

# Using Analogy to Transfer Manipulation Skills

**Frank Guerin (1) [f.guerin@abdn.ac.uk](mailto:f.guerin@abdn.ac.uk)**

Paulo Abelha Ferreira (1) [p.abelha@abdn.ac.uk](mailto:p.abelha@abdn.ac.uk)

Bipin Indurkha (2) [bipin@agh.edu.pl](mailto:bipin@agh.edu.pl)

(1) University of Aberdeen,

King's College, AB24 3UE Aberdeen, Scotland.

(2) AGH University,

al. A. Mickiewicza 30, 30-059 Krakow, Poland.

## Abstract

We are interested in the manipulation skills required by future service robots performing everyday tasks such as preparing food and cleaning in a typical home environment. Such robots must have a robust set of skills that can be applied in the unpredictable and varying circumstances that arise in everyday life. To succeed in such a setting, a service robot must have a strong ability to transfer old skills to new varied settings. We are inspired by the strong transfer ability demonstrated by infants and toddlers on simple manipulation activities, and we are motivated to try and replicate these abilities in an artificial system. We treat this as a problem of making analogies, and describe a theoretical framework which could account for it. We sketch the ideas of a computational model for implementing the required analogical reasoning.

## Introduction

We are motivated by the problems facing future service robots, in particular robots that can prepare food, cook, clean, and maintain a house. We focus on the manipulation skills required. Such robots must have a robust set of skills that they can apply even when circumstances vary, for example the tools used in food preparation may vary, the materials to be used may vary, the target tasks (e.g. what dishes to prepare) may vary. To succeed in such a setting a service robot must have a strong ability to transfer old skills to new varied settings.

We take inspiration from the way in which even very young humans (infants and toddlers) can effectively transfer old skills to a wide variety of new settings. While toddlers may make quite a mess if asked to make a cake for example, they always approach the task with an array of skills which they have used in different settings, and can re-purpose. For example, if asked to transfer the contents of a bowl to a tin, a toddler who has never done such a task will find ways to transport the mixture by pouring, or scooping with a spoon, and to scrape remaining bits from the bowl. The toddler will draw on previous personal experience with similar tasks such as pouring water, scraping food from his/her own feeding containers, shaking or scraping sticky food from a spoon, previous play with dough etc.

This is in stark contrast to existing robots, who when faced with a novel setting sometimes can find no effective match between their old skills and the new task demands. Robots tend to be programmed quite specifically for a target task, and not for generic transferable skills. Infants in contrast spend a great deal of time acquiring generic transferable skills. In fact it is impossible to get an infant to do a task like transferring a cake mix from a bowl to a tin; they will instead play with the mixture, bang a spoon up and down in it, experiment with it with their fingers. It is almost as if the infant knows that it is pointless to jump to advanced tasks before the necessary substrate skills are built. The consequence of this difference is described quite succinctly by Rod Grupen (Fagard et al. 2013):

“In general, skills and abilities in infants and robots are still acquired in quite different ways. Infants build layer upon layer of support skills by exploration that seems to be independent of any other purpose than to acquire comprehensive mastery of increasingly sophisticated relationships to the world. No task is required. The state of the art in robotics, however, typically starts with a target task and is reduced into pieces that are described algorithmically. Typically, a designer anticipates all the events and intermediate states and therefore, the robot is unsupported by the same breadth of contingencies that the infant spends all of its time constructing during the sensorimotor stage of development.”

In other words, when faced with a new setting which is broadly similar to a previous example, but has some new aspect which causes difficulty (new “contingency”), the infant is likely to be able to relate the new aspect to one of the skills it has practised in the past, and to borrow that past skill and find a way (with some groping) to apply that knowledge in the new situation. We believe that dealing with contingencies in this way requires analogy; a low level kind of sensorimotor analogy which can find similarities between motor programs, visual scenes, and properties of objects/materials. In this paper we sketch our ideas about how the required analogical mechanism might be implemented in an artificial system.

