Processing-based concept kinds for actor-agent communities

Manuela Viezzer[1,2] and **Niek Wijngaards** [2] and **Masja Kempen**[2]

Abstract. Engineering complex highly-interactive systems consisting of both human and artificial agents (actor-agent communities) requires insight in the use and role of concepts by individuals and in interaction between individuals (both human and artificial). In this philosophically-oriented paper, the distinction between *concept kinds* is found to depend on processing differences for these kinds, rather than content-based or structural differences. In addition, this leads to the characterization of a new concept kind: affordance concepts. Our next step is to a) experiment with acquisition of affordance concepts and b) investigate the role of affordances for strategic management of teams.

1 INTRODUCTION

In the near future, humans and artificial systems are foreseen to engage in (even closer) collaboration. These organisations are named *actor-agent communities*, in which collaboration between multiple participants, including humans and artificial systems, takes place for the realization of a common mission or for the support of a shared process [45]. Within a community there are social rules that members adhere to, and there is communication, sharing of responsibility and a certain distinct identity among the community members. From a human perspective, an actor-agent community is not unlike any conventional human community and the same traits apply. From a technical perspective, an actor-agent community contains distributed systems and processes that have autonomous and anticipatory capabilities and software systems that can be referred to as agents or agent systems, which are not intended to be 'artificial humans'. An important characteristic is that agents can compensate for human shortcomings and amplify human competencies [23]. In addition, they can inhabit more worlds (virtual ones, for example) than the physical 3D world where humans operate. Actor-agent communities are typically involved in complex collaborative decision making processes, such as the day-to-day (air) traffic management or crisis response and management.

Actor-agent communities (AACs) require analysis as a whole, including the triple: actors, agents and their 'niche' [23, 36], where niche is characterised by requirements, constraints and opportunities. The intense collaboration between actors and agents places additional requirements on the development of artificial agents: the agents need to 'understand' humans - as communication is a pre-requisite for collaboration. A central concern in cognitive engineering is to provide artificial systems with work domain representations that are compatible with those elaborated by humans. Such a compatibility should span the hierarchical organisation of the domain as well as its conceptual building blocks[41]. The identification of concept kinds that can be common to both human and artificial minds is a means via which shared situation awareness and shared meaning can be facilitated. For example, the cognitive architecture of an agent can be designed to support the use and manipulation of such concept kinds. Since agents in AACs are placed in a dynamic environment, they need to adapt to unforeseen circumstances, changing goals, plans, organisational structures and behaviour of individuals. This requires not only adaptation of behaviour, but also the ability to acquire new concepts (and relations between concepts) to remain effective (albeit to a certain degree) when their environment changes.

Unfortunately, despite the wide agreement on the importance of concepts as major constituents of human cognition, no unified, comprehensive and well-established theory of concepts exists, and different empirical findings seem to support different, when not incompatible, views on the matter [26]. Moreover, even though research on concepts has been abundant in many fields, notably psychology, philosophy, cognitive science and AI, the question of whether there are distinct kinds of concepts has been rarely addressed, because most of the research on concepts has focused almost exclusively on natural object concepts (chair, bird, tool, etc.), within categorisation tasks [37, 38].

The purpose of this paper is to further our understanding of conceptual structures by addressing the question of whether distinct kinds of concepts can be individuated, and on which basis. The importance of a framework for distinguishing among different concept kinds, including models for acquisition, reasoning, and other manipulations, thereof should not be underestimated. When developing (intelligent) agents in general, it is imperative to choose the right processing mechanisms for the right purposes. If it is known which concept kinds can be distinguished and if the distinction is based upon the existence of processing differences between those kinds, important parts of an agent's cognitive architecture can be specified. Although it is compelling to look for only one processing mechanism 'which suits all needs', pragmatic experience dictates that multiple processing mechanisms within one agent are mandatory as the 'needs' can be extremely diverse. It is expected that the same holds for processing mechanisms for manipulation of concepts. Our analysis shows that processing-based distinctions facilitate the identification of concept kinds, basically affordance concepts and goal-based concepts on one side, and taxonomic concepts on the other. The processing mechanisms taken into consideration include reasoning with and acquisition of concepts.

The paper is structured as follows. Section 2 broadens the notion of context to accommodate for agents that are more than reasoners.

