, 2006), the latter is being shown as the more plausible. The conceptual approach is also the one taken in our work. In the next Section we introduce an example that shall be a thread through the remainder of this article.

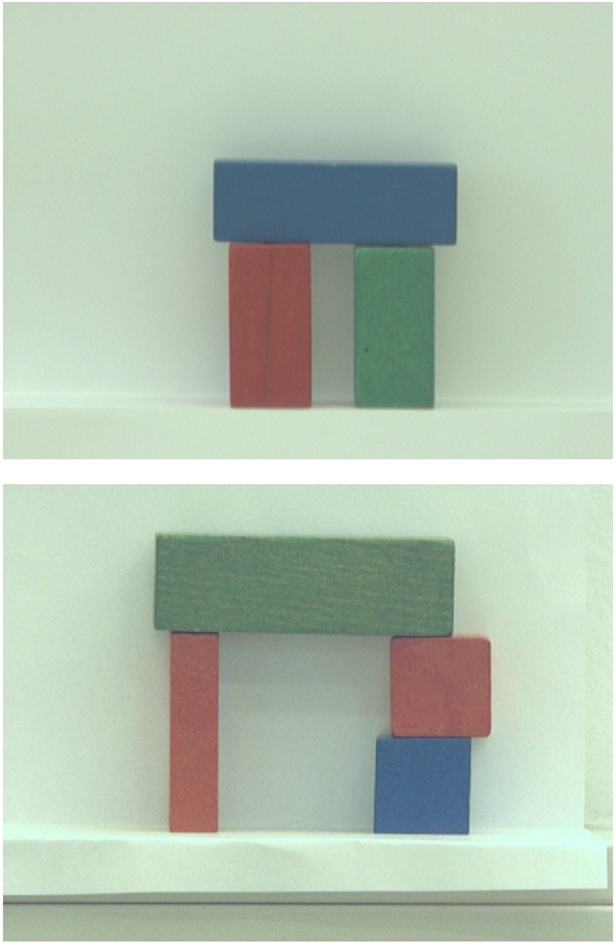

Figure 1: A simple blockworld example: different arches.

Finally, doing something means acting in a world which further implies a rough understanding of the objects in this world at the least. Being guided by a task, which is always given through intentionality, we need to have semantically grounding of those objects. Furthermore, the ability to abstract objects and classify them, is inherently connected to an understanding of what these objects are for. A famous example is the search for a chair, which can be considered as rather searching for a function, namely a sit-able surface. This function-based approach requires semantics while classifying and acting in one's environment. To link perceived objects with semantics is therefore a necessary precondition for a cognitive system.

2.1 A simple example of human abstract object knowledge

Consider an arch, like the one in Fig. 1. No matter whether to look at the upper or the lower image, we are able to understand in both cases that these blocks form something like an arch.

Without having seen these specific arches, we are able to classify them. So the point is: How are we able to do that? Obviously, we do not only store all arches that we have ever seen in our life – most likely, none of the readers have ever seen *these* arches. It rather seems that we have some kind of abstract representation what the constituting substrate – the substantial core – is that an arch is made of. This is what we will call *object concept* later on and what is known in philosophy with the notion of the thing-in-itself (see Sec. 3.). In the case of humans, this abstracted object concept is very likely to be learned – along with its semantic content.

2.2 Hard issues

Some of the hard problems to be tackled when designing a cognitive system have already been outlined. Especially the consciousness towards objects – real or ideal, the representation of generalised concepts and the long standing problem of symbol anchoring/grounding. The litmus test for a cognitive system in our view might be whether it has the following capabilities: abstraction, generalisation and prediction. This is, of course, not an exhaustive list. For example, to use abstraction in a meaningful way, the system must be able to derive from these abstractions when encountered with a specific instance. We deliberately don't use the term *recognition* here because of its connotation in connection with computer vision.

Most of the just mentioned capabilities are in need for a thorough notion of objects, and more precisely, the ability to recognise what makes an object an object.