## Some Examples of Analogy in Infants and Toddlers

We now present some examples of sensorimotor analogy in infants or young toddlers for the purposes of (i) convincing the reader that analogy is widely used at this age, and (ii) to give some example analogies which we would like an artificial system to be able to replicate.

In general, to reinforce the claim that infants/toddlers make frequent use of sensorimotor analogy we note that: (i) Infants seem to acquire object manipulation skills which are very robust to changes in situations and materials, and to acquire a basic commonsense knowledge of physics. This has been determined by testing infants (e.g. (Uzgiris and Hunt 1975; Brown 1990)). (ii) If one looks at the experiences the infants have had prior to the testing they are almost never exposed to the same objects or exactly the same situations as in the tests. The infants seem to have a strong ability to “transfer” from similar situations. It seems to be impossible to account for these observations without attributing them to analogical mechanisms of thought.

The problem of retrieving an out-of-reach toy across a table using a rake tool is exceedingly difficult for infants up until about 18 months of age (Rat-Fischer, O’Regan, and Fagard 2012). One of the significant difficulties involves the need to get the “business-end” of the rake behind the toy. The typical environment of an infant in the home does not usually involve situations where a rake-tool would be used to retrieve an object. The typical care-giver does not train an infant on this task. The fact that infants without training tend to succeed at 18 months suggests that they had other experiences in the preceding weeks/months which helped. It is possible that an infant can make an analogy between this manipulation and the manipulations infants typically practice in the 12-18 month period when they are learning to self-feed with a spoon. This often involves getting the spoon behind a piece of food in order to scoop it up. Recent research has tested infants on a task of retrieving an out of reach object where the tool to be used is a large spoon, and the object is a piece of cake or a toy. It was found that half of the infants who failed in a standard rake-retrieval task succeeded when the tool was a spoon and the target was a piece of cake [spoon + cake], and a third of infants succeeded when the tool was a spoon and the target was a regular toy [spoon + toy]. Furthermore only the infants who succeeded with [spoon + toy] were able to generalize to [rake + toy] afterwards (100% of them), while none of the infants who only succeeded with [spoon + cake] were able to generalize to [spoon + toy] or [rake + toy].<sup>1</sup>

Brown (Brown 1990) tested toddlers for transfer ability in retrieval tasks with a variety of stick-like tools, some of which had a hook or rake at the head. Tools varied in length, rigidity, colour pattern, and type of head. Children showed a good knowledge of the properties required of an effective tool for the task. Piaget (Piaget 1936) also observed similar transfer in this task “So it is that, in Observation 160, Jacqueline, knowing how to use a stick without hesitation, also uses

<sup>1</sup>Personal communication from Lauriane Rat-Fischer about yet-to-be-published results.

a book and a banana as instruments.” Brown makes a strong case for the ability to transfer being very much domain specific, and related to the child’s understanding of causality in the particular task. This means that even young infants could reasonably be expected to be capable of analogical reasoning provided that it is in a domain where they have a good deal of experience, e.g. in pulling cloths. Brown makes the point that by later ages children may have more domain-general abilities, where more traditional analogy tasks can be solved.

Piaget describes (OI(Piaget 1936), p. 216-219, obs. 122-123) how an infant learns, in the period from 7½ to 8 months, how to remove an obstacle blocking him from grabbing a visible object. He then describes how the child subsequently uses this discovery to find a solution to the problem of attempting to grasp something when the hand is already full; the child discards the object held (OI, p. 221, obs. 125 “Finally at 0;7 (28) Laurent finds the solution ...”). Note that Piaget states that this second behaviour “apparently derives from” the previous.). Later (at 9, and then 9½ months) the infant is faced with a situation where a cushion covers a desired object, and the infant executes a means action to remove the cushion by analogy with the situation of removing an obstacle, displacing, pushing away barrier (OI, p. 222, obs. 126). It is noteworthy that the infant had failed to remove a covering object prior to these examples. In the two examples above, the child seems to exploit an analogy with a previous situation in order to recognise the appropriate behaviour to be executed.

