Logique & Analyse **206** (2009), 99–124

ARTEFACTS AND CATEGORIES

LIEVEN DECOCK

Abstract
This article analyzes the relation between artefacts and categories within the embodied cognition paradigm. An evolutionary account of human cognition, in which artefacts play a central role, is put forward. It is argued that artefact creation and categorisation are entwined and have co-evolved. Some ontological corollaries with regard to artefacts' creator's intentions, artefact functions, artefact kinds, and the kind "artefact" are discussed. Possible limitations on the scope of the view are dismissed.

The aim of this paper is to explore the role of artefacts in modern epistemology. Traditionally, artefacts have only received marginal attention in epistemology or cognitive science. Few theories of human cognition would assign an important role to artefacts. The epistemological relevance of artefacts is usually restricted to the design stage, while their material implementation is only an extraneous realisation without epistemic repercussions. Against this received (Aristotelian[1]) view, I will reflect on the role of artefacts from the perspective of embodied cognition theories. I will argue that this new perspective may lead to artefacts having a more central role in epistemology. In particular, I will focus on the relation between artefact creation and categorisation.

As a starting point, I will present the current state of the debate on artefact kinds and artefact categorisation. In the second section, I will put forward the embodied cognition approach in cognitive science, and explain how it can provide an account of categorisation. In the third section, I will reconsider the relation between artefacts and categories from this perspective. Artefact creation and categorisation mutually influence one another and coevolve over time. In the fourth section, I will briefly hint at the ontological implications of the new epistemological perspective, in particular with regard to artefacts' creator's intentions, artefact functions, artefact kinds, and

[1] The best exposition of Aristotle's views can be found in *Physics*, Book II.

the kind "artefact". In the last section, I discuss some possible limitations of the scope of the view presented here; further, I respond to objections that this view might not apply to abstract and innate categories or to technical artefact kinds.

1. *Artefact essentialism*

The issue of artefact kinds was taken up in the debate on natural kinds in the 1970s. In the context of a theory of possible worlds semantics, Kripke and Putnam put forward the idea that direct indexical reference to natural kinds is possible. The theory replaced the theory of descriptions, in which a kind was said to be characterised by a set of properties. Putnam elaborated this theory into a full semantic theory for normal languages. In "The Meaning of 'Meaning"' [1975] he argues that words such as "gold", "tiger", or "water" directly refer to kinds in the outside world. This reference is possible because these kinds are believed to be determined by an essence, which is an underlying structural feature. For the examples given, this feature is a chemical structure or a biological trait such as gene structure.

In his pivotal paper, Putnam indicated that the indexical theory of kinds could also be applied to artefacts. He gives the example of pencils:

> It follows that 'pencil' is not *synonymous* with any description — not even loosely synonymous with a loose description. When we use a word 'pencil,' we intend to refer to whatever has the same *nature* as the normal examples of the local pencils in the actual world. 'Pencil' is just as *indexical* as 'water' or 'gold'. [Putnam 1975, 243]

The obvious next question is what the essence of artefacts could be. Putnam remains silent on this issue; he believes artefacts have a hidden essence that remains to be discovered.[2] This view is not very satisfactory, and in a rejoinder, Stephen Schwartz argued that artefact kinds could not be natural kinds, since they have no essences:

> As it is, Putnam does not even give a hint as to what such an underlying nature might be in the case of pencils. I believe, of course, that there is no such underlying nature of pencils, nor is there a presumption of such nature. What makes something a pencil are superficial

[2] For Putnam this question is entirely open. He even countenances the possibility that this essence could be organic, see Putnam [1975, 242–244].

> characteristics such as a certain form and function. There is nothing underlying about these features. [Schwartz 1978, 571]

In a subsequent rejoinder, Hilary Kornblith [1980] inverted the argument, and thus ended the discussion.[3] He argued that there are various sorts of kinds, each determined by a particular essence; there are chemical essences, biological essences, and also artefact essences. For artefact kinds, Schwartz had already given an important hint, namely the characteristics form and function. Kornblith suggests that the essence of an artefact kind is the function of its artefacts. The final solution consisted in function essentialism.

For more than a decade, the issue was scarcely debated. Since the beginning of the nineties, however, there has been a renewed and growing interest in artefact kinds (or rather artefact categories) in experimental psychology, especially developmental psychology. Psychologists study the way people use categories, and more specifically, how infants learn to use categories. The outcome of several psychological experiments was that artefact categorisation is more problematic than one would expect. There are several explanations of categorisation, most of them involving some form of psychological essentialism. In contrast with Putnam's metaphysical realism, in which the essence of a natural kind is located in the outside world, psychological essentialism claims that categorisation, i.e. grouping objects in categories, is based on psychological essences, namely concepts [Medin & Ortony 1989]. Psychological essentialism claims that people believe that objects have essences, and that such essences provide a basis for their judgements of kind membership, or in other words, that people believe that objects belong to stable kinds, and that there are concepts determining these kinds.

Malt and Johnson [1992] undermined the most straightforward artefact essentialism, namely function essentialism.[4] They carried out several experiments in which they compared categorisation of artefacts on the basis of intended function with categorisation on the basis of physical shape. It became clear that identifying an artefact function is not sufficient for judging an artefact to belong to a kind. Even worse, the results indicate that physical features are more important than functional features. In their experiments, subjects were presented with a functional description of a boat ("manufactured and sold to carry one or more people over a body of water for the purposes of work or recreation"), and a deviant physical shape ("spherical and made of rubber, is hitched to a team of dolphins, and has a large suction cup

[3] In Denkel [1995] a similar argument is elaborated.

[4] See also Bloom [1996, 3–6] and Sloman & Malt [2003, 564–569] for further critique of function essentialism.

that can keep it in one place"), and were asked whether they would classify it as a boat. Most subjects would deny this. The inverse case, an appropriate physical shape ("wedge-shaped, with a sail, an anchor, and wooden sides") and a deviant functional description ("manufactured and sold as a holding area for dangerous criminals or persons in exile by detaining them a certain distance off-shore") was less problematic, and the described object was judged to be a boat. These and similar results [Hampton 1995, Landau 1994, Landau, Smith & Jones 1998] have discredited psychological function essentialism.

In response to these difficulties, Paul Bloom has put forward a new form of psychological essentialism [1996, 1998]. Based on Jerrold Levinson's ideas on works of art, he elaborated an intentional-historical theory of artefact concepts:

> We construe the extension of artifact kind X to be those entities that have been successfully created with the intention that they belong to the same kind as current and previous X's. [1996, 10]

Bloom's account of artefact categorisation is author-based. He believes that people are able to infer the original intention of the creator of the artefact, and that on the basis of this ascribed intention, artefacts are grouped into kinds. For example, a certain wooden object will be judged to be a chair, in case the subject believes that the carpenter has constructed the object with the intention it to be a chair. This theory is a form of essentialism, since it assumes that artefact creators start with the idea of creating a new object of a certain kind; i.e., the existence of artefact kinds is implied in the creation of artefacts. Hence, artefact categorisation is based on the recognition of the artefact's creator's intention.

This artefact essentialism has also been severely criticised. Sloman and Malt [2003, 569–572] indicate some problems with creator's intention essentialism. First, it is argued that the characterisation of essences is circular; the kind is determined by the creator's intention, while the (first) creator's intention implies the existence of the kind.[5] Furthermore, some empirical studies[6] indicate that if people are given privileged access to the creator's

[5] Thomasson [2003, 597] therefore proposes that this account should be slightly modified, so that "something is of artifactual kind K if and only if it is the product of a successful intention to create something of kind K, where the relevant intentions involve a substantive intensional concept of the nature of things of kind K rather than a transparent reference to an historical sample."