3. The Thing-in-itself

Let's get to the question of abstraction, specifically how object concepts are formed and – even more important – retrieved when encountered in a specific situation. The discussion about what constitutes an object is practically as old as philosophy itself. Plato established *ideas* that we have seen in heaven and that we can recall. Those ideas are given the individual in an a priori way. Aristotle introduced categories, starting from ουσια, the substance (e.g. being human) and then going further on to quantity, quality, relation, place, time, position, state, action and affection[1]. Additionally, by defining that e.g. all humans have the same „human-ness", he introduced the notion of accidentals, meaning that beside the same essence, other properties form a specific instance (e.g. Socrates). Therefore, an arch is an arch because of his taking part in the ουσια of arch-ness. Additional accidentals might be that it is red, high,

[1]The first four categories are the most important. The last four differ in *Categories* (Aristotle, 2006) from the ones in *Metaphysics* (Aristotle, 2007).

dirty or whatever.

At a later date, Immanuel Kant did not only point to the fact that the perception of the world around us is biased by the subject that is perceiving, but he was also the first one to widely introduce the notion of the *thing-in-itself*. By that he described those actual existing categories of the world that we cannot directly perceive; and as the thing-in-itself is inherently not ascertainable, we cannot even know its existence as well. We can only *guess* its existence by perceiving the things-for-us, which are the actual phenomena. Our (subjective) reality is consequently made up of appearances that come from but are not the same as the things-in-themselves. The latter only affect our senses (Kant, 1787). However, in this work we are heading for implementing some sort of thing-in-itself in order to provide certain knowledge about the world, therefore the Kantian position is not fully satisfactory.

Nicolai Hartmann's *critical realism* purges the Kantian position even more from metaphysical touches (Hartmann, 1921). Knowledge (in the cognition sense) is nothing that is created in the human's mind, but rather something that does exist no matter whether it is perceived or not. This was also Kant's opinion, however, the main difference is that Hartmann believes that this knowledge is captured actively and we are in fact able to perceive important traits of the thing-in-itself and not only the thing-for-us. For him, this is the reason for us being able to talk about things, that we did not actually perceive. In this argumentation, Hartmann tries to strike a balance between realism and idealism. This is in line with constructivist approaches, such as (Varela et al., 1993), where the traditional dichotomy of realism and idealism is broken up by neither focusing only on the subject (idealism) nor on the object (realism), but instead on the *relation* between subject and object. Hartmann, as well, emphasises on subject and object being elements of an ontological interrelation that has acausal character.

This difference in Kant's and Hartmann's notions of the thing-in-itself seems to be an unnecessary subtleness, but the latter is a lot more useful for us having a synthetical modelling of the things-in-themselves in mind. Kant thought of the thing-in-itself as something never to be ascertainable, Hartmann, however, presumes that importaint traits are in fact perceivable. Therefore, from a pragmatical point of view, we are pleading for the latter notion as, of course, we would not only like to describe those entities but also be able to *use* this representation later on from the system's perspective. The notion *concept of an object* that we would like to suggest in this paper is therefore philosophically the thing-in-itself in Hartmann's sense.

The idea now is to use such a description of an object concept in a cognitive system. This means storing things-in-themselves and relating them to current perceptions. This way, a qualitative description of known object concepts in the world can be separated from the quantitative localisation of current appearances. As vision is the richest of the human senses due to its capability of delivering distal information of spatial arrangements of objects in the world, this sensor is a good choice for a lot of robotic applications. Especially in home robots, navigation and activity is in need of spatial information of objects. A first approach of such a combination of vision and object concept knowledge will be shown later on.

4. On the way of synthesising

The philosophical positions might help us so far in explaining how human knowledge – from a rather high-level view – could be built up, but they are clearly lacking the details for giving us a recipe of how to find or even come close to an object's constitutive properties. The phenomenological suggestion would be to apply the technique of *eidetic reduction*, meaning that we are varying all properties of an object in order to detect those that are really constituting.