As with Piaget, many examples of analogy in infancy are anecdotal. There is very little research on analogy in infancy, so that in some cases anecdotes are the best evidence we have. We have one further anecdote from our own observations, which concerns a two-and-a-half-year-old using a long light plastic stick to retrieve a bunch of keys across a table. The child had no prior training for this task. The child did a successful retrieval showing skill that demonstrated foresight of the problems involved. Firstly, the child put downward pressure on the stick. If there were no downward pressure the light stick would have passed over the keys. The child did not discover the need for downward pressure through groping on the task, but applied it from the outset. Secondly, when the stick was successfully dragging the keys, but slipped slightly to one side the child readjusted to put the stick more in the middle. The child foresaw that the stick would slide off one side. The essential point we highlight from this example is that children have prior transferable knowledge from experience with other objects in analogous situations. The idea of pressing down to keep contact is like a schema of knowledge, which potentially can also be applied to scraping food from the inside surface of a bowl. Toddlers have many of these schemas of knowledge, and when faced with a new task they will find ways to apply the appropriate schemas. This means they need to make analogical mappings.

### Theoretical Framework for Analogy

We describe now a framework for how analogy can work. There are two important processes in analogy, firstly (i) the

process of determining how best to represent or interpret two situations (often called source and target), and sometimes this means imposing a representation (or point of view) that makes them similar. Secondly, (ii) there is the process of identifying similarities between two situations; creating a mapping between component parts. Process (ii) can work in a pure form (without needing the distorting effect of process (i)), if we assume that the situations have fixed unchanging representations and the cognitive agent is navigating this landscape of representations *noticing* this and that similarity. But if we look at how analogy typically works in perceptual and cognitive processes like recognition, problem solving, concept formation, and so on, we find that analogy alters the landscape by changing representations of the situations. In this regard, it may be best to think of analogy as a process that *renders* situations similar. Such an approach requires that one makes a distinction between before-the-analogy and after-the-analogy representations, and study how representations change as a result of analogy.

Note that the dominant approaches to analogy in the literature can be characterised by their focus on one or other of the two processes sketched above. The need for an integrated theory of analogy to encompass both views has been brought out clearly in an article on “The Mistaken Fight Over The Explanation of Analogy” (Morrison and Dietrich 1995), which describes the main bodies of research on the two approaches as either horizontal (corresponding to (ii) above) or vertical (corresponding to (i) above).

As our objective here is on replicating the type of analogies that occur in infant and toddler cognition, it is crucial to consider the representation-changing role of analogy. Infant and toddler cognition develops (implying change) at a very rapid rate: infants are constantly acquiring new skills and new ways to act with familiar objects and situations, hence acquiring new perspectives on situations.

### Knowledge Structures and Functions

We follow the model of Indurkha ((Indurkha 1992) 1992, see Sec. 5.8) for a layered cognitive system. Layer 0 is the external world, layer 1 is the first level of information from sensors, and there may be any number of layers above. For example, in the human visual system, layer 0 is the actual scene a person is looking at; layer 1 is the information picked up by the retina; layer 2 may contain descriptions of edges; a higher layer may contain concepts such “chair” and “lectern”, and higher still “lecture” (these connect with more than visual perception at lower layers). The entities (e.g. “edge” or “chair” or “lecture”) at any layer are called the “symbols” of that layer.