[6] Sloman and Malt [2003, 571] argue that Chaigneau's [2002] results refute claims to the contrary made by Gelman, Bloom, Keil, Matan, Carey and Rips.

intention, this does not necessary lead to less ambiguous categorisation. The actual categorisation of subjects deviates from the categorisation based on the creator's intention. At a more general level, the anti-essentialists argue that no context-independent categories exist.[7] The use of categories depends on the task at hand allowing for different ways of categorising. Analogously, they point out that artefact categorisation in a community can change over time.

The present situation in cognitive science is thus confused. Some psychologists try to explain artefact categorisation by invoking essences, more particularly, function essentialism and creator's intention essentialism, while others try hard to demonstrate that the proposed forms of essentialism fail to comply with empirical data, and argue that universal, a-temporal, context-independent categories are undesirable. At the moment, the debate is quite open, and this results in an increasing and refined empirical research into artefact categorisation in different disciplines, such as developmental psychology[8], cognitive ethology, cultural anthropology, and neurology.[9]

In recent years, the debate on artefact kinds has also been taken up anew in philosophy. Several authors have proposed alternatives to function essentialism. Risto Hilpinen defined the notion of artefact by means of authorship: "An object o is an artefact if and only if o has an author." [1993, 156] Hilpinen's position is akin to Bloom's view, and is also inspired by Levinson. In "Realism and Human Kinds", Amie Thomasson defends a similar view. She explicitly addresses the realism issue, namely whether artefacts kinds have a mind-independent existence. Since substantive intensional concepts are essential for the artefact kinds, the intentional-historical account of artefact kinds is less unproblematic for the realist than the functional account [2003, 603–604]. In "Artifacts and Human Concepts" [2007, 72–73], she

[7] This anti-essentialism in cognitive science is analogous to the anti-essentialist position in philosophy. In recent decades, philosophers have argued that empirical data in biology and chemistry contradict essentialist claims concerning natural kinds. Biological categorisation turns out not to be universal, but very context-dependent. In general, biologists avow that they use merely stipulative definitions (e.g. common ancestors, possibility of interbreeding, common genetic elements, etc.) in order to demarcate biological kinds, which are to an unexpected extent arbitrary (for a clear exposition, see Dupré [1995]). Similarly, it has been argued that chemical kinds, even H_2O, are less clear-cut than Putnam and Kripke assumed [see van Brakel 1986; 2000]. Malt [1994] has argued the same from a cognitive point of view.

[8] In recent studies, developmental changes in artefact categorisation are intensively studied. How artefact categorisation changes from say 4 year olds to 6 year olds to adults is investigated. Important changes have been found, see e.g. Matan & Carey [2000], Gutheil et al. [2004].

[9] Excellent overviews of recent results in these disciplines can be found in Margolis & Laurence [2007].

argues that complete mind-independence may be too narrow as a criterion of existence, and concludes that artefact kinds can be real kinds. Crawford Elder [1989, 2007] sticks to mind-independence, and offers a functional-historical account of copied kinds, of which artefact kinds are a subgroup. Copied kinds are characterised by an actual copied shape, a proper function, and a historically proper placement, and other properties may cluster around this combination [2007, 37]. However, Elder's critical notion of "proper function" is not uncontroversial. Various philosophers have tried to formulate a precise definition of an artefact's proper function, or have criticised such attempts [Preston 1998a; Millikan 1999; Vermaas & Houkes 2003; Preston 2003; Houkes & Vermaas 2004; many articles in Kroes et al. 2006]. As yet, there is no generally accepted notion of artefact function. More generally, there is no generally accepted philosophical notion of artefact kind.

2. *Categorisation in the embodied cognition paradigm*

In the 1970s, the debate on artefact kinds took place against the backdrop of the prevailing cognitivist-computational paradigm in cognitive science. The Putnam-Kripke theory of natural kinds, which was part of the causal theory of reference, became generally accepted, and concurred with other developments in cognitive science.[10] These cognitivist theories rather naturally lead to some form of essentialism, be it metaphysical essentialism or the weaker[11] psychological essentialism. In recent decades, however, we have witnessed several developments leading to alternative views on cognition and categorisation, such as the prototype and exemplar theory, decentralised or modular thought [e.g. Minsky 1986, 200*ff*], connectionism [e.g. Rogers & McClelland 2004], and most recently embodied cognition. Of course, such alternative views on the nature of cognition, and in particular categorisation, must lead to a different relation between artefacts and categories.

A full analysis of the repercussions of the various alternative paradigms in cognitive science on the relation between artefacts and categories would be commendable, but impossible within the scope of this article.[12] Rather, I

[10] Kornblith [2007] clearly highlights the relevance of this cognitive background in the discussion of artefact kinds. He contrasts the Putnam-Kripke causal, historical or direct theory of reference to the earlier description-based theory of reference.

[11] Psychological essentialism is closely related to the 'theory theory' of concepts, and is arguably closer to definitionism than the Putnam-Kripke theory; see Prinz [2002, ch. 4.1].

[12] Since the collapse of the "classical view", i.e. the Aristotelian definitional view, on concepts and categories in the 1970s, there is no generally accepted overarching theory now [Murphy 2002, 4].

will take some first steps towards a description of the relation between artefacts and categorisation within the embodied cognition paradigm, which is arguably the most radical competitor of traditional cognitivism that is firmly based on empirical research. But even for a single paradigm, it is not feasible to spell out all the details, so the best I can do is to put a new position on the map. Nevertheless, presenting this alternative view will be illuminating, because it casts a different light on various philosophical topics, and most notably on artefact essentialism.[13]

A paradigmatic example of embodied cognition[14] is Rodney Brooks's work in robotics.[15] In the mid-eighties, Brooks proposed a radically new approach in artificial intelligence and robotics. He side-stepped the traditional SMPA-approach in robotics, in which physical action and cognition were clearly separated. SMPA stands for "sense-model-plan-act." In building intelligent robots, the engineer uses a cognitive architecture with different modules operating consecutively. First, the robot uses its sensory devices to gather information about the environment. Subsequently, this information is processed into a complete model of the environment. In view of this information, and the task at hand, the robot calculates into all details a plan of action that is subsequently executed by activating the robot's motor devices. The traditional model clearly reflects traditional conceptions on cognition. There is a clear separation of the conceptual and the physical; all interesting work is done at the conceptual level by the central processor in the M- and P-phases; and there are clear bridges between the conceptual and the physical, upwards via the sensory apparatuses, and downwards via the activators of the motor devices. In Brooks's "subsumption architecture", however, action and cognition are merged, and there is no longer a clear separation between cognitive and physical modules in the robot, nor is there a separation between the robot's cognitive and active periods.

Brooks's approach was a real challenge for the traditional approach in artificial intelligence. The new ideas led to surprisingly successful results in

[13] In this paper, I especially focus on artefact essentialism, but one may also consider topics such as the normativity of artefacts, non-intended uses of artefacts, the dual nature of artefacts, the relation between artefact kinds and biological kinds, etc. from the perspective here presented. A full elaboration of such themes is relegated to further articles.

[14] Other terms are "situated cognition" [Clancey 1997], "existential cognition" [McClamrock 1995], ecological cognition [referring back to J.J. Gibson's work, see Gibson 1979], or in artificial intelligence "EAI", which stands for embodied artificial intelligence [Chrisley 2003].