---

[1] School of Computer Science, The University of Birmingham, B15 2TT Birmingham, United Kingdom

[2] DECIS Lab, Thales Research and Technology Netherlands, PO BOX 90, 2600 AB Delft, The Netherlands. Contact author: Manuela.Viezzer@decis.nl

Section 3 discusses different criteria to distinguish among concept kinds, and argues for the superiority of a processing-based distinction. We examine the proposed categories based on structural and processing differences, and show that the only acceptable structural-based distinction actually reduces to a distinction based on how the involved concept kinds are processed. We then turn to the processing-based criterion, and account for two distinct kinds of concepts: taxonomic concepts on the one hand, and goal-derived concepts on the other. Section 4 describes affordance concepts as a special kind of concepts, and shows how affordance concepts can be characterised in processing-based terms, along similar lines as goal-derived concepts. Finally Section 5 contextualises our work and describes our intended future research.

2 BROADENING THE NOTION OF CONTEXT

The problem of contextual reasoning has been stated as the problem of understanding and then formalise the reasoning mechanisms employed by people when dealing with information "such that (i) its representation depend on a collection of contextual parameters, and (ii) is scattered across a multiplicity of different contexts." [9] Following this intuition, the notion of context is given a cognitive flavour and is understood as "a partial and approximate representation used by an agent to interact with the environment and with other agents." An agent taking part in an actor-agent community however, is usually more than a database of knowledge and requires the term 'reasoning' to be interpreted in many ways, ranging from the application of logical rules to more general adapting, learning and problem solving abilities. For the sake of clarity, we assume 'reasoning' to imply the first characterisation, and 'processing mechanisms' to imply the latter; in specific to include both acquisition (or learning) and reasoning.

In this larger perspective, the 'box metaphor' proposed in [9] to illustrate the notion of context-dependent representation can be broadened (see Figure 1) and applied to behaviours in general (not only semiotic ones). Inside the box we have any behaviour of an agent; outside the box we have a collection of processing mechanisms capable of generating that behaviour. The intuition is that the characterisation of what is inside the box (whether the agent is expressing an emotion, or whether it is acting intentionally, etc.) is determined by the kinds of processing supporting and generating it. Instead of context-dependent representations, we will therefore speak of context-dependent behaviours and contextual processing. Insofar as cognitive behaviours depend heavily on the use of concepts, our assumption is that a theory of contextual processing should accommodate, among others, the existence of conceptual structures which requires radically different processing mechanisms for their acquisition and use.

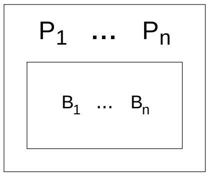

Figure 1. A 'box metaphor' for contextual behaviour.

In the next section we examine different criteria to distinguish among concept kinds, and argues for the superiority of a processing-based distinction. According to this last criterion, we account for two distinct kinds of concepts: taxonomic concepts on the one hand, and goal-derived concepts on the other.

3 COMPARING CRITERIA FOR CONCEPT KINDS

As already noticed in the Introduction, even though research on concepts has been extensive and varied, the question of whether there are distinct kinds of concepts has been rarely addressed. One of the few attempts to bring together candidates for kinds of concepts that have emerged across the different theoretical perspectives is to be found in [33], where Medin *et al.* propose criteria based on structural differences, processing differences, and content differences for distinguishing concept kinds.

Already in [33] the idea that kinds of concepts can be individuated by referring to content-laden differentiating principles is criticised. The content-based criterion in fact is grounded on the claim that different domains of conceptual knowledge can be individuated (e.g. naïve biology, naïve psychology, naïve physics). Such a claim however, though having some empirical foundations [13], is far from being validated, and it is very difficult to give a precise definition of domain [22]. For this reason we maintain that the actual choice for providing a characterisation of concept kinds is between structural and processing differences.

3.1 Structural differences

Criteria for kinds of concepts based on structural differences would distinguish among different kinds of concepts by looking at differences in the kinds of their constituent features, and in the relations among them. The implicit assumption is therefore that concepts can be analysed into constituent features. Usually, evidence is first provided for a category distinction. Then the hypothesis of structural differences among the corresponding concepts is stated, assuming that concepts must be distinct if the associated categories are distinct.

In the literature, distinctions have been proposed that rely on the following oppositions: i) necessary and sufficient defining features *vs.* defining features that are only probable; ii) compositional features *vs.* functional features; and iii) perceptual features *vs.* relational features.