We introduce “abstraction operations” by which a higher layer interprets the symbols in the layer below. Abstraction operations define rules for how lower-layer symbols can be combined to make a higher-layer symbol. This is easy to explain in visual processing: e.g. at layer  $n$  some surface patches with a particular orientation could qualify as a “concave” surface at layer  $n + 1$ ; another example: a car in a particular orientation at layer  $n$  is expected to have two wheels (at layer  $n - 1$ ) within a particular distance. In general, abstraction operators are not restricted to only drawing from

one layer below; they may draw on multiple layers: consider a complex concept like embarrassment for example; its presence could be inferred based on high-level indicators, such as behaviours in a particular social context and also lower-level markers like blushing.

A set of symbols in a layer can be used in a “concept network”. In a concept network the symbols are renamed as a set of object-concepts and function-concepts. Function-concepts define relationships over object-concepts: for example, if object-concepts include a pipe, and a rubber membrane blocking the middle, and the pressure of water on each side, then a function-concept might describe the relationship between pressure and the amount of water on one side. The layered architecture allows second- and higher-order functions. However, in the vision examples above there are only object-concepts (e.g. surfaces) and no function-concepts.

A concept network is an uninstantiated set of object-concepts and function-concepts; it describes the possible entities and relationships among them. An instance can then be defined using a concept network; e.g. in a concept network for car safety, we can have object-concepts for such things as tyres and tread, and various road surfaces, and function-concepts for things such as skidding and rolling, and aquaplaning. An analysis of a particular accident can provide an instance within this concept network. When an instance is constructed within a concept network from lower level perception, it will typically be partial as not all aspects are directly perceived; function-concepts can help to flesh out the missing details; e.g. in a hydraulics example, if the pressure is less than expected in a pipe, one might infer the presence of a leak. A concept network can be created for any “slice” of reality. It is like what Barsalou calls a “simulator” (Barsalou 1999). Our concept network is like the “schema” of Piaget.

### Projection (Assimilation) and Accommodation

Recognition is perhaps the most basic cognitive act. To recognize an object or a situation means to recognize it as something, and that something comes from the past experience. Here, we can consider object recognition, as ‘this is my mother’; recognizing an object as an instance of a concept, as ‘this is a cat’, and recognizing an instance of an abstract concept, as ‘they are in love,’ or ‘this is a robbery.’ Recognition is an example of analogy because situations or objects are rarely exactly the same, so some notion of similarity has to be used.

Recognition works by observing data from the environment and making a connection with a concept network inside the agent’s mind; if the data fits sufficiently well with the concept then the agent may be said to “recognise” that it encountered an instance of that concept, or the agent may “judge” that the concept is present. For example in recognising a car, components of the car concept network are mapped onto what is perceived in a lower layer (e.g. wheel symbol maps to a recognised shape combination, door maps to another, etc.). This is what Piaget called *assimilation*. We use the term projection instead of Piaget’s assimilation in order to distinguish it from the broad range of things that fall under the umbrella of assimilation in Piaget’s writings (Indurkha,

1992, p.134).

In accommodation, on the other hand, some aspects of the perceptual stimuli at the lower level are used to change the concept network in the higher level. For example, on seeing a convertible car, the car concept network may be adjusted to reflect that a roof is not necessary in a car; and later, on encountering a few more instances of convertibles, a new subcategory of cars may be created. Needless to say, accommodation is related to learning.

Both projection and accommodation may be invoked in recognising an object as an instance of a concept, e.g. a car. An instance of car is constructed (via accommodation) such that the parameters of the car in the instance have some degree of correspondence with the perceived components in the lower layer. Some details of the car as detected by lower layers in the system may be ignored at the higher layer; the car instance is approximate. The agent has some degree of choice over how much to accommodate; less accommodation would mean ignoring some of the details of this particular car (e.g. size of wheel or distance between wheels) and simply assuming typical values. The same applies to an infant grasping a new cup; it could accommodate its hand in detail to the shape of the handle of this cup, or it could just assume it is a typical cup and perform a typical grasp.