Here is the point. We are arguing against rebuilding humans as artificial cognitive systems with respect to the *way* things are abstracted, represented and retrieved. What we *do* argue for, is providing those systems with a similar *kind* of representation in a high level sense. Therefore, we are proposing that some sort of thing-in-itself is provided to a cognitive system by the designer as we want the system to perform in a way that it is able to support and – to a certain extent – *understand* us when interacting in *our* world. This further implies, that we do not want it to learn abstract object concepts in its own way – which unquestionably could be interesting for some disciplines as well – in order to form some particular representations itself. The aim is that such a system is able to work and bind objects to semantics in *our* understanding. Otherwise we are in fact on a rebuilding trail, which is – at least for the domain of home robots – of minor importance.

In evolutionary epistemology, e.g. (Vollmer, 1975), it is argued that the kind of knowledge that we are able to build up throughout our lifetime is strongly biased by the specifics of our senses that have evolved over thousands of years. The goal was always to provide us with the necessary means to survive. This reminds of the phenomenological theory of the *life-world* (Husserl, 1986), which states that all cognition ultimately refers back to our surrounding environment[2]. This is something that should hold for robots as well. It should survive, meaning, it should

[2]The critical realism mentioned in Sec. 3., as well, is of the opinion that all knowledge is bound to the real being.

perform in a way that we users would judge to be meaningful. There is one advantage that those system have in comparison to our evolution: We are able to provide those systems with those sensors and capabilities that they are able to circumvent thousands of years of evolution. And the theory of the lifeworld is just another word for the paradigm that robotics turned to some years ago: embodiment and situatedness. In (Agre and Horswill, 1997), we find a definition for the domain of artificial intelligence: „A lifeworld, then, is not just a physical environment, but the patterned ways in which a physical environment is functionally meaningful within some activity.“. The point we are arguing here is that the ecological niche of a home robot is right at the user's side, and this should be accounted for. With other words, the lifeworld of the robot is to a large extent overlapping to the human's. This further implies that functional meaningfulness will be judged by humans.

Let's get to the domain of computer vision for robotics, for which the following system is built. A classical image processing technique is re-cognition of objects or object parts that were learned, e.g. (Agarwal et al., 2004). Another frequently used approach is using models, i.e. shape-cues, to re-cognise objects (Lowe, 1987, Solina and Bajcsy, 1990). Difficulties in detecting these fixed model cues often lead to failure. However, all of those approaches – as we just outlined – are probably not the recognition technique performed by humans. We can detect a novel arch also – and obviously only detecting learned objects is not very cognitive, already because there are no semantics linked with that. This is the old-fashioned warehouse metaphor for knowledge representation. Humans, instead, seem to be able to store a generalised model of objects that allow for redetection – not in a purely geometrical sense, but with a lot of other properties as well. Yet another technique is perceptual grouping (Zillich, 2007): finding low-level features and organising them into Gestalts, eventually ending at a proto-object level, e.g. grouped lines forming a closed contour. Here, our approach steps in: out of these grouped proto-objects, concepts of objects are to be detected, easing the bridging to the semantic object level.

Our proposal is therefore to implement the things-in-themselves in a knowledge base, i.e. an ontology. By that we are clearly able to define the ecological niche of the system, its – phenomenologically speaking – lifeworld, which is in the example proposed a simple blockworld scenario. In other words, by providing the system with information about its lifeworld, we are modelling its *Umwelt* in von Uexkuell's sense (von Uexkuell, 1909). We are having in mind to extend this approach for the domain of home robots in the future. Moreover, a vision part retrieves closed contours from a camera and tries to link them with the high-level description in the ontology. As in robotics, we are always having a *task*, thus the intentionality of the system is modelled inherently and is therefore able to reduce the search space. Finally, the symbol anchoring problem is circumvented by being able to store some functionalities in the ontology as well. This way, the concept of an object whose instance is found in the camera frame, immediately also helps us in retrieving its function. Sir Karl Popper stated that in humans the consciousness is marshalling the empirical material and is itself this way sense-constituting (Popper, 1972). This is what we want to achieve by separating object concept descriptions – the high-level knowledge in our „consciousness“ with all their semantic comments – from the empirical data processing – image processing in the system proposed.