Case i): The classical view on concepts, according to which concepts are structured entities encoding necessary and sufficient defining features for their instances, is opposed to the prototype theory, according to which only probable features are encoded [26]. The opposition between the two views is cast in terms of how concept constituent features are to be used in the process of identifying members of a given category: an all or none matching process is used in the classical case, while a partial or probabilistic matching process is used in the prototype case. However, in [3] it is argued that, since graded responses are found also for well-defined categories such as PRIME NUMBER, the distinction between well-defined and fuzzy categories is ill-grounded. This suggests to discard the idea of a distinction between classical (or well-defined) concepts on the one hand, and prototype concepts on the other.

Case ii): Another category distinction that has been proposed [25, 8, 34] opposes natural kinds to artifacts on the basis that people consider features referring to internal structure to be criterial for membership within natural kinds, while they consider functional features to be criterial for membership within artifact categories. For example, [34] shows that external transformation (e.g. due to chemical hazards) is less likely to change natural kinds' identity than internal transformation (such as internal maturation), while artifacts membership tends to change if the transformation involves the function of the artifact (e.g. a tire that cannot roll after transformation).

However, a successive study [29] shows that in certain cases physical properties of artifacts (for example, for boats, the fact that they are wedge-shaped, with a sail and anchor, etc.) were judged to be more important than, or just as important as functional properties (such as carrying people over a body of water for purpose of work or recreation), thus making the alleged distinction less clear-cut. A deeper explanation of the various data, which provides a coherent picture of the phenomena, is to be found in [2]. In this work Ahn proposes a causal status hypothesis, according to which the centrality of a feature to a category does not depend on the kind of feature but rather on the causal role that the feature plays (relative to the other features in the category). In other words, given a set of causally related features, the experiments done by Ahn show that in general people are biased toward considering features that serve as causes of other features more central to category membership than their effects. These findings considerably reduce the claim that natural kinds and artifacts are two separate domains and that different features are more central in categorisation as a result of this domain distinction.

Case iii): In linguistics and psychology a distinction is made between object (or noun) concepts, corresponding to clusters of perceptual features, and relation (or verb) concepts, corresponding to relational properties. Cross-linguistical studies such as [27] show that lexical entries corresponding to perceptual chunks are consistent through different languages/cultures, while lexical entries corresponding to relational structures tend to vary. Also, studies on how children acquire language, such as [40, 10, 44, 46], seem to support an early bias towards noun acquisition and production.

In order to explain which words are earliest learnt by children and why, in [16] Gentner *et al.* show that the word acquisition process can be characterised by either *cognitive dominance* or *linguistic dominance*. In the case of cognitive dominance, aspects of perceptual experience form inevitable conflations that are conceptualised and lexicalised, while in the case of linguistic dominance, the world presents perceptual bits whose clumping is not pre-ordained, and it is the linguistic interaction that guides how the bits get conflated into concepts. The authors claim that dominance is divided along the continuum from open-class terms to closed-class terms: “At one extreme, concrete nouns — terms for objects and animate beings — follow cognitive-perceptual dominance. They denote entities that can be individuated on the basis of perceptual experience. At the other extreme, closed-class terms — such as conjunctions and determiners — follow linguistic dominance.” The consequences in term of early word learning is an early predominance of names for objects and individuals, and a later increase in the proportion of relational terms. This is so because noun referents are easier to individuate than verb and other closed-class terms referents. The child’s task when learning language is one of attaching words in the stream of speech to their referents in the stream of experience, via the conceptual space. Given the fact that concrete objects have already been individuated prelinguistically, as reported in [4], the acquisition of concrete nouns reduces to finding the correct linguistic label. In contrast, as argued in [16] “for verbs and other relational terms, isolating the word is only part of the job. The child must also discover which conflation of the available conceptual elements serves as the verb’s referent in her language.”

In the light of evidences from the fields of psycholinguistics and language development, the distinction between object concepts, based on the clustering of perceptual features, and relation concepts, based on the clustering of relational features, appears to be granted. However, the hypothesis of the division of dominance refers to the acquisition process of the proposed concept kinds (object concepts are acquired pre-linguistically while relation concepts are acquired through linguistic interaction), and therefore establishes a processing-based distinction rather than a structural one.

3.2 Processing differences

When processing differences matter, the candidates for concept kinds that have been discussed in the literature are taxonomic concepts *vs.* goal-derived concepts.