In projection (used in recognising a situation) we have a *target* which is the data of the situation, coming bottom up from our senses. We also have a *source* which is the concept network we attempt to fit. The fit happens from a higher layer to a lower layer. If the source concept network is layer  $n$ , then the target data to be interpreted is at layer  $n - 1$ . Sometimes we need to try different sources before finding a satisfactory fit. Sometimes we attempt to see things from a different perspective when we attempt to ignore the standard interpretation and force a fit with an alternative source. In this case we may have a conventional source concept network at layer  $n$  which interprets the data of layer  $n - 1$ , but we wish to interpret it in a new way by projecting an alternative source also at layer  $n$ .

Visual recognition is a special case of recognition which we focus on as a first step for our analogical system; we speculate that the same process may be at work in recognition in general. Just seeing an object as a concept  $x$  seems to use analogical machinery. For example, to recognise a teapot, perhaps an unusual one, one may recognise components such as two apertures, a spout, a handle, a flat base to set on a hot surface. These components at layer  $n - 1$  and are mapped to concepts already stored in the agent. If there are sufficient matches with the teapot concept at layer  $n$ , then a teapot has been recognised. Computer vision uses this approach via modelbases of objects. Based on a bottom up analysis of features from the image it makes a first guess (hypothesis) at the most likely object to match, and then it renders a model of that object in the right pose (this is the top-down process), and attempts to verify its hypothesis. It searches for data to confirm the possibility of the existence of expected parts which might not have been detected in the bottom up analysis (see e.g. (Shanahan 2005)). This is also called categorisation. We expect this process to be ubiquitous in the everyday life of infants, as soon as they

have acquired models for various objects (which is a long and gradual process).

## A Computational Model for Analogy in Tool Use

In this section we restrict attention to the problem of finding a suitable tool, from a set of tools presented in a scene, for a particular manipulation action. We are considering service robotics scenarios where the robot has prior knowledge of a set of tools and possible manipulations, but now needs to use a tool which is not available, and must find an object which could serve the same function as that tool.

The system must be able to generate its own creative solutions to problems instead of just expected ones based on the semantics of the human programmer. Furthermore, flexible and hierarchically structured representations would facilitate creativity and allow bottom-up and top-down reasoning. Bottom-up reasoning gives the system the needed variability of trying different approaches, and top-down reasoning keeps the system in track by driving its overall attention to one or another aspect of the problem.

Let us consider the specific case that the required manipulation is to crush an object with a force, and the hammer is the tool of choice. However the hammer is not available, so the robot looks for other objects that could substitute. A frying pan has a flat base, and a handle, and is of suitable size and weight to deliver a blow to crush something. How can we make an artificial system which could see this? The system needs to project the concept of hammer (source) down onto the components which compose the frying pan (target).

## Representations

Our specific hammer / frying pan example is illustrated in Fig. 1. The environment in which our system is immersed consists of visual scenes of 3D objects commonly found in a house. This visual data can be gathered through a Kinect camera in a simulated environment (Ellekilde and Jorgensen 2010) and a visual system, such as the Early Cognitive Vision (ECV) ((Krüger, Pugeault, and Wörgötter 2010) and (Olesen et al. 2012)), can be used to calculate 3D point clouds, which are commonly used in state-of-the-art vision systems (Rusu and Cousins 2011).

The agent's perceptual mechanism has different levels or representation. Progressing upwards, we have increasingly abstract layers: point clouds; surface patches; geons or superquadrics; relationships between geons or superquadrics; and a symbol for the object associated with a complete 3D model. With respect to using superquadrics relationships for segmentation, in (Biegelbauer and Vincze 2007), there is evidence for the importance of modelling relationships among superquadrics when detecting an object. The authors tried to detect a rubber mallet in a cluttered scene using three alternatives: just the head of the mallet; only the shaft of the mallet; and with a composite object of the hand and the shaft, and the relationship between their positions. Results showed an increase of about 25 percent when using the composite representation.