[15] For a larger exposition of both the technical details and the philosophical ideas, see Brooks [1999], a collection of earlier articles, including the influential 'Intelligence without reason'.

robot navigation. Navigation had vexed roboticists for years, and it seemed close to impossible to program a robot so that it can move in a real world situation in real time. Brooks managed to build robots that can perform simple tasks while smoothly moving in an arbitrary environment, while avoiding other objects and persons standing and moving there. This technical progress proved decisive in gaining credibility for his innovative ideas on the nature of cognition.

The new approach is based on a few basic ideas that are closely related. First, cognition is believed to be "situated" in an environment. This means that cognitive processes are essentially determined by the environment in which they are located, and implies that a cognitive process would be altered if relocated in a different environment. There is no separate realm of cognition, but cognition is a process in the real outside world, and is determined by its place in the outside world. A second tenet is the "embodiment" of cognition. It is believed that cognition requires a physical body, and that the make-up, history, and deficiencies of this body are essential constituents of the cognitive processes. Changes in the body of the cognitive agent would result in changes in the cognitive processes. A third basic idea concerns active cognition. The boundary between cognition and action, typical for the SMPA-approach, is wiped out, and cognition thus becomes a special form of bodily action in a physical world. Cognitive processes all involve real-time interaction with the environment. A fourth ingredient is the demise of central processing. The traditional view on cognition was largely based on the computer metaphor; a central processor computes information available from internal memory and external input, and subsequently presents output. In embodied cognition, it is no longer believed that there is an area in the brain where the central processing takes place. Instead, there are several modules that co-operate with a minimum of co-ordination. A fifth element is minimal representation. While the traditional view on cognition assumes a large conceptual or mental realm with a vast stock of representations of the real world, i.e. ideas, concepts, internal images, internal sounds, memories, etc., in embodied cognition theories internal storage is believed to be quite limited, and the representations are believed to contain much less information about the environment than usually taken for granted. It is readily seen that this approach has originated in cognitive theories of navigation. This work also helped to revitalise the ecological approach to perception,[16] and was later extended to other domains of cognition, even to memory.[17] Of

[16] The ecological view on perception was first proposed by J.J. Gibson [1979], and has become influential again through the work of Varela [et al. 1991], and more recently O'Regan and Noë [2001], Noë [2005].

[17] Memory is not only the internal storage of information, but that it also involves external storage in the real world, see e.g. Rowlands [1999, ch. 6], or Donald [1991, ch. 8].

more immediate importance, in combination with connectionist networks, embodied cognition can also provide a new outlook on categorisation.

In a first step, categorisation can be studied in a connectionist setting. In the mid 1980s, scientists in various disciplines realised that the computational or cognitive conception of categorisation faced immense problems. Recognition of objects or written signs in the visual field or of phonemes in the auditory module was hardly computable. The invention of the so-called backward propagation algorithm for neural networks, i.e. networks consisting of parallel coupled processors, connected by means of links of varying strength, was far more successful. It became possible to train networks for categorisation tasks that were hitherto only feasible for human beings. Moreover, the use of artificial neural networks had the further advantage that it could also explain human categorisation. The idea of neural networks was roughly based on the working of the brain. Neurones are connected through axons and dendrites with varying strengths. One could thus reasonably believe that categorisation in the brain and in artificial neural networks is structurally the same.

However, this connectionist account of categorisation, which was a major breakthrough, is still not the full picture, since it is too much wedded to traditional ideas of categorisation. It is generally assumed that there is a pre-existing set of categories, and networks are then trained to classify sensory input into one or more of these categories.[18] Typically, networks are "supervised" in the training period — a network is provided with sensory input, tries to give the appropriate category, and after being corrected reshuffles the strengths of the connections. Therefore, connectionist artificial neural networks are not autonomous. Of course, this situation resembles infant learning, in which feedback by adults is also extremely important for categorisation. However, the account leaves no room for the evolution of categories, and for the creation of new categories. Moreover, in the connectionist account of categorisation, categorisation is an independent faculty. This is not necessarily problematic, if one countenances a modular view on the mind and the human brain. However, the combination of connectionism and embodied cognition yields another picture on mental architecture.[19]

Categorisation is not a part of the "model"-step in the SMPA cycle, but is integrated in the global behaviour of an agent. Avowedly, a neural network or a brain is necessary for categorisation tasks, but this does not mean that

[18] For a more elaborate theoretical exposition of the traditional account, see Barsalou [1992], and for a practical application for the case of object recognition, see Ullman [1996].

[19] The here presented account of categorisation in embodied cognition is largely based on Scheier & Pfeifer [1995]; Lambrinos & Scheier [1996]; Scheier & Lambrinos [1996a; 1996b]; Pfeifer & Scheier [1999, ch. 12]; Beer [2003]; Nolfi [2005], Poirier et al. [2005].

one can easily locate clear-cut categories in these neural nets. Rather, categorisation is a process that is spread over the neural network, the physical agent, and the environment. In recent years, various artificial agents have been developed in which categorisation is an "emergent"[20] phenomenon, tightly related to sensorimotor coordination.[21] Instead of relying on sensory stimuli only, the agents actively use their physical movements. These agents were developed so that the implementation of an internal set of categories or the external control of category learning could be avoided. For example, the robots SMC I and SMC II were built to collect bins and bring them to a dustbin. Some bins were too large for the robot to transport, while others could easily be brought to the dustbin. After a few rounds through the environment, the robot could be seen to move towards the objects, and as it came closer to the bins it avoided the large bins, and went directly to the others to pick them up. The robots were seen to develop categories related to the task at hand, and adapted to the environment in which they were located.

In conclusion, the embodied cognition framework provides a new account of categorisation that is in important ways different from more traditional views. Categorisation is tightly related to embodiment; it is a process that is dependent on the physical make-up of the agent because of its relation to sensori-motor control. It can be ascribed to an agent, but is not easy to locate within the agent, or defined in terms of physical processes. In other words, there is no clear separation between categorisation and other forms of (cognitive) behaviour. Categorisation is also situated; it is highly dependent on the environment in which it takes place.

Of course, it may be doubted whether this theory of categorisation, which is based in robotics, is also suited to model human cognition. Admittedly, future research in psychology and neurology may refute the theory, but at present this is not the case. On the contrary, there are several characteristics of human cognition that lend credibility to the model. First, the human brain has the structure of an immense neural network, so this accords well with connectionism. Second, it seems reasonable to describe the activity

[20] One should be careful with the use of the word "emergence". The notion is central in contemporary metaphysics of mind, e.g. in the work of Kim, Jackson, and many others. In cognitive science, however, the term is especially used in the embodied cognition literature to describe the spontaneous creation of simple behaviour or processes in an initially quite complex situation. For a broader overview, see Clark [2001].

[21] This fits with other empirical findings on categorisation in other disciplines, such as neurobiology [Edelman 1987, 210: "perceptual categorization depends upon the interplay between cortical sensory maps and local motor maps"], developmental psychology (Ruff [1984] stresses the importance of object rotation in human infants), or comparative biology and archaeology (e.g. Donald [1991, 147] speculates on the relation between hand control and the growth of cognition in hominids).

of human beings in models that resemble Brooks's subsumption architecture. Navigation and action in an environment are then the basic actions, and higher cognitive/behavioural functions are superposed on them. As a working hypothesis, I will assume that the here presented view is an adequate description of human categorisation. In the remainder of the paper, I will take the analysis a step further, and relate embodied categorisation to artefact creation.