Taxonomic concepts: Taxonomic categories and the correlated concepts have been proposed and studied by Rosch [35]. According to Rosch, two general principles underlie the human ability to form taxonomies: the cognitive economy principle, asserting that the task of a category system is to provide maximum information with minimum cognitive effort; and the perceived world structure principle, asserting that the perceived world comes as structured information rather than as arbitrary or unpredictable attributes.

This last principle in particular explains the characteristic exemplar-based formation process of taxonomic concepts, which can be summarised as follows [30]: It is an empirical fact that certain combinations of attributes are quite probable, in the sense that they often appear together (for example, wings co-occur with feathers more than with fur, shape co-occur with similarity of movement more than with similarity of colour, etc.) Humans seem to be predisposed to take certain attributes as reliable indicators for kind membership (‘kind syndrome’). Therefore, they tend to group together objects that are similar under those attributes as belonging to the same kind. This provide sets of clustered exemplars that can be successively analysed to extract cognitively more economical representations for the clusters.

Taxonomic concepts and the related categories are therefore created to capture the correlational structure of the physical world, and their acquisition process is a case of exemplar-based learning. Since category knowledge is induced from experiences with exemplars (and then internally represented in a more economical form), clusters of exemplars must be readily available before the actual learning process starts. In particular, the clustering is done according to the ‘kind syndrome’ as discussed by Margolis in [30].

Goal-derived concepts: In [5, 6, 7] Barsalou contrasts taxonomic categories (also referred to as object categories) and goal-derived categories on the basis of several differences: unlike object categories, goal-derived categories may activate context-dependent features of category members; their members are not especially similar to one another, so they violate the correlational structure of the environment that is exploited by taxonomic concepts; their typicality is based on proximity to ideals rather than on central tendency[3]; they are acquired through conceptual combination rather than exemplar-based learning. It is this last issue that grounds the distinction between taxonomic concepts and goal-derived concepts in terms of processing mechanisms.

---

[3] An ideal is a characteristic that exemplars should have to serve a goal optimally. For example, *zero calories* is an ideal for the category *food* in the context *foods to eat on a diet*. Central tendency, instead, is the average characteristics of a category’s exemplars. Both ideals and central tendency information may provide content for a prototype, and proximity to ideals works also for taxonomic concepts, e.g. in expert judgements about typicality. Actually, it seems that central tendency plays a role only for relative novices [28].

According to Barsalou, goal-derived concepts and the corresponding categories are built when constructing plans to achieve goals, for example the concepts ‘people to visit in California’ or ‘foods to eat on a diet’. The methodology proposed in [7] to study goal-derived concepts and categories is therefore based on a protocol analysis of planning. When planning an event (e.g. a vacation), the first thing that seems to be done by people is to retrieve a frame for that event, that is a set of attributes that can take different values across different instantiations of the same event. Usually, the frame is retrieved only partially and not as a rigid structure: sometimes important attributes are just ‘forgotten’ by one subject; the order in which attributes are considered varies with the subject; etc. However, once the appropriate frame has been retrieved, people begin to instantiate the various attributes, adopting particular values for use in the current frame. This instantiation process normally follows a process of successive refinements: first, a general class of instantiations is produced, and then this class is refined down to more specific instantiations. For example, when trying to instantiate the temporal attribute of a vacation, a subject might begin with the concept ‘in October’; that same subject might then continue with the refinement ‘towards the end of the month, possibly a week-end’, ending up with the more specific concept ‘a week-end towards the end of October’.

These successive sets of instantiations for frame attributes are what Barsalou calls ‘goal-derived concepts’ and the previous example shows both how goal-derived concepts are involved in planning, and also how their formation exploits a process of conceptual combination. Such a process is substantially different from the exemplar-based learning involved in the acquisition of taxonomic concepts. Goal-derived concepts (and their corresponding categories) are acquired through conceptual combination because the process of frame attributes instantiation giving rise to them, actually involves the combination of increasingly specific properties with the existing, current attribute description. This implies that exemplars for these concepts are not known before hand, but are actually formed while the concept and its corresponding category are formed.

In the next section we discuss affordance concepts and show how this kind of concepts can be characterised in processing-based terms, along similar lines as goal-derived concepts.