Regarding the bottom-up aspect, on each level, there are different expert algorithms trying possible segmentations and configurations of the entities. This gives rise to multiple alternative interpretations. As for the top-down pressures, we have a projection downwards trying to find matches for its lower-level counterparts, e.g. to find some surface which correspond to the hitting part of the hammer. The matching between different levels of representation does not have to be exact, but rather needs to go through some metrics of acceptability.

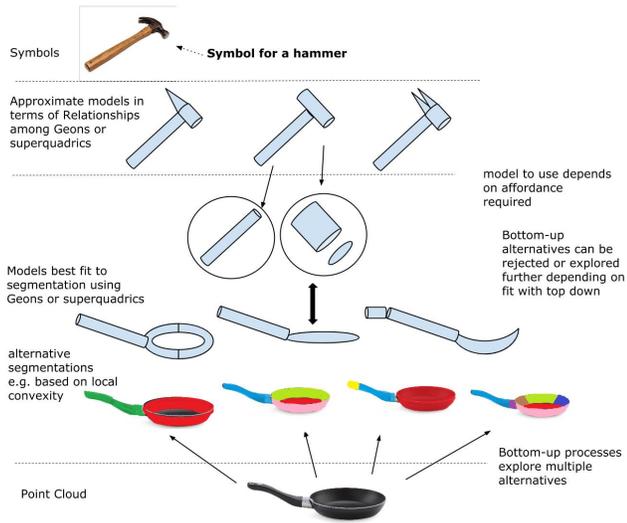


Figure 1: Different layers of representations from lower to upper level: point clouds; alternative segmentations; geons or superquadrics; approximate models in terms of relationships among geons or superquadrics; and symbols. An affordance of an object directs the decomposition process. In the middle, we see a decomposition for smashing which is able to find counterparts on segmentations for a frying pan.

The agent is given prior knowledge about a set of tools and their typical manipulation actions. For each tool, it has a cognitive structure relating the object to its known parts. For example, a hammer may have a possible decomposition as a cylinder (handle) attached to another transversal smaller cylinder (head). Then, the head of the hammer may have two different decompositions: one as a cylinder with a flat surface (face) and another as a cylinder with a flat surface attached to two cones (claw). An object's affordance should influence the decisions of matching between levels, during projection. For instance, if the agent wants a hammer to smash something, then the claw on the head of the hammer should receive a lower level of necessity in comparison to the face and the handle. Therefore, the relationship between an object and its affordance determines the possible decompositions and the degrees to which certain aspects of the representation may be relaxed: When I do not have a hammer and I want to remove a nail from a board of wood, I need to project a certain aspect of the hammer onto the available objects, namely, the characteristics of a hammer claw which

enables one to pull a nail out.

## Functioning

Now that we have the representations and the relationships between the different entities, we need to describe how the agent would go about using them to come up with a solution to its problem. Our architecture follows the ideas of the parallel terraced scan (Hofstadter 1996) of FARG (Fluid Research Analogy Group). This is a stochastic parallel architecture composed of microagents scanning the environment and building transient relationships between concepts. Each microagent does some small work and a regular structure arises from their interactions (a helpful picture to bear in mind is that of an ant colony, where you have different specialised agents building a larger and more regular structure).

There are microagents on each level working in parallel trying to match the top-down pressure with the bottom-up possibilities. Top-down pressures influence the level of activity of certain microagents, which, in turn, when finding appropriate matches, send a bottom-up message reasserting a top-level representation. The system's goal is converging to settle down on a satisfactory projection. For instance, imagine an agent who becomes interested in finding a hammer to smash a fly; this will increase the level of attention on the hammer schema and certain decompositions, increasing attention on the handle and on the hammer head as having a flat surface. If there is only a frying pan available, the agent will have to balance its attention and activity on certain decompositions and level matching in order to be able to temporarily project the hammer schema onto the frying pan.