3. *The co-evolution of artefacts and categories*

In the embodied cognition literature, artefacts are seldom discussed. This is not surprising, since one of the basic rules is real-time implementation in physical robots, and at present, we are nowhere near the construction of robots that are able to create artefacts.[22] Nevertheless, it is worthwhile to extend the embodied cognition perspective to the creation of artefacts, and to reflect on the role of artefacts in human cognition.[23] I will argue that categorisation and the creation of artefacts are essentially entwined. Artefacts thus play a crucial role in the development of human cognition.

It is obvious that categories play a role in the creation of artefacts, both in traditional cognitive theories and in embodied cognition. First, the construction of artefacts fits perfectly in the SMPA cycle, or in a loop of such cycles. Especially the plan and act steps are important. This is in accordance with the widespread conviction that artefacts are the material realisations of plans conceived by their creators. Moreover, the categories that are available to the artefact designer are essential in devising a plan. Devising artefacts involves various categories, at least categories of physical parts, function categories, and categories of artefact kinds. Second, embodied cognition theories will also highlight the role of categories in artefact creation. These categories

[22] In the literature on artificial agents, it is sometimes indicated that situated agents not only navigate in an environment, but may actually modify the environment for certain tasks, see e.g. Beer [2003, 211]: "[...] a situated agent can utilize and manipulate the physical and functional organization of the space around it, as well as the social organization in which it exists, to offload problems to its environment." It is thus suggested that there is a strong agent-environment coupling, whence it is a small step to argue that the agents create an 'artificial' environment.

[23] In a remarkable paper, Preston [1998b] discusses the role of artefacts in philosophy starting from a Heideggerian perspective, and reaches conclusions that are quite similar. In view of the similarities between embodied cognition and some forms of phenomenology, this is not too surprising. For a more elaborate contemporary Heideggerian perspective on artefacts and technology, see Verbeek [2005].

differ from the categories in traditional cognitive theories; they are not really clear-cut and fixed, and they are related to the manipulative behaviour or plans[24] of the creator. Nevertheless, any sensible description of the construction of artefacts will invoke categories.

The major difference between traditional cognitive theories and embodied cognition theories is that the latter also involve an inverse relation, while the former do not. In the former, there is a strict separation and hence a clear order between planning and acting. Planning precedes the construction of artefacts, and the actual artefact creation has no obvious feedback loops to the cognitive framework in which the planning takes place. More specifically, building artefacts has no influence on the categories used in the planning.

In the embodied cognition view, however, the relation between artefacts and categories is different. The construction of new artefacts or the modification of old artefacts have subtle but important repercussions on categorisation. In the embodied cognition view, categorisation is related to action, and dependent on the environment in which the agent is located. This environment is always changing, and thus categories change through the gradual change in the neural connections that underlie the categories. A most important characteristic of human beings is that they actively modify their environment, especially when building artefacts. The result is that categories constantly change through the introduction of new artefacts and the change in existing artefacts. For example, one could say that the Bauhaus revolution in cutlery gave us new categories of "knife", "spoon", etc. These objects had a slightly different form after this cultural innovation, and also the sensorimotor handling of these artefacts had slightly shifted.[25] The result is a rewiring of some neural connections,[26] which will result in a slightly altered

[24] Some forms of planning, i.e. cognitive processing that is not in real-time, are compatible with the embodied cognition paradigm. Some authors countenance processes such as emulation and simulation, which are representational, but still thoroughly grounded in embodied, situated action. See e.g. Poirier et al. [2005, 752*ff*].

[25] A more didactical example might be Petroski's description of the historical development of forks and the related sensorimotor changes in western eating habits [1994]. I prefer the given example because it highlights the continuous subtle change in categorisation. Less fine-grained examples are too often explained in terms of discrete "meaning changes" or "category changes", in order to avoid the account of a continuous gradual change in categorisation.

[26] The here presented view is inspired by connectionist theories of categorisation [e.g. Rogers & McClelland 2004]. It is possible to construct networks in which categories are not clearly located but spread over a multitude of nodes. Giving input or feedback to a network may lead to small adjustments in some connection weights, so that the overall categorical structure, which is determined by the connection weights, slightly changes. However, one ought not to neglect the embodiment of the network, see Pfeifer & Scheier [1999, 433]: "The

categorisation of cutlery.[27] Furthermore, the modification or creation of artefacts leads to changes in other categories. The use of Bauhaus cutlery at the beginning of the previous century led to a slightly shifted sensorimotor experience of activities such as eating or cutting, and thus — it is an implication of the present view — to alterations in the verbal categories "to eat" or "to cut". It is not immediately clear how far these rather small changes will reverberate through the neural network,[28] and how they will affect the overall activity of an agent in a changed environment, yet it is inevitable that such changes will in fact occur. There is no independent cognitive framework of categories in which artefact design takes place; the construction of artefacts leads to changes in the environment, and thus to a change in the categories the creator uses in his environment.

Moreover, one should not assume that shifts in categorisation behaviour are only due to neuronal reorganisation. By stressing coordination dynamics and hybrid representational forms, there is room for complementarity between biological and artefactual cognitive contributions. [Clark 2006, 300]. On this view, artefact categorisation is a complex hybrid interplay between material objects in the external world and the internal biological processes. Changes in categorical behaviour may result from material changes in the outside world (new design processes, the creation of new artefacts, etc.) and from neuronal reorganisation of (thin, sensorimotor) representations. Both changes can interlock, and lead to looping processes in which material culture, the neuronal organisation, and hence the categorisation behaviour, radically change over a period of time.

categories in our robots are represented in a distributed way in the weights of certain connections in the neural networks, ... These weights make sense, though, only if they are embedded in a real-world physical agent. Take the same network and put it into different agents: The connections then mean something entirely different. If you take the same network initially and put it into a different agent, the new agent will acquire very different categories. Thus once more, we see that categorisation is not what is done by a particular module, but is a property of the complete agent."

[27] It is noteworthy that the embodied cognition view does not equate categories with their extension, so that these gradual changes do not necessarily lead to changes in the classification of existing artefacts into kinds.

[28] I am not committed to precise models of the neural organisation of categorisation behaviour. The view is compatible with highly structured models of categories, based on taxonomic hierarchies [Quillian 1968]. Rogers & McClelland [2004, Ch. 3] have demonstrated how hierarchies can be latent in distributed networks. Minor local changes in one category then leave the structure almost intact, and only affect categories nearby in the taxonomic structure. Nevertheless, it is important to note that these taxonomic hierarchies are not intrinsic, and can change in learning processes.

The embodied cognition perspective thus leads to a close relation between artefacts and categories. Artefacts and categories have a mutual influence, and in a historical perspective they can be seen to co-evolve. The construction of artefacts has had a decisive impact on the environment, in which human beings live, which leads to quite particular ways of categorising this 'artificial' environment. On the other hand, categories are crucial in human cognition and hence in the design of artefacts. Therefore, it is unsurprising that artefact categories can change dramatically over time, though in a series of small steps. For example, the category "mill" is now used in a different way than in the Middle Ages. Now one thinks in the first place of windmills that provide energy in the form of electricity. In the Middle Ages, they were used to turn grain into flour, or in more watery region such as Holland to drain superfluous water to the sea. Historians may thus reconstruct the parallel development of the category "mill" and the actual mills that have been used over the last centuries.

One can even argue that this co-evolutionary account of the relation between artefacts and categories is already an overidealisation. Namely, it assumes a rather clear separation between artefact creation and categorisation, but, on the assumptions made above, there is no clear separation between action and cognition. Therefore, artefact creation and cognition (and in particular categorisation[29]) are two sides of a single coin. When one looks at someone's creative activity, one can describe this behaviour more materialistically as artefact creation, or one can focus on the cognitive perspective, and thus describe the activity as involving plans and designs. In principle, a materialist account of the interplay of an agent, his neural wiring, and the environment could even suffice as a complete description of the cognitive activities, although the "emergent" cognitive phenomena are a more appropriate description level of what is happening. This possibility highlights the strong link between material and cognitive activity.