4 AFFORDANCES AND AFFORDANCE CONCEPTS

The notion of *affordance* is due to J. J. Gibson, but has never been given an ultimate definition by its creator [24]. Along Gibson’s writing in fact there is a tension between considering affordances as ‘perceptual invariants’ [21] and considering them as ‘possibilities for actions’ [19, 20]. In order to define affordance *concepts*, we refer to the second connotation.

4.1 Introducing affordance concepts

According to Gibson, affordances become apparent when perception is approached from an ecological perspective. A basic assumption of the ecological perspective is that the animal and the environment have been co-designed through evolution and are therefore ‘mutually compatible’. This implies that the animal and the environment can only be adequately described when considered in relation to one another. So, to say that animals can detect affordances actually is to say that animals have been designed by evolution to detect properties of the environment that are directly *relevant* to them, usually because their survival depends on these properties (for example, surfaces, textures, substances, but also properties such as ‘graspable’, ’stand-on-able’, sit-on-able’, etc.)

To make this notion of relevance more precise, we say that given an autonomous agent inhabiting an environment where different kinds of objects exist, these objects can be partitioned into relevant and irrelevant (for that agent) depending on whether they can modify the agent’s internal states, and hence be useful or dangerous, when acted upon. Slightly differently, if the agent has goals pertaining to its internal states (e.g. the goal of increasing its internal energy), we also say that a category of objects is relevant for an agent either because the objects can be used to achieve the agent’s goals, or because they prevent such an achievement. For example, if the agent has the goal of increasing its internal energy and eats an object O to achieve such goal, then O is said to belong to a relevant category, the category ‘Food’ or ‘Edible items’.

Our proposal is to consider an affordance concept to be a concept denoting a relevant category of objects, in the sense defined above. Affordance concepts are tightly related to the notion of action, insofar as the objects subsumed by affordance concepts are individuated by acting on them. Given an agent and an action a performed by that agent, we can in fact distinguish the objects that satisfy the agent’s expectations on the results of acting with a upon them, and those that do not. This leads to the notion of *positive* and *negative* affordance concepts, respectively. For example, if we call Food the class of all the objects that satisfy an agent’s expectations on the results of eating (e.g. an increase in its energy level), then the concept associated to it is a positive affordance concept, and we say that ‘food items’ afford eating.

4.2 Processing mechanisms

How can we tell whether an agent can think of an affordance concept, for example one that corresponds to the category Food? Assume we observe an agent and see that, when it is hungry, it does not just try to eat whatever happens to be near it, but continues moving around until it finds something that has a certain appearance, something that ‘looks like food to it’. To generate the observed behaviour, the agent’s architecture might contain a condition/action rule expressing something like *If hungry and sense item with features* $f_1, ..., f_n$ *in visual field, then approach it and eat it; else move on*. Shall we say that such an agent has an affordance concept FOOD that corresponds to the category Food, as described above? Our answer is no, or at least not if this is the whole story about our agent. Probably, something like the rule we sketched above should be present in the agent’s architecture, but what really matters in order to establish whether the concept FOOD is present, is *how* such a rule happened to be part of the architecture, in particular whether it is innate or acquired, and whether it has special relations with other rules. Our intuition is that we should say that the concept FOOD is present if i) the rule has been acquired and ii) the acquisition process exploited the relation between the action of eating and the expectation of an increase of the internal energy as a result of eating.

Many studies in the past quarter century have examined the physical basis for perception of affordances, such as the climbability of steps [43], the safety of gaps [12, 32] and the suitability of different surfaces for locomotion by infants at different stages [1]. These studies address the issue of unveiling the invariants grounding the affordance in question (e.g. a stair raiser being less than 88% of a person’s leg length), but the issue of whether and how affordances (and affordance concepts) are learned is rarely addressed. One exception is the theory proposed by E. Gibson in [17, 18], according to which learning affordances is a process of differentiation and selection which involves learning the relations between the organism's power of control and some opportunities or constraints offered by the environment. The first step in affordance learning is an exploratory activity, whose main output is learning control of an event. This requires learning how to orchestrate a sequence of self-generated actions in preparation for an anticipated outcome. An important part of the exploratory activity is the observation of consequences of actions. Human perceptual systems are designed for cycles of perception-action, and when some kind of environmental contact is achieved through action, the contact will be perceived in its turn, and its usefulness (uselessness) can be evaluated.

