Rev.Phil.Psych. (2013) 4:465–481
DOI 10.1007/s13164-013-0148-1

A Taxonomy of Cognitive Artifacts: Function, Information, and Categories

Richard Heersmink

Published online: 25 May 2013

Abstract The goal of this paper is to develop a systematic taxonomy of cognitive artifacts, i.e., human-made, physical objects that functionally contribute to performing a cognitive task. First, I identify the target domain by conceptualizing the category of cognitive artifacts as a functional kind: a kind of artifact that is defined purely by its function. Next, on the basis of their informational properties, I develop a set of related subcategories in which cognitive artifacts with similar properties can be grouped. In this taxonomy, I distinguish between three taxa, those of family, genus, and species. The family includes all cognitive artifacts without further specifying their informational properties. Two genera are then distinguished: representational and non-representational (or ecological) cognitive artifacts. These genera are further divided into species. In case of representational artifacts, these species are iconic, indexical, or symbolic. In case of ecological artifacts, these species are spatial or structural. Within species, token artifacts are identified. The proposed taxonomy is an important first step towards a better understanding of the range and variety of cognitive artifacts and is a helpful point of departure, both for conceptualizing how different artifacts augment or impair cognitive performance and how they transform and are integrated into our cognitive system and practices.

1 Introduction

Situated cognition theorists argue that in order to properly understand a substantial part of our cognitive capacities, we should take into account the artifacts we deploy to help us perform our cognitive tasks. We use maps to navigate, notebooks to remember, rulers to measure, calculators to calculate, sketchpads to design, agendas to plan, textbooks to learn, and so on. Without such artifacts we would not be the same cognitive agents, as they allow us to perform cognitive tasks we would otherwise not

R. Heersmink (✉)
Department of Cognitive Science, Macquarie University, Sydney, Australia
e-mail: richard.heersmink@gmail.com

be able to perform. Cognitive artifacts are thus important to study, not only because they make us more powerful and versatile thinkers, but also because they shape and transform our cognitive system and cognitive practices, ontogenetically (Menary 2007; Kirchhoff 2011) and phylogenetically (Donald 1991).

Numerous philosophical and empirical studies conceptualize how we deploy artifacts to perform cognitive tasks (Kirsh and Maglio 1994; Kirsh 1995; Hutchins 1995; Clark and Chalmers 1998; Clark 2008; Heersmink 2013). These studies focus on well-chosen examples and demonstrate that *particular* artifacts play a crucial role in performing *particular* cognitive tasks. But, although there is a great variety in cognitive artifacts, there has been no substantial or detailed attempt to categorize different types of such artifacts. Some brief attempts have been made, but these focus only on a segment of representational artifacts (see e.g. Norman 1993; Nersessian et al. 2003; Brey 2005) and neglect non-representational cognitive artifacts. Due to this variety and neglect of non-representational artifacts in current categorizations, a taxonomy providing a systematic understanding of the different categories of cognitive artifacts and their different cognition-aiding properties would be of great help for situated cognition theorists. Andy Clark has aptly put it as follows:

> "The single most important task, it seems to me, is to better understand the range and variety of types of cognitive scaffolding, and the different ways in which non-biological scaffoldings can augment or impair performance on a task.... The Holy Grail here is a taxonomy of different types of external prop, and a systematic understanding of how they help (and hinder) human performance" (Clark 2004, p. 32–33).

Clark's suggestion has two interrelated aspects. In order to acquire a systematic understanding of how different types of cognitive scaffolds augment or impair performance on different cognitive tasks, it is helpful to first have an understanding of the range and variety of such scaffolds. In other words, a first possible step to obtain Clark's "Holy Grail" consists of a taxonomy of cognitive scaffolds, outlining their distinctive cognition-aiding properties. In this paper, I focus on one (important) type of such scaffolds–namely, artifacts–by developing an elaborate and detailed taxonomy in which cognitive artifacts with similar informational properties can be grouped. To the best of my knowledge, such a taxonomy does not yet exist in the literature. Although the functional role of token artifacts in performing cognitive tasks has been widely discussed in the philosophical and cognitive science literatures under the heading of situated cognition (Dourish 2001; Clark 2003, 2008; Robbins and Aydede 2009), metaphysical issues concerning (the categories of) cognitive artifacts have not received much attention by those authors and are in some cases even deliberately ignored (Hutchins 1995).