This rather materialist and evolutionary account of cognition, and more specifically artefact creation and categorisation, ties in with contemporary theories of the evolution of human cognition, especially in anthropology. There is still a large fog over the origins of human cognition, but several tentative attempts to address this question render plausible the claim that cognition and the use of tools are tightly related in an evolutionary progress. There is a lot of investigation into the development of tool use from primates until *Homo sapiens sapiens* [see Gibson & Ingold 1993; Mithen 2007]. In the 1920s, Wolfgang Kohler already demonstrated that some primates are able to use sticks to obtain food; furthermore, it has been shown that some primates use sticks to catch termites, or that they can combine tools in a

[29] Harnad [2005] even defends the claim that cognition *is* categorisation.

hammer and anvil fashion; in the Oldowan culture (Tanzania, 2.500.000 BC–1.500.000 BC) primitive stone tools are found, namely sharpened pebbles; in the Acheulian culture of *Homo erectus* one finds more sophisticated polyfunctional handaxes and choppers that were in use for more than a million years; there is an outspoken refinement in tool manufacture with the advent of *Homo sapiens* around 100.000 years ago, which led to a spectacular increase in tool use in the Mesolithic and Neolithic period. This development in tool use goes together with cognitive developments, such as increased encephalisation, lateralisation in the brain, more complex social behaviour, increased visual categorisation,[30] an improved episodic memory, mimetic gestures, and the development of a vocal tract, which led to the development of language. This evolution is rather haphazard and jumpy, with large periods of stability and sudden changes. It is most likely that this evolution was a tangled web in which the several developments interacted. Of more immediate interest, it is reasonable to believe that the material and cognitive developments did not occur independently.

Even if we are left in the dark on the origin of artefacts and cognition, we can learn a lot from subsequent developments. There are enough archaeological and historical data on the development of many common artefacts we use everyday. Even radical novelties are less a result of planning, than the result of modification of an existing environment. In the case of architecture, it wasn't the case that at the end of the latest Ice Age, some architect devised a bright plan to build a house, collected the materials and thus built the first house in history. On closer scrutiny,[31] one sees that humans started to hide in caves, to collect material in the caves, to make fire in them, and to provide the walls with painted figures. Later, they started building movable sheds, such as tents, subsequently started building temporary residences such as earthen houses, igloos, wooden sheds (in response to environmental conditions), and later started to divide them into rooms, or started living permanently in them. At no point in this development is there a clear planning in an architectural or technological sense. Building houses has originated from a series of small modifications of the environment, and a gradual improvement of the already 'artificial' environment. It is often very hard to judge to what extent these modifications are really planned or just emerge from contingent variations in

[30] See e.g. Donald [1991, 153–157] on event perception. Donald compares the visual abilities of humans and other animals, especially primates.

[31] Mithen [2004] offers a readable overview of the evolution of mankind including its cognitive development at the start of the material or technological culture at the end of the last Ice Age. One finds a varied account of the development of housing conditions in the period of 20.000 B.C. until 5.000 B.C. in the various continents. If the reader would not agree that houses are artefacts, one could repeat the example for the case of pottery.

the actual construction. From the perspective of embodied cognition, there is no real answer to this question. Cognition and action are essentially entwined, also in the construction of artefacts. One can conclude that there is little difference between thinking and tinkering. Archaeological evidence on the co-evolution of artefacts and categories is easily encompassed within the embodied cognition framework.[32]

4. *The ontology of artefacts*

Hitherto, artefacts have been discussed from an epistemological point of view; the place of artefacts in modern theories of cognition has been sketched. The here presented account also affects the ontological status of artefacts. I will briefly sketch some ontological consequences of the embodied cognition perspective, namely for creators' intentions, functions of artefacts, artefact kinds, and the kind "artefact." I will heavily draw on Pfeifer and Scheier's discussion of the frame-of-reference problem [2001, 112–117; 387].[33] This 'problem'[34] in robotics is based on the different perspective of an outside observer and the perspective within an artificial cognitive agent.[35] The observer can attribute several cognitive properties to the emergent behaviour of a robot in a certain environment, but these properties cannot be identified at the physical or mechanical level. For example, one can observe that a robot "avoids objects" or "follows walls"; but this is only clear to the observer, and nowhere implemented or programmed in the robot. The cognitive features only emerge through the interaction of the robot and its environment. Moreover, there is no reason to believe that in these cases of emergent behaviour suddenly new ontological realms come into existence. One only has a new

[32] This is of course not to say that the archaeological evidence can only be explained within the embodied cognition paradigm. Mithen has elaborated a view on the origins of human cognition [1996] and artefact creation [2007] that heavily draws on Fodor's modularity thesis.

[33] A more rigorous account goes beyond the scope of this paper. The implicit metaphysical assumptions and the relation between ontology and epistemology should be made more explicit. Roughly speaking, the account presented here is based on the Quinean view that ontology is the handmaiden of epistemology; it is epistemology-driven ontology. Hence, philosophers who believe in a strict separation of ontology and epistemology, could still believe in the existence of 'metaphysical' artefact kinds.

[34] Strictly speaking it is not a genuine problem, but rather an ambiguity. This "frame-of-reference problem" should not be confused with the notorious "frame problem" in A.I.

[35] Nolfi [2005, 876] uses the terms 'distal description of behavior' for the observer's point of view, and 'proximal description of behavior' for the point of view of the agent's sensorimotor system.

level of description, which is moreover related to specific interests of an external observer. I will extend the frame-of-reference problem to the more general case of human cognitive agents designing artefacts.

a. *Creator's intentions*
The intention of the artefact creator exists only for the outside observer.[36] At the physical level, one has only the interaction of an agent and its environment. The observer can take an "intentional stance",[37] and thus attribute intentions to a creator of artefacts. However, there is no ontological realm to which these intentions belong. The intentions are relative to the specific interests of the observer.

b. *Artefact functions*
Artefact functions had a clear role in the cognitivist account.[38] The cognitive agent has a clear model of its environment, certain goals, and on this basis develops a plan for further action. The function of an artefact is related to the expediency of this artefact to the ultimate goal of the plan, and artefacts are developed with this purpose. However, if one merges cognition and action, and decentralises cognitive processes, one can attribute planning behaviour to certain agents, but there is no longer a central plan in which the agent and his artefacts have a precise role. In embodied cognition theories, the cognitive agent carries out various tasks at the same moment in real-time response to the environment, so that he is in various ways interacting with artefacts at a single instance of time. This also implies that artefacts can assume different roles in human behaviour, so that there are no unique clear-cut proper artefact functions.[39]

[36] In principle, the observer could be the creator. The underlying physical processes that determine the interaction of the creator and its environment are to a large extent opaque to the creator. Embodied cognition does not exclude self-reference, but this discussion goes beyond the scope of this paper.

[37] This term has been put forward by Daniel Dennett [1987], and is also the most discussed emergent property that can be ascribed at the observer's level in the embodied cognition literature. Dennett [1990] explicitly discusses creator's intentions in the case of artefacts.

[38] The transparency of artefact functions is already reduced as a result of a philosophical comparison of etiological theories of artefact and biological functions, see e.g. Preston [1998a]; Millikan [1999]; Houkes & Vermaas [2003]; Preston [2003]; Lewens [2004]. It would be most interesting to unravel the interplay between the biological, cognitive and technological evolutions, but this analysis cannot be encompassed within the scope of this paper.