## Conclusion

We have emphasised that analogy-making is a necessary cognitive skill that must be incorporated in robotic systems in order that they can handle a variety of tasks without explicit coaching for each one. In this regard, we find that it may be very useful to draw upon the insights from the evolution and development of tool use in infants and toddlers (Guerin, Kraft, and Krüger 2013). Of course, there is over thirty years of research on incorporating analogy into robotic and AI systems, but most of these models have focused on either the horizontal aspect of analogy-making, which works by mapping between the given representations of the source and the target, or on the vertical aspect, which works on building representations from scratch. In this paper, we have outlined an architecture for combining these two approaches. The key features of this architecture is a layered representation, and mechanisms of projection and accommodation that work in consort for recognising that an object is an instance of a concept. We outlined a computational model for this architecture, and sketched how it would work using a hammer / frying pan example. We are working on implementing this architecture.

Though we have not elaborated on this in this paper, we believe that playfulness and imagination are two major aspects of cognition that emerge from our proposed architec-

ture, and that certain aspects of creativity can be modelled in this as well.

## References

- Barsalou, L. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22:577–660.
- Biegelbauer, G., and Vincze, M. 2007. Efficient 3D Object Detection by Fitting Superquadrics to Range Image Data for Robot's Object Manipulation. In *Proceedings of the Fourth IEEE International Conference on Robotics and Automation*, 1086–1091. Rome: IEEE.
- Brown, A. L. 1990. Domain-specific principles affect learning and transfer in children. *Cognitive Science* 14(1):107–133.
- Ellekilde, L., and Jorgensen, J. 2010. Robwork: A flexible toolbox for robotics research and education. In *Proceedings of the Joint Forty First International Symposium on Robotics and Sixth German Conference on Robotics*, 1–7. Munich: VDE.
- Fagard, J.; Grupen, R. A.; Guerin, F.; and Krüger, N. 2013. Mechanisms of Ongoing Development in Cognitive Robotics (Dagstuhl Seminar 13072). *Dagstuhl Reports* 3(2):55–91.
- Guerin, F.; Kraft, D.; and Krüger, N. 2013. A survey of the ontogeny of tool use: from sensorimotor experience to planning. *IEEE Transactions on Autonomous Mental Development* 5(1):18–45.
- Hofstadter, D. R. 1996. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York, NY: Basic Books, Inc.
- Indurkha, B. 1992. *Metaphor and Cognition*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Krüger, N.; Pugeault, N.; and Wörgötter, F. 2010. Visual primitives: local, condensed, semantically rich visual descriptors and their applications in robotics. *International Journal of Humanoid Robotics* 07(03):379–405.
- Morrison, C. T., and Dietrich, E. 1995. Structure-mapping vs. high-level perception: The mistaken fight over the explanation of analogy. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, 678–682. Pittsburgh, PA: Psychology Press.
- Olesen, S.; Lyder, S.; Kraft, D.; Krüger, N.; and Jessen, J. 2012. Real-time extraction of surface patches with associated uncertainties by means of Kinect cameras. *Journal of Real-Time Image Processing* 1–14.
- Piaget, J. 1936. *The Origins of Intelligence in Children*. London: Routledge & Kegan Paul. (French version 1936, translation 1952).
- Rat-Fischer, L.; O'Regan, J. K.; and Fagard, J. 2012. The emergence of tool use during the second year of life. *Journal of experimental child psychology* 113(3):440–6.
- Rusu, R. B., and Cousins, S. 2011. 3D is here: Point Cloud Library (PCL). In *Proceedings of the Eighth IEEE International Conference on Robotics and Automation*, 1–4. Shanghai, China: IEEE.
- Shanahan, M. 2005. Perception as abduction : Turning sensor data into meaningful representation. *Cognitive science* 29(1):103–134.
- Uzgiris, I. C., and Hunt, J. M. 1975. *Assessment in infancy : ordinal scales of psychological development*. Champaign, IL: University of Illinois Press.