The next step is a selection process whose output is the selection or recognition of an affordance relation. E. Gibson proposes two main principles of selection: i) selection for an affordance fit, and ii) selection for unity, order and economy. What has to be differentiated and then selected is a sequence of actions. The selection is determined (at least partly) by the kind of fit between actions performed and the ensuing consequences. In most cases, good fits are those obviously serving functions that are of biological value to the organism.

Since economy of action and reduction of perceptual information also works as principles of selection for increasing specificity, over development perception-action cycles becomes more and more refined and better differentiated thanks to constant practice. An example is the achieving of economy by generating a characterisation of classes of objects presenting a certain affordance fit, such as graspable or edible, in terms of external features.

Some recent findings reported in the AI literature appear to corroborate E. Gibson's insight. In particular, [39] shows that given the desired effect (in this case, the attachment of an object to a robot's body, so that the object's movements can be controlled) and information about an object shape, the appropriate behaviours to be performed to obtain the desired effect can be induced. This corresponds to the exploratory activity step individuated by E. Gibson. Examples of the second step individuated by E. Gibson, the selection process, are given in [14, 15, 42] where the authors show that given an hard-wired association between survival-related internal variables, motivations and behaviour execution, it is possible to infer whether a particular object in the environment is suited to a particular interaction. This actually corresponds to learning information about the objects' potential for action, that is learning affordance concepts in the sense described above. To summarise, the architectural feature that enables an agent to acquire affordance concepts is the fact that any internal state of the agent which deviates from stability triggers both the intention to perform a certain action, and an expectation, with a definite content, on the result of the action (in terms of a resulting new internal state). For any agent that has these initial characteristics there is a *space of possible affordance concepts* that the agent can learn.

4.3 An additional concept kind

We can now compare affordance concepts and goal-derived concepts with respect to the processing mechanisms that needed to acquire them.

In both cases, exemplars are not known before the concept and its corresponding category are formed. In the case of affordances, the subject does not know in advance whether or not a certain exemplar is useful (or useless) in a given situation, so the first step of the learning process is the formation of a notion of usefulness, which constitute the functional aspect of the affordance concept; this notion is then used to judge in any particular situation whether an encountered exemplar is or is not a good candidate for that particular notion of usefulness. Affordance learning is not therefore an exemplar-based learning. However, it does also not rely on a process of conceptual combination, as it was the case for goal-derived concepts, because the formation of the notion of usefulness is grounded on the actual interaction with the environment, and not on some previously established knowledge of other concepts.

To conclude, we can say that affordance concepts and goal-derived concepts can be grouped together as concepts that are learnt in a mediated way, in the sense that clusters of exemplar are not available before the learning process; instead, a prior formation of clustering criteria is required. Both affordance and goal-derived concepts can therefore be opposed to taxonomic concepts, which are exemplar-based. On the other hand, affordance concepts and goal-derived concepts are distinct with respect to the intermediate clustering step that is involved: goal-derived concepts require a process of conceptual combination, while affordances require direct interaction with the environment.

5 DISCUSSION

An important finding in this paper is that processing-based distinctions facilitate the identification of concept kinds. For goal-related concepts different processing is required (usually without the use of exemplars) than for taxonomic concepts. An additional finding by distinguishing processes is that affordance-related concepts can be identified as a separate concept kind. The proposed characterisation of those concept kinds can be summarised as follows: Taxonomic concepts accounts for regularities in the environment, independently of the agent, and rely on observation. Affordances and goal-based concepts account for regularities in the coupling of the agent with its environment, and rely on interaction. The distinction of different learning processes (for goal-related concepts: conceptual combination; for affordance-related concepts interaction with the environment; for taxonomic concepts: exemplar-based learning) gives additional insight in the role and use of concepts in agents.

Our main research effort is focused on actor-agent teams AACs. With respect to the role and use of concepts and concept kinds in AACs, we plan the following subsequent research activities. On the one hand we aim to experiment with affordance-based learning approaches to enable an agent to adapt its behaviour to its (changing) simulated environment, albeit at the level of interaction with (physical) objects and properties thereof. In addition, we intend to explore the notion of affordances, to the extent that it can apply to strategic notions, such as those found in the management of teams. This is to be grounded by experimentation with human teams, as well as by experimenting with agents capable of applying affordances at a strategic level.

ACKNOWLEDGEMENTS

The research reported here is part of the Interactive Collaborative Information Systems (ICIS) project (`www.icis.decis.nl`), supported by the Dutch Ministry of Economic Affairs, grant nr: BSIK03024.