[39] It is worth noting that the first human artefacts were polyfunctional. For more than a million years, the only human artefact was the handaxe. Archaeologists are still puzzled about the various ways in which handaxes were used, and do not understand why it took so long before other stone tools (or wooden or bone tools) were found [Mithen 1996].

c. *Artefact kinds*
For two reasons, the account presented here leads to an anti-essentialist account of artefact kinds. First, the two dominant proposals of artefact essentialism are function essentialism and creator's intention essentialism. As we have argued, neither creator's intentions nor artefact functions can be adequately pinned down. Since the alleged essences only appear in the observer's perspective, relative to his place and interests in the environment, they can only determine kinds relative to the description of an observer.[40] Second, the theory of categorisation presented here does not lead to a theory of cognition in which kinds have a natural place. From an epistemic point of view, we can suffice with categories determining classes of objects in our cognitive framework. We need not, as in traditional theories of cognition, invoke essences that determine (natural,[41] artefact, or psychological) kinds. But in view of the account of categorisation we have given, categories only appear at the observer's level, and are relative to the observer's interests. Therefore, there is no context-independent use of artefact categories.

d. *The kind "artefact"*
A related problem is the determination of the kind, or rather category "artefact". In the above account, I have used the term 'artefact' rather casually, as if the demarcation of artefacts from other physical objects poses no special problem. However, on deeper reflection, this imprecision is not warranted. Apart from the overall context-problem with categories, there are more specific problems with this special category. In the embodied cognition framework, the boundary between action and cognition is blurred, and artefact use and artefact constructions are basically the observer's description of the active modification of the environment by a cognitive agent. No clearer characterisation of artefacts is readily at hand.[42] Since there is only a

[40] See e.g. Pfeifer & Scheier [1999, 433]: "We made a strong point of saying that the categories the robot acquires are, in fact, in the observer's head rather than in the agent's."

[41] As indicated in footnote 7, there is reason to doubt that natural kinds exist; after all, embodied categorisation implies context-dependence. Nevertheless, the case of natural kinds may be more intricate. It may be the case that some features of a fixed environment are especially salient so that they nearly all cognitive agents in a wide variety of cognitive tasks are responsive to these features. For artefact kinds on the hand, such *permanent* salient features are unlikely on the present account, because cognitive agents change their environment continuously. However this may be, the discussion of natural kinds requires further elaboration that goes beyond the scope of this paper.

[42] This is not to say that every object is an artefact. Objects or parts of the environment that have not in any way been modified by cognitive agents (e.g. the planet Saturn) are not artefacts.

gradual difference between instinctive, random, and planned behaviour (for the observer), there is no principled demarcation between the instinctive, random, or planned use and construction of artefacts. In other words, there is no principled reason to exclude bird's nests or Jackson Pollock paintings from the category "artefact". Moreover, it is unclear which modifications of the environment can count as artefacts, and which not. There seems to be no *principled* difference in modifying living organisms and modifying metallic ore. In other words, there seems to be no principled way of excluding horses, seedless grapes, or the OncoMouse from the category "artefact". There are no principled limits on the size, duration, or tangibility of the part of the environment one is modifying, so there is no real reason to exclude streets, Central Park, Holland, perfumes, and songs from the category "artefact".

5. *Some possible limitations*

Limitation 1. In the previous sections, the role of artefacts in Cambrian intelligence (i.e. the insectlike intelligence in modern robotics [Brooks 1999]) and Pleistocene intelligence (i.e. the intelligence of early humans) has been highlighted. A common complaint about embodied cognition has been that it is not well suited as an account of higher cognitive processes. One could thus argue that the here presented view is not likely to explain the origin and use of abstract categories (or concepts).
Reply: Andy Clark's influential work provides the essential elements for at least one forceful rejoinder. Clark argues that in linguistics, instead of elaborating on grammar and semantics, one can focus on symbolic notations. He describes language as "the ultimate artefact"; the symbolic notations are regarded as material artefacts useful in cognitive activity [1997, ch. 10; 2006]. He suggests that even mathematical reasoning is based in mathematical symbol manipulation [2001, 146; see also De Cruz forthc.], and speculates that the typically human higher cognitive functions such as self-evaluation and self-criticism, i.e. the second-order cognitive functions, may critically depend on public language [1997, 208]. In brief, it is a bit of an understatement to state that a strong coupling between the use of material symbolic artefacts and abstract categories can be accommodated within the embodied cognition perspective.[43]

[43] In a sense, the view presented in this article is an extension of the extended mind hypothesis [Clark & Chalmers 1998, Clark 2006]. Whereas the extended mind hypothesis primarily focuses on the complementarity of biological representations on the one hand and material symbols, or "sociocultural artifacts, including gestures, diagrams, external text, software applications, and more" [Clark 2006, 299] on the other hand, the view here put forward highlights the complementarity of biological representations and non-linguistic, non-symbolic

Limitation 2. One might argue that some categories are so basic that they are not affected by any action in the outside world. Immutable, innate categories will not take part in the coupling here described. Moreover, *prima facie*, there is evidence that at least some categories are indeed innate. For example, it is widely believed that Berlin & Kay [1969] demonstrated that colour categories are universal, whence it is a small step to the belief that they are innate.
Reply: Since colour is the ubiquitous example in this context, let me elaborate on colour categorisation. In the last decade, the Berlin & Kay thesis has come under attack [see e.g. Saunders & van Brakel 1997; 2002 and references therein, especially Roberson]. Instead of viewing colours as innate categories, it is possible to develop a view in which colour categories are related to material dyes.[44] Saunders and van Brakel [2002] have described colour as an "exosomatic organ". By this they mean, that the colour is a tool that extends the powers of the human body.[45] Colour vision and colour categorisation are thus based on a coupling of the human body and the (technologically) transformed environment:

> What explains the pragmatic success of colour science? One answer is that the environment is remade by paints and dyes, reproductive technologies, coloured lights, advertisements, fashion, television, cinema, computer screens and so on, in such a way that it reproduces the theory of colour that provides the techno-material base for the production of the coloured world. [2002, 342]

In brief, it is possible to develop a plausible view in which colour categorisation and artefact creation are tightly related. In principle, the coupling between artefact creation and categorisation can be so strong as to affect even the most basic categories.

artefacts in a particular, but crucial, cognitive activity, namely categorisation. Clark is congenial to giving a important role to 'ordinary' artefacts in human cognition [2006, 292; 300; private communication].

[44] For a popularising, but very illuminating history of the relation between colour words and material dyes, see Finlay [2002].

[45] For the present purposes, I am not interested in the body-environment boundary [for related views see Brooks 2002; Clark 2003]. What is relevant is that both the cognitive agent and its environment are important in categorisation, and less where the exact boundary between the two lies.

Limitation 3. One might object that the here presented view is not well suited to deal with contemporary technical artefacts.[46] Even if one would believe that there is a strong coupling between categorisation and manipulation of the environment, it can hardly be denied that artefacts such as satellites, coffee machines or scissors[47] do make up rather precise kinds and have been designed with a particular intended function.
Reply: First, at a larger time-scale, the co-evolution of categories and artefacts has been illustrated by means of historical analyses of the development of modern artefacts [see e.g. Basalla 1988, Petroski 1994, Ziman 2000]. In other words, the rather precise delineation of technological artefact categories is dependent on a background structure of categories that remains stable only within a limited time-span. Second, on a small time-scale, the categorical structure can be explained as resulting from social factors.[48] The social division of labour may explain the gaps between users' functions and producers' functions, or between the design, production, and testing stages. The normativity of artefacts [Hilpinen 2004, ğ3; Fransen 2006], i.e. the distinction between appropriate functioning and malfunctioning, may originate in social regulations. The uniformity of artefacts within a certain artefact kind is mostly a result of standardisation procedures.[49] These social factors need not be incompatible with the view here presented. In this article I have mainly focused on the relation between artefacts and categorisation from the perspective of an individual cognitive agent. Adding social factors may be a sensible extension of the view presented here,[50] but would largely exceed

[46] For various philosophical issues concerning technical artefacts, see Kroes & Meijers [2006].