5. First implementation

For the implementation of the ideas just proposed, it is necessary to define what to store in the ontological knowledge repository as the qualitative high-level description of the object concepts that we derived earlier from the philosophical notion of the thing-in-itself. A famous example from the domain of computer vision is the one by Stark and Bowyer (Stark and Bowyer, 1991) that is an approach to detect only necessary properties of objects. They thin out every aspect of a chair to the functional description of having a sit-able surface along with stable support. Our first approach is motivated by having a robot in mind that tries to gain knowledge about its environment, starting from a simple task such as detecting arches like those of Fig. 1. *Every* arch shall be detected this way, not only the one that simply consists of two supporting and one supported object. The starting point is the modelling of spatial relations between closed contours that are found by a perceptual grouping tool (Zillich, 2007).

The main difference to usual feature-based recognition lies in the usage of the ontology as top-down guidance for the search of objects. This way, we are trying to combine the bottom-up data that is delivered by the perceptual grouping tool with the top-down information of ontological world knowledge. This is biologically plausible, as there are a lot of pathways and feedback-loops from higher to lower levels of the neuronal brain, cp. e.g. (Hawkins, 2004). Second, due to the applied abstraction technique, we are trying to be as general as possible so to detect as much variations of the objects as possible. The approach is therefore comparable with part-based recognition, but with the additional benefit of separating facts from current image processing. This allows for an explicit storage of the world's properties without the need to consider vision expertise. Theoretically, any sensor modality should be able to use the ontology.

5.1 Ontology

The word ontology actually stems from philosophy as well, meaning „[...] the study of what there is.“ (Stanford University, 2007). In this approach, we are modelling *what there is* for the system in an ontology in the computer science sense. There, an ontology is „an explicit specification of a conceptualization“ as the popular definition of (Gruber, 2007) states. The emphasis lies on *explicit*, stating that in an ontology we model the knowledge of our domain (i.e. the world our robot as well as the user act in) clearly in natural language, making all assumptions visible.

This approach enables a step-wise abstraction of the objects and parts in the world around us. Starting with the identification of some part (e.g., due to attention), we might come up with a bigger object to which this part belongs and that finally might again be part of an even larger object. This should also accelerate the whole computation – dealing with grouped objects instead of all subparts highly reduces the search space.

For a start, our ontology is very slim and holds the classes *Stack* and *Arch* as known object concepts (composed of polygons) as well as the definition of the following spatial relations: left_of, right_of, on_top_of and below_of. The latter two additionally have subproperties defining whether the object on top is stably supported by the subjacent object (in a naïve physics way). Of course, this is very basal and needs to be refined in future work. A stack, for example, is currently defined as being made up of two polygons one resting stably on the other. Due to the stepwise abstraction technique, we are, of course, also able to find a column made up of more parts as each step connects two polygons to one, new larger object.

5.2 Combining Vision and Ontology

As vision pre-processing, the perceptual grouping framework *vs2*, developed by (Zillich, 2007), is used. With the closed contours extracted this way, we have two possibilities depending on the task that the system is given: Either a specific object concept is tried to be detected or all relations in the image are sought to comply with the concepts known.

Given the task of finding a specific concept, the arrangement of found polygons is now checked to comply with the abstract qualitative description that is stored in the ontology. This is done by checking every two polygons whether they fulfill a specific relation to each other, e.g. one being on top of the other in a stable manner. In the case of a naive definition of an arch, three of those restrictions, involving three polygons, have to be met: two polygons being adjacent and supporting a third, but being not adjacent to each other. Technically, this is a simple neighbouring check, using geometric translations.