Whilst other authors working in philosophy of technology have addressed metaphysical issues concerning artifacts, they focus on technological artifacts in general and have little, if anything, particular to say about cognitive artifacts (Preston 1998, 2009, 2013; Houkes and Vermaas 2004, 2010; Kroes 2012). So there is a gap in the literature which this paper addresses by conceptualizing functional and informational properties of cognitive artifacts. By doing so, this paper brings into contact concepts and theories in (philosophy of) cognitive science and philosophy of technology and strengthens the rather thin ties between those fields. The taxonomy proposed in this paper thus builds intradisciplinary and interdisciplinary bridges.

This paper is structured as follows. It starts by identifying some of the components of situated cognitive systems and argues for the usefulness of a taxonomy of cognitive artifacts (section 2). Next, the target domain is identified and a method for taxonomizing is proposed (section 3). A taxonomy is then developed in which cognitive artifacts with similar informational properties can be grouped into categories (sections 4, 5 & 6).

2 Components of Situated Cognitive Systems

Situated cognitive systems have components that interact, transform each other, and are (in varying degrees) integrated into a larger cognitive system. A possible strategy to better understand the larger system is by decomposing it into its components and by conceptualizing the particular cognitive properties of those components. This strategy is helpful because those properties largely determine how the situated system performs its cognitive tasks and to what degree the components are integrated and transform each other. But before those properties can be conceptualized, we first need to identify the components. Some situated cognitive systems consist only of human agents such as, e.g., transactive memory systems in dyads (Harris et al. 2011) or larger groups (Theiner 2013); others consist of humans and artifacts; and yet others consist of humans and non-artifactual objects. The focus in this paper is on situated systems that consist of human agents and cognitive artifacts, while acknowledging the existence of systems consisting of other components.

2.1 Agents, Techniques, Artifacts, and Naturefacts

It may seem obvious that artifacts are easily identifiable components of certain situated systems. However, cognitive anthropologist Edwin Hutchins is critical about identifying and developing a distinct category of cognitive artifacts. In his seminal book, *Cognition in the Wild,* he writes:

> "We are all cognitive bricoleurs–opportunistic assemblers of functional systems composed of internal and external structures. In developing this argument I have been careful not to develop a class, such as cognitive artifacts, of designed external tools for thinking. The problem with that view is that it makes it difficult to see the role of internal artifacts, and difficult to see the power of the sort of situated seeing that is present in the Micronesian navigator's images of the stars" (Hutchins 1995, p. 172).

Hutchins' worries here are twofold: If we focus on cognitive artifacts as the most important component of situated cognitive systems, it makes it difficult to see the functional roles of (a) internal artifacts and structures, and (b) of external structures that are not human-made. To make his case, Hutchins points out that Micronesian navigators use the stars as perceptual anchors to navigate at sea. He argues that the interaction between internal artifacts, or perceptual strategies, and the external stars makes it possible for the Micronesians to navigate. Thus, in order to explain the navigational capacities of the Micronesians, we have to take into account their learned perceptual strategies and the stars, neither of which are proper artifacts. If we only focus on situated systems consisting of embodied agents and cognitive artifacts, we might overlook interesting cases like these, which would reduce explanatory power. And, furthermore, if we only focus on the artifactual component in such situated systems, we might overlook the functional role of what Hutchins calls internal artifacts. I am sympathetic to Hutchins' worries, but they can be overcome by acknowledging and emphasizing that we should be aware that artifacts are only one possible component of particular situated cognitive systems that interact and are integrated with other components. Developing a category of cognitive artifacts does, in my view, not mean that other components of situated systems are ignored or overlooked. Ultimately, we should study situated cognitive systems, rather than their components, but this does not mean that we cannot develop categories and vocabularies for their components, as this would equally reduce explanatory power.