[47] For present purposes, a "technical artefact" can be any artefact designed and produced in a technological, industrial process.

[48] This strategy would narrow the gap between 'artefact kinds' and 'social kinds'. In Thomasson [2003], these kinds are treated separately.

[49] Elder [2007, 46*ff*] argues that the category of Eames desk chairs, 1957 design, is a more plausible candidate as an artefact kind than the general category of chairs. It would seem that for Elder a standardised process of copying a certain model is a crucial characteristic of an artefact kind.

[50] In my discussion, I have mainly focussed on individual agents. It is natural to extend this methodological individualism to social contexts, with many interacting individuals mutually influencing one another's behaviour. Proponents of embodied cognition often stress the social dimension of cognition [e.g. Hutchins 1995]. Of particular interest is the theory of "stigmergy", which describes how individuals can affect the behaviour of others and themselves through modification of the external environment or artefacts, see e.g. Susi & Ziemke [2001]; Marsh & Onof [2007].

the scope of this article. Hence, there is no reason to believe that my general view is not applicable to modern technological artefacts. Nonetheless, it is reasonable to believe that the view can be complemented by invoking social factors.

In conclusion, there are no immediate limitations on the account of the co-evolution and interplay of artefacts and categories. I thus hope to have demonstrated the viability, soundness, and plausibility of a cognitive theory in which artefact technology and human cognition are strongly and inextricably entwined. At the present stage, this is the best that can be hoped for, because an argumentation for the correctness or preference of the view would require a considerable amount of extra empirical evidence. Nonetheless, the multifarious philosophical ramifications are sufficient justification for a further exploration of the view.

ACKNOWLEDGEMENTS

I thank the participants of the conferences *Philosophy of Artefacts* in Delft and *Extended Cognition* in Antwerp, and the members of the research group 'Ontology and Epistemology' at the Vrije Universiteit Amsterdam, for discussions on earlier versions. I especially thank Beth Preston, Pieter Vermaas, Jaap van Brakel and Igor Douven for their helpful comments on the final drafts.

Faculty of Philosophy
Vrije Universiteit Amsterdam
De Boelelaan 1105, 1081 HV Amsterdam
E-mail: LB.Decock@ph.vu.nl

REFERENCES

Barsalou, L.W. (1992). *Cognitive Psychology. An Overview for Cognitive Scientists.* Hillsdale: Lawrence Erlbaum.

Basalla, G. (1988). *The Evolution of Technology.* Cambridge: Cambridge University Press.

Beer, R. (2003). 'The Dynamics of Active Categorical Perception in an Evolved Model Agent'. *Adaptive Behavior* 11, 209–243.

Berlin, B. & Kay, P. (1969). *Basic Color Terms.* Berkeley: University of California Press.

Bloom, P. (1996). 'Intention, History, and Artifact Concepts'. *Cognition* 60, 1–29.

Bloom, P. (1998). 'Theories of Artifact Categorization'. *Cognition* 66, 87–93.
Brooks, R. (1999). *Cambrian Intelligence.* Cambridge MA: MIT Press.
Brooks, R. (2002). *Flesh and Machines.* New York: Pantheon Books.
Chaigneau, S.E. (2002). *Studies in the Conceptual Structure of Object Function.* (Ph.D. Thesis, Department of Psychology, Emory University, Georgia)
Clark, A. (1997). *Being There.* Cambridge MA: MIT Press.
Clark, A. (2001). *Mindware.* Oxford: Oxford University Press.
Clark, A. (2003). *Natural-Born Cyborgs.* Oxford: Oxford University Press.
Clark, A. (2006). 'Material Symbols'. *Philosophical Psychology* 19, 291–307.
Clark, A. & Chalmers, D. (1998). 'The Extended Mind'. *Analysis* 58, 7–19.
Chrisley, R. (2003). 'Embodied Artificial Intelligence'. *Artificial Intelligence* 149, 131–150.
Clancy, W.J. (1997). *Situated Cognition.* Cambridge: Cambridge University Press.
De Cruz, H. (forthc.). 'An Extended Mind Perspective on Natural Number Representation', *Philosophical Psychology*.
Denkel, A. (1995). 'Artifacts and Constituents'. *Philosophy and Phenomenological Research* 55, 311–322.
Dennett, D. (1987). *The Intentional Stance.* Cambridge MA: Bradford.
Dennett, D. (1990). 'The Interpretation of Texts, People and Other Artifacts'. *Philosophy and Phenomenological Research* 50, 177–194.
Donald, M. (1991). *Origins of the Modern Mind.* Cambridge MA: Harvard University Press.
Dupré, J. (1995). *The Disorder of Things.* Cambridge MA: Harvard University Press.
Edelman, G.E. (1987). *Neural Darwinism.* New York: Basic Books.
Elder, C. (1989). 'Realism, Naturalism, and Culturally Generated Kinds'. *Philosophical Quarterly* 39, 425–444.
Elder, C. (2007). 'On the Place of Artefacts in Ontology'. In E. Margolis. & S. Laurence (Eds.). *Creations of the Mind.* Oxford: Oxford University Press, 33–51.
Finlay, V. (2002). *Colour. Travels Through the Paintbox.* London: Sceptre.
Fransen, M. (2006). 'The Normativity of Artefacts', *Studies in the History of Philosophy of Science* 37, 42–57.
Gibson, J.J. (1979). *The Ecological Approach to Perception.* Boston: Houghton Mifflin.
Gibson, K. & Ingold, T. (1993). *Tools, Language and Cognition in Human Evolution.* Cambridge: Cambridge University Press.
Gutheil, G. et al. (2004). 'The Role of Historical Intuitions in Children's and Adults' Naming of Artefacts'. *Cognition* 91, 23–42.