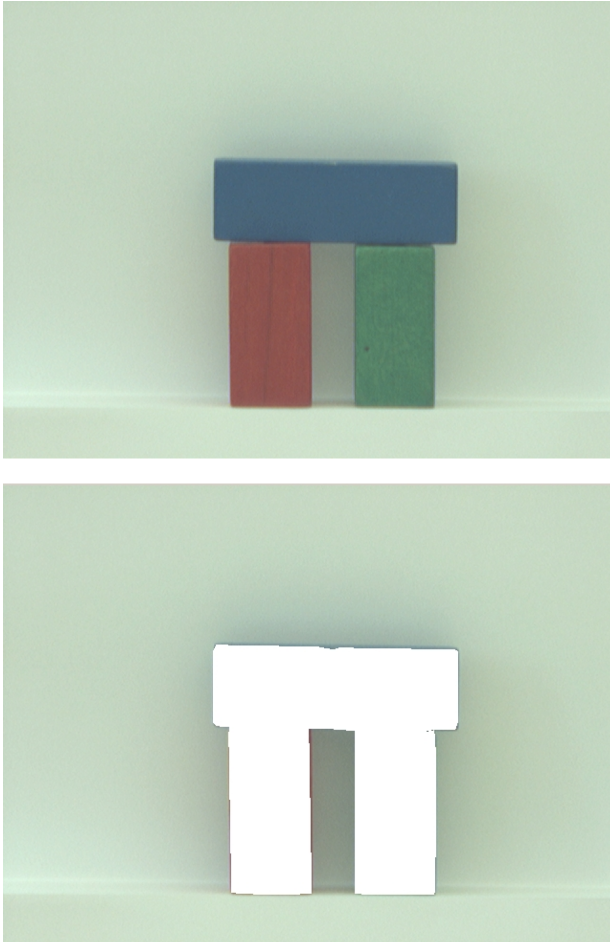

Figure 2: A simple arch (on top) and the result after checking against the ontology (below).

This might lead to a finding like in Fig. 2, which shows the first example arch along with the result after checking the relations between the single blocks in the ontology.

In the context of abstraction, the second example is a bit more interesting: What we are arguing is that different specific arrangements of objects (e.g. having a different number of contributing polygons) shall also be assigned to an instance of the concept of an arch. Fig. 3 on top shows the second example arch, where the right pillar is itself composed of two objects. In other words, we have a different appearance, yet still the same concept.

A direct query for location of an arch without previous simpler abstraction steps, leads to a finding like the second image of Fig. 3. If we are first locating the stack instead (see Fig. 4 on top), we are able to

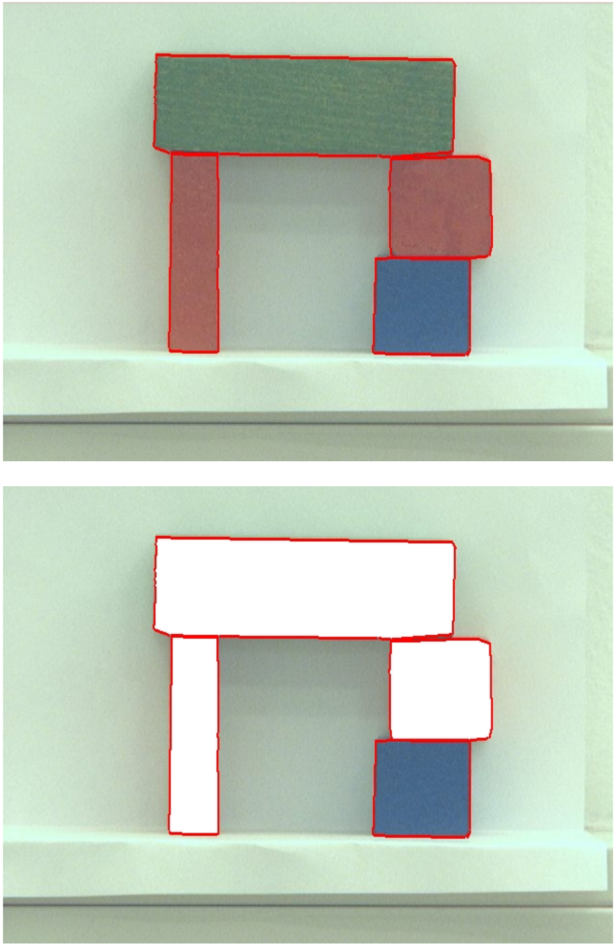

Figure 3: An arch consisting of differently structured columns. Input image (on top) and direct query for arch.