Thanks to the anonymous referees for their comments which led to a number of improvements.

Part of this work was conducted at the School of Computer Science, The University of Birmingham, under the supervision of Prof.

A. Sloman, to whom goes a special thanks for his guidance and advice.

REFERENCES

[1] K. E. Adolph and M. A. Eppler, 'Development of visually guided locomotion', *Ecological Psychology*, **10**, 303–322, (1998).

[2] W. Ahn, 'Why are different features central for natural kinds and artifacts?: the role of causal status in determining feature centrality', *Cognition*, **69**, 135–178, (1998).

[3] S. L. Armstrong, L. R. Gleitman, and H. Gleitman, 'What some concepts might not be', In Margolis and Laurence [31].

[4] R. Baillargeon, 'The object concept revisited: new directions in the investigation of infants' physical knowledge', in *Carnegie Mellon Symposia on Cognition: Visual perception and cognition in infancy*, ed., C. E. Granrud, volume 23, 265–315, Lawrence Erlbaum, Hillsdale, NJ, (1993).

[5] L. W. Barsalou, 'Ad hoc categories', *Memory and Cognition*, **11**, 211–227, (1983).

[6] L. W. Barsalou, 'Ideals, central tendency, and frequency of instantiation as determinants of graded structure of categories', *Experimental psychology: learning, memory and cognition*, **11**, 629–654, (1985).

[7] L. W. Barsalou, 'Deriving categories to achieve goals', in *The psychology of learning and motivation: Advances in research and theory*, ed., G. H. Bower, volume 27, 1–64, Academic, San Diego, CA, (1991).

[8] M. E. Barton and L. K. Komatsu, 'Defining features of natural kinds and artifacts', *Journal of Psycholinguistics Research*, **18**, 433–447, (1989).

[9] M. Benerecetti, P. Bouquet, and C. Ghidini, 'Contextual reasoning distilled', *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)*, **12**(3), 279–305, (2000).

[10] P. Bloom, E. Tinker, and C. Margulis, 'The words children learn: evidence against a noun bias in early vocabularies', *Cognitive Development*, **8**, 431–450, (1993).

[11] *Language Acquisition and Conceptual Development*, eds., M. Bowerman and S. Levinson, Cambridge University Press, Cambridge, UK, 1999.

[12] G. Burton, 'The role of the sound of tapping for nonvisual judgement of gap crossability', *Journal of Experimental Psychology: Human Perception and Performance*, **26**, 900–916, (2000).

[13] S. Carey, *Conceptual Change in Childhood*, Bradford Books, Cambridge, MA, 1985.

[14] I. Cos-Aguilera, L. Cañamero, and G. Hayes, 'Motivation-driven learning of object affordances: First experiments using a simulated khepera robot', in *The Logic of Cognitive Systems. Proceedings of the Fifth International Conference on Cognitive Modelling, ICCM*, eds., F. Detje, D. Dörner, and H. Schaub, pp. 57–62. Universitäts-Verlag Bamberg, (2003).

[15] I. Cos-Aguilera, L. Cañamero, and G. Hayes, 'Using a sofm to learn object affordances', in *Proceedings of the Workshop on Physical Agents (WAF04)*, University of Girona, (2004).

[16] D. Gentner and L. Boroditsky, 'Individuation, relativity and early word learning', In Bowerman and Levinson [11].

[17] E. J. Gibson, 'Perceptual learning in development: some basic concepts', *Ecological Psychology*, **12**(4), 295–302, (2000).

[18] E. J. Gibson, 'Where is the information for affordances?', *Ecological Psychology*, **12**(1), 53–56, (2000).

[19] J. J. Gibson, *The perception of the visual world*, Houghton Mifflin, Boston, 1950.

[20] J. J. Gibson, *The senses considered as perceptual systems*, Houghton Mifflin, Boston, 1966.

[21] J. J. Gibson, *The ecological approach to visual perception*, Laurence Erlbaum, Hillsdale, NJ, 1979.

[22] L. A. Hirschfeld and S. A. Gelman, 'Domain specificity: an introduction', in *Mapping the Mind: Domain Specificity in Cognition and Culture*, eds., L. A. Hirschfeld and S. A. Gelman, 3–55, Cambridge Univ. Press, Cambridge, UK, (1994).