Hampton, J.A. (1995). 'Testing the Prototype Theory of Concepts.' *Journal of Memory and Language* 34, 686–708.
Harnad, S. (2005). 'To Cognize is to Categorize: Cognition is Categorization'. In H. Cohen & C. Lefebre (Eds.), *Handbook of Categorization in Cognitive Science*, Amsterdam: Elsevier, 19–43.
Hilpinen, R. (1993). 'Authors and Artifacts'. *Proceedings of the Aristotelian Society* 93, 155–178.
Hilpinen, R. (2004). 'Artifact' In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, available at http://plato.stanford.edu/entries/artifact/.
Houkes, W. & Vermaas, P. (2004). 'Actions versus Functions. A Plea for an Alternative Metaphysics of Artefacts'. *Monist* 87, 52–71.
Hutchins, E. (1995). *Cognition in the Wild.* Cambridge MA: MIT Press.
Kornblith, H. (1980). 'Referring to Artefacts'. *Philosophical Review* 89, 109–114.
Kornblith, H. (2007). 'How to Refer to Artifacts'. In E. Margolis & S. Laurence (Eds.). *Creations of the Mind.* Oxford: Oxford University Press, 138–149.
Kroes, P. & Meijers, A. (Eds.) (2006). *The Dual Nature of Technical Artefacts.* Special issue in *Studies in the History of Philosophy of Science* 37, 1–158.
Lambrinos, D. & Scheier, C. (1996). 'Building Complete Autonomous Agents: A Case Study on Categorization'. In *Proceedings of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems.* Piscataway: Institute of Electrical and Electronics Engineers, Inc., 170–177.
Landau, B. (1994). 'Where's What and What's Where: The Language of Objects in Space'. *Lingua* 92, 259–296.
Landau, B., Smith, L. & Jones, S. (1998). 'Object Shape, Object Function, and Object Name'. *Journal of Memory and Language* 38, 1–27.
Lewens, T. (2004). *Organisms and Artefacts.* Cambridge MA: Bradford.
Malt, B.C. (1994). 'Water is not H_2O'. *Cognitive Psychology* 27, 41–71.
Malt, B.C. and Johnson, E.C. (1992). 'Do Artifact Concepts Have Cores?' *Journal of Memory and Language* 31, 195–217.
Malt, B.C. & Johnson, E.C. (1988). 'Artifact Category Membership and the Intentional-historical Theory'. *Cognition* 66, 79–85.
Margolis, E. & Laurence, S. (Eds.) (2007). *Creations of the Mind.* Oxford: Oxford University Press.
Marsh, L. & Onof, C. (2008). 'Stigmergic epistemology, stigmergic cognition'. *Cognitive Systems Research* 9, 136–149.
Matan, A. & Carey, S. (2001). 'Developmental Changes within the Core of Artifact Concepts'. *Cognition* 78, 1–26.
McClamrock, R. (1995). *Existential Cognition.* Chicago: University of Chicago Press.

Medin, D.L., & Ortony, A. (1989). 'Psychological Essentialism'. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning.* New York: Cambridge University Press.
Millikan, R. (1999). 'Wings, Spoons, and Quills: A Pluralist Theory of Function'. *Journal of Philosophy* 96, 191–206.
Minsky, M. (1986). *The Society of Mind.* New York: Simon & Schuster.
Mithen, S. (1996). *The Prehistory of the Mind.* London: Thames & Hudson.
Mithen, S. (2004). *After the Ice.* Cambridge MA: Harvard University Press.
Mithen, S. (2007). 'Creations of Pre-Modern Human Minds: Stone Tool Manufacture and Use by *Homo habilis*, *heidelbergensis*, and *neanderthalensis*'. In E. Margolis. & S. Laurence (Eds.). *Creations of the Mind.* Oxford: Oxford University Press, 289–311.
Murphy, G. (2002). *The Big Book of Concepts.* Cambridge MA: Bradford.
Noë, A. (2005). *Action in Perception.* Cambridge MA: Bradford.
Nolfi, S. (2005). 'Category formation in Self-Organizing Embodied Agents'. In H. Cohen & C. Lefebre (Eds.), *Handbook of Categorization in Cognitive Science*, Amsterdam: Elsevier, 869–889.
O'Regan, K. & Noë, A. (2001). 'A Sensorimotor Account of Vision and Visual Consciousness'. *Behavioral and Brain Sciences* 24, 939–972.
Petroski, H. (1994). *The Evolution of Useful Things.* New York: Vintage.
Pfeifer, R. & Scheier, C. (1999). *Understanding Intelligence.* Cambridge MA: MIT Press.
Poirier, P. et al. (2005). 'Embodied Categorization'. In H. Cohen & C. Lefebre (Eds.), *Handbook of Categorization in Cognitive Science.* Amsterdam, Elsevier, 739–765.
Preston, B. (1998a). 'Why Is a Wing Like a Spoon? A Pluralist Theory of Function'. *Journal of Philosophy* 95, 215–254.
Preston, B. (1998b). 'Cognition and Tool Use'. *Mind & Language* 13, 513–547.
Preston, B. (2003). 'Of Marigold Beer: A Reply to Vermaas and Houkes'. *British Journal for the Philosophy of Science* 54, 601–612.
Prinz, J. (2002). *Furnishing the Mind. Concepts and Their Perceptual Basis.* Cambridge MA: Bradford.
Putnam, H. (1975). 'The Meaning of "Meaning"'. In *Mind, Language and Reality.* Cambridge: Cambridge University Press.
Quillian, M. (1968). 'Semantic Memory'. In M. Minsky (Ed.), *Semantic Information Processing.* Cambridge MA: MIT Press, 227–270.
Rogers, T. & McClelland, J. (2004). *Semantic Cognition. A Parallel Distributed Processing Approach.* Cambridge MA: Bradford.
Rowlands, M. (1999). *The Body in Mind.* Cambridge: Cambridge University Press.
Ruff, H.A. (1984). 'Infants' Manipulative Exploration of Objects'. *Developmental Psychology* 20, 9–20.

Saunders, B. & van Brakel, J. (1997). 'Are There Non-trivial Constraints on Colour Categorization?'. *Behavioral and Brain Sciences* 20, 167–228.
Saunders, B. & van Brakel, J. (2002). 'The Trajectory of Color'. *Perspectives on Science* 10, 302–355.
Scheier, C. & Lambrinos, D. (1996a). 'Categorization in a Real-world Agent Using Haptic Exploration and Active Perception'. In P. Maes et al., *From Animals to Animats.* Cambridge MA: Bradford, 65–75.
Scheier, C. & Lambrinos, D. (1996b). 'Adaptive Classification in Autonomous Agents'. In R. Trappl (Ed.), *Proceedings of the European Meeting on Cybernetics and Systems Research (EMCSR '96).* Vienna: Austrian Society for Cybernetics Studies, 1011–1023.
Scheier, C. & Pfeifer, R. (1995). 'Classification as Sensori-motor Coordination'. In F. Moráu et al. (Eds.), *Advances in Artificial Life.* Berlin: Springer, 656–667.
Schwartz, S.P. (1978). 'Putnam on Artefacts'. *Philosophical Review* 87, 566–574.
Sloman, S.A. & Malt, B.C. (2003). 'Artifacts are Not Ascribed Essences, Nor Are They Treated as Belonging to Kinds'. *Language and Cognitive Processes* 18, 563–582.
Susi, T. & Ziemke, T. (2001). 'Social Cognition, Artefacts, and Stigmergy'. *Cognitive Systems Research* 2, 273–290.
Thomasson, A. (2003). 'Realism and Human Kinds'. *Philosophy and Phenomenological Research* 67, 580–609.
Thomasson, A. (2007). 'Artifacts and Human Concepts'. In E. Margolis. & S. Laurence (Eds.). *Creations of the Mind.* Oxford: Oxford University Press, 52–73.
Ullman, S. (1996). *High-level Vision. Object Recognition and Visual Cognition.* Cambridge MA: MIT Press.
Van Brakel, J. (1986). 'The Chemistry of Substances and the Philosophy of Mass Terms'. *Synthese* 69, 291–324.
Van Brakel, J. (2000). *Philosophy of Chemistry.* Leuven: Leuven University Press.
Varela, F., Thompson, E., Rosch, E. (1991). *The Embodied Mind.* Cambridge MA: MIT Press.
Verbeek, P.-P. (2005). *What Things Do. Philosophical Reflections on Technology, Agency, and Design.* Penn State: Pennsylvania State University Press.
Vermaas, P. & Houkes , W. (2003). 'Ascribing Functions to Technical Artefact'. *British Journal for the Philosophy of Science* 54, 261–289.
Ziman, J. (ed.) (2000). *Technological Innovation as an Evolutionary Process.* Cambridge: Cambridge University Press.