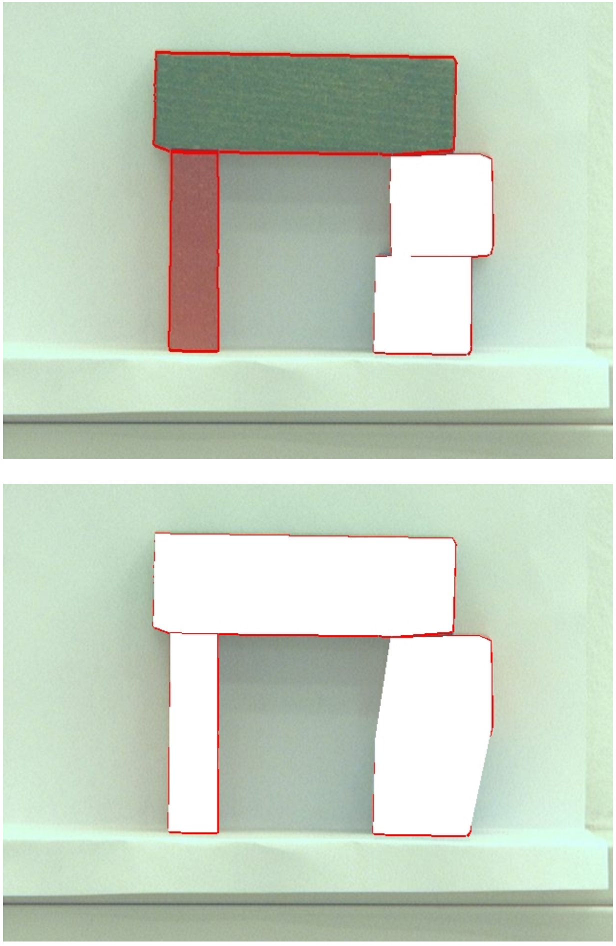

Figure 4: Result of a query on Fig. 3 for stack and the result for a query for arch *after* querying for stack.

use this – abstracted – knowledge for retrieving the whole structure: an arch (see Fig. 4 below). The location of a stack works hereby in the same manner as the one for the arch: First, the restrictions defining a stack are checked for all polygons in the image (via geometric translations), then complying polygons are grouped together to form one new object, namely a stack – for the later retrieval of arch, this stack is seen as a simple polygon again.

As mentioned, modelling some sort of thing-in-itself in an ontology and using this knowledge for reasoning about the actual perceptual input, we are able to detect concepts of objects that could not be detected by classical object recognition approaches. The modelling of the essential properties is crucial yet tricky. If we were able to identify those properties for the most important objects in our domain, we are thus able to define those *qualitative relations* between objects in the ontology and later on to retrieve a *quantitative description* of their existence in the current view without having a representation of the object concept (stack, arch) in the vision part. With this combination we can discover object classes such as stack or arch independent of how many pieces contribute.

6. Conclusion and further work

In this paper we tried to approach artificial cognitive systems from the view of philosophical epistemology. Key concepts of epistemology, namely the notion of an object and the thing-in-itself along with their representation on the one hand and intentionality of consciousness on the other hand were shown to tackle some of the questions needed to be answered in cognitive system engineering.

The implementation of those philosophical theories is as follows. We introduced the thing-in-itself as a good means to have an abstraction for recognition of object concepts. Additionally, we tried to roll up the field from the side of always having a task – which is in the broadest sense the focus of evolutionary epistemology and the intentionality of consciousness. Finally, we gave a first qualitative abstraction method which is able to assign different appearances to the same object concept.