Drawing on Hutchins' quote, I now distinguish some of these components. First, it is clarifying to make a distinction between technology and technique. A technology (or artifact) is usually defined as a physical object intentionally designed, made, and used for a particular purpose,[1] whereas a technique (or skill) is a method or procedure for doing something. Both technologies and techniques are intentionally developed and used for some purpose and are in that sense *artificial*, i.e., human-made. However, it is important to note, or so I claim, that they are not both *artifactual*. Only technologies are artifactual in that they are designed and manufactured physical objects and in this sense what Hutchins refers to as internal artifacts, such as perceptual strategies, can best be seen as cognitive techniques, rather than as internal artifacts. Moreover, given that these cognitive techniques are learned from other navigators and are thus first external to the embodied agent, it is perhaps more accurate to refer to them as *internalized* cognitive techniques, rather than as *internal* cognitive techniques.

Second, Hutchins writes that Micronesian navigators use the stars in the same way as manufactured navigational artifacts are used, loosely implying that they are a kind of cognitive artifact. Whilst stars are neither artificial nor artifactual in the sense just explained, they are nevertheless the perceptual object of a cognitive technique and have an important functional role in navigation. Using natural objects or structures for some purpose is not uncommon, for example, we might use a dead branch of a tree as a walking stick, a stone as a hammer, or, indeed, the stars to navigate. When doing so, the branch, stone, and stars are not intentionally made for those purposes and may be seen to form a bridge between natural objects and artifacts. Risto Hilpinen (2011), in his entry on artifacts in the *Stanford Encyclopedia of Philosophy*, refers to such objects as naturefacts. So, following Hilpinen's terminology, I suggest referring to natural objects that are used for cognitive purposes as cognitive naturefacts. This does not make them less important for performing certain cognitive tasks. My point is that because stars are not intentionally designed, manufactured, or modified for some purpose, they do not belong in a category of artifacts. Hutchins' example is apt in that it shows that humans as cognitive agents not only intentionally construct and modify

[1] This definition is sufficiently broad as to include less prototypical cases of artifacts such as domesticated animals (e.g., guide dogs) and genetically modified organisms (e.g., biofuel producing algae). Guide dogs and biofuel producing algae are intentionally modified (or trained) by humans to perform a particular function, i.e., guiding blind people or producing biofuel. Thus the material of which (cognitive) artifacts are made can be biological or non-biological and in some cases (cognitive) artifacts may even be alive, e.g., in the case of a guide dog.

their environment for cognitive purposes, but even exploit natural objects for cognitive purposes. Thus, to conclude this subsection, a preliminary and high-level taxonomy of cognitive scaffolds consists of embodied agents (or social scaffolds), cognitive techniques, cognitive artifacts, and cognitive naturefacts.

2.2 Why Focus on (Taxonomizing) Artifacts?

As Hutchins rightly pointed out, we are opportunistic assemblers of functional systems that are composed of internal (or internalized) structures and external structures. From an agent-centered perspective, it does not matter whether these external structures are artifactual or natural. What matters is that they functionally contribute to performing a cognitive task. So, in one sense, artifacts and naturefacts are continuous in that they can both function as external cognition-aiding structures. There is, however, one relevant difference between artifacts and naturefacts that justifies paying more attention to (taxonomizing) artifacts: our intentional control over the informational content and functions of cognitive artifacts is considerably larger and results in significantly more variety, as compared to the intentional control over cognitive naturefacts and the resulting variety. We have intentional control over the content and functions of cognitive naturefacts only insofar as we can choose which natural objects or structures to use for some cognitive task. We cannot actually modify naturefacts to help us perform a cognitive task, because when we do, they enter the realm of the artifactual.