[23] R. R. Hoffman, P. Hayes, K. M. Ford, and P. Hancock, 'The triples rule', *IEEE Intelligent Systems*, **17**(3), 62–65, (2002).

[24] K. S. Jones, 'What is an affordance?', *Ecological Psychology*, **15**(2), 107–114, (2003).

[25] F. C. Keil, *Concepts, Kinds, and Cognitive Development*, MIT Press, Cambridge, MA, 1989.

[26] S. Laurence and E. Margolis, 'Concept and cognitive science', In Margolis and Laurence [31].

[27] S. C. Levinson, 'Covariation between spatial language and cognition and its implications for language learning', In Bowerman and Levinson [11].

[28] E. B. Lynch, J. D. Coley, and D. L. Medin, 'Tall is typical: Central tendency, ideal dimensions and graded category structure among tree experts and novices', *Memory and Cognition*, **28**(1), 41–50, (2000).

[29] B. C. Malt and E. C. Johnson, 'Do artifacts concepts have cores?', *Journal of Memory and Language*, **31**, 195–217, (1992).

[30] E. Margolis, 'Concepts: core readings', In Margolis and Laurence [31].

[31] *Concepts: core readings*, eds., E. Margolis and S. Laurence, The MIT Press, Cambridge, Massachusetts, 1999.

[32] L. S. Mark, Y. Jiang, S. S. King, and J. Paasche, 'The impact of visual exploration on judgements of whether a gap is crossable', *Journal of Experimental Psychology: Human Perception and Performance*, **25**, 287–295, (1999).

[33] D. L. Medin, E. B. Lynch, and K. O.Solomon, 'Are there kinds of concepts?', *Annu. Rev. Psychol.*, **51**, 121–147, (2000).

[34] L. J. Rips, 'Similarity, typicality and categorisation', in *Similarity and Analogical Reasoning*, ed., A. O. Stella Vosniadou, 21–59, Cambridge University Press, New York, (1989).

[35] E. Rosch, 'Principles of categorization', in *Cognition and categorization*, eds., E. Rosch and B. Lloyds, Lawrence Erlbaum, (1978).

[36] A. Sloman, 'Interacting trajectories in design space and niche space: a philosopher speculates about evolution', in *Lecture Notes in Computer Science*, volume 1917, 3–16, Springer-Verlag, Berlin, (2000).

[37] E. Smith and D. Medin, *Categories and concepts*, Harvard University Press, Cambridge, Mass., 1981.

[38] K. O. Solomon, D. L. Medin, and E. B. Lynch, 'Concepts do more than categorize', *Trends in Cognitive Sciences*, **3**(3), March, (1999).

[39] A. Stoytchev, 'Toward learning the binding affordances of objects: A behavior-grounded approach', in *Proceedings of AAAI Symposium on Developmental Robotics*, March 21-23, Stanford University, (Mar 21-23 2005).

[40] M. Tomasello, *First verbs: a case study of early grammatical development*, Comabridge University Press, New York, 1992.

[41] K. J. Vicente and J. Rasmussen, 'Ecological interface design: Theoretical foundations', *IEEE Transactions on Systems, Man, and Cybernetics*, **22**(4), 589–606, (1992).

[42] M. Viezzer, 'Autonomous concept formation: an architcture-based analysis'. Draft PhD. Thesis, 2003.

[43] W. H. Warren, 'Perceiving affordances: Visual guidance of stair-climbing', *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 683–703, (1984).

[44] S. R. Waxman and D. B. Markov, 'Words as invitaions to form categories: evidence from 12-months old infants', *Cognitive Psychology*, **29**, 257–302, (1995).

[45] N.J.E. Wijngaards, M. Kempen, A. Smit, and K. Nieuwenhuis, 'Towards sustained team effectiveness', in *Selected revised papers from the workshops on Agent, Norms and Institutions for Regulated Multi-Agent Systems (ANIREM) and Organizations and Organization Oriented Programming (OOOP) at AAMAS'05*, ed., G. Lindemann, volume 3913 of *LNCS*, 33 – 45, Springer-Verlag, Berlin, (2006). In press.

[46] A. L. Woodward and E. M Markman, 'Early word learning', in *Handbook of child psychology*, eds., W. Damon, D. Kuhn, and R. Siegler, volume 2, 371–420, Wiley, New York, (1997).