Though we argued that the situatedness of the robot (it's ecological niche) could and for a start should be defined by its operators, humans, it would be an interesting field of research to investigate how the thing-in-itself could be derived from appearance-based prototypes. This would lead to an even more human-like approach (if whished), as from the physicalist view, everything we know is somehow built up and stored in our neurons, in other words: learned. Quoting Quine once again: „all inculcation of meaning of words must rest ultimately on sensory evidence" (van Orman Quine, 1969). This would, of course, afford something like Husserl's *eidetic reduction* to be functionally implemented in the system. And besides, a representation of the semantics would have to be really built up by the system itself. This, however, is future work.

Acknowledgements

The work described in this article has been funded by the European Commission's Sixth Framework Programme under contract no. 029427 as part of the Specific Targeted Research Project XPERO („Robotic Learning by Experimentation") as well as by the Austrian Science Foundation under the grant #S9101 („Cognitive Vision").

References

Agarwal, S., Awan, A., and Roth, D. (2004). Learning to Detect Objects in Images via a Sparse, Part-Based Representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(11):1475–1490.

Agre, P. and Horswill, I. (1997). Lifeworld analysis. *Journal of Artificial Intelligence Research*, 6:111–145.

Aristotle (2006). *The Categories*. Dodo Press.

Aristotle (2007). *Aristotle's Metaphysics*. Read Books.

Feldman, J. (2006). An algebra of human concept learning. *Journal of Mathematical Psychology*, 50(4):339–368.

Gruber, T. (2007). What is an Ontology? online at *http://www-ksl.stanford.edu/kst/what-is-an-ontology.html*.

Hartmann, N. (1921). *Grundzüge einer Metaphysik der Erkenntnis (Fundamentals of metaphysics of knowledge)*. Vereinigung wissenschaftlicher Verleger, Berlin.

Hawkins, J. (2004). *On Intelligence*. Times Books. ISBN 0-8050-7456-2.

Husserl, E. (1986). *Phänomenologie der Lebenswelt (Phenomenology of the lifeworld)*. Reclam.

Kant, I. (1787). *Critique of Pure Reason*. Cambridge University Press, 1999.

Lowe, D. (1987). Three-Dimensional Object Recognition from Single Two-Dimensional Images. *Artificial Intelligence*, 31(3):355–395.

Popper, K. R. (1972). *Objective Knowledge*. Oxford University Press.

Schneider, N. (1998). *Erkenntnistheorie im 20. Jahrhundert – Klassische Positionen (Epistemology in the* 20^{th} *century)*. Reclam.

Solina, F. and Bajcsy, R. (1990). Recovery of Parametric Models from Range Images: The Case for Superquadrics with Global Deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(2):131–147.

Stanford University, M. R. L. (2007). Logic and Ontology. Entry of the *Stanford Encyclopedia of Philosophy*. online at: http://plato.stanford.edu.

Stark, L. and Bowyer, K. (1991). Achieving Generalized Object Recognition through Reasoning about Association of Function to Structure. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 13, pages 1097–1104.

van Orman Quine, W. (1969). *Epistemology Naturalized. In: Ontological Relativity and Other Essays*. Columbia University Press.

Varela, F. J., Thompson, E., and Rosch, E. (1993). *The Embodied Mind: Cognitive Science and Human Experience*. The MIT Press. ISBN 0-2627-2021-3.

Vollmer, G. (1975). *Evolutionäre Erkenntnistheorie (Evolutionary Epistemology)*. Hirzel, 2002.

von Uexkuell, J. B. (1909). *Umwelt und Innenwelt der Tiere*. Julius Springer Verlag.

Ziemke, T. (2001). The Construction of 'Reality' in the Robot. *Foundations of science*, 6(1-3):163–233.

Zillich, M. (2007). *Making Sense of Images: Parameter-Free Perceptual Grouping*. PhD thesis, Vienna University of Technology.