One could argue that we also do not have full intentional control over the content and functions of all cognitive artifacts, as some are designed and made by others. Maps, timetables, textbooks, manuals, encyclopaedias, and roads signs, for example, are designed and made by agents outside the situated system and a user typically has no control over the content of such artifacts. The informational content of other cognitive artifacts, however, *is* designed and made by the user of the artifact. Notebook and agenda entries, to-do lists, shopping-lists, PowerPoint slides, and an architect's sketch, for example, are typically made by the user of the artifact. So the intentional control an agent has over the content and functions of a cognitive artifact differs, depending on the kind of artifact (Heersmink 2012). But, in either case, cognitive artifacts are intentionally designed and made to aid human cognition. The intentional design and making of cognition-aiding artifacts (either by designers or users) results in a broad range of cognitive artifacts exhibiting different kinds of informational properties that are specifically geared towards realizing a broad range of cognitive tasks, including navigating, calculating, remembering, measuring, planning, designing, etc. This kind of intentional control not only results in a much richer variety of cognitive artifacts, as compared to cognitive naturefacts, but also results in external artifactual structures that can be integrated much deeper into the onboard cognitive system, because they are functionally and informationally malleable. Consequently, the transformative impact of artifacts on our cognitive system and practices, both ontogenetically and phylogenetically, seems much more substantial as compared to naturefacts.

There are at least two other reasons why it is important to pay attention to (taxonomizing) cognitive artifacts. First, it gives us a much deeper conceptual understanding of a particular kind of artifact, namely, cognitive artifacts. This is important because it contributes to expanding and further developing a relatively small and emerging subfield in the philosophy of technology, which is sometimes referred to as philosophy of artifacts. As Randall Dipert (1993) has argued, an adequate philosophical theory of artifacts is largely lacking in the history of analytical Western thought (compare Houkes and Vermaas 2010; Kroes 2012; Preston 2013). Given the ubiquitousness of artifacts and their fundamental role in our lifeworld, culture, and cognition, an adequate philosophical theory of artifacts would be very much welcome. Developing a taxonomy of cognitive artifacts contributes to such a theory.

Second, the reason cognitive artifacts are important to better understand is because they have different properties and affordances as compared to cognitive techniques. Although cognitive techniques are also physical in that they supervene on neural and sometimes bodily structures, the physical material of artifacts is very different and allows operations that are very difficult to perform in the brain. The specific physicality and operations external artifacts allow, gives them particularly distinct functional and informational properties, which are important to study in their own right. As Merlin Donald (1991) has pointed out, exograms (or external representational systems) have properties that are different from engrams (internalized information in biomemory). Engrams are internalized and realized in the medium and format of the brain, whereas exograms are external and much less restrained in their format and capacity. The storage capacity of exograms far exceeds the storage capacity of both single entries and clusters of entries in biological memory. Exograms are flexible in that they can be reformatted and easily transmitted across different media, whereas engrams are less flexible. These differences certainly do not always apply, but when they do apply, they are enabled by the particular physicality, malleability, and format of external artifacts. *Because* exograms have such different functional and representational properties, as compared to engrams, they have the capacity to *complement* the shortcomings and limitations of engrams (Sutton 2006, 2010). This is arguably why we have developed exograms in the first place and if we want to understand how they complement the working of engrams and cognitive techniques, it is very helpful to have a taxonomy that distinguishes between different types of cognitive artifacts, conceptualizing their distinctive cognition-aiding properties.

3 Taxonomizing Cognitive Artifacts

3.1 Identifying the Target Domain

Having identified some important components of situated cognitive systems, I now focus on one of those components, i.e., cognitive artifacts. To get some clarity on what they precisely are, it is helpful to start with discussing some existing definitions:

- "Cognitive artifacts are artificial devices that maintain, display, or operate upon information in order to serve a *representational* function and that affect human cognitive performance" (Norman 1991, p. 17, emphasis added).
- "Cognitive artifacts are physical objects made by humans for the purpose of aiding, enhancing, or improving cognition" (Hutchins 1999, p. 126).

- "Cognitive artifacts are a special class of artifacts that are distinguished by their ability to *represent*, store, retrieve and manipulate information" (Brey 2005, p. 385, emphasis added).
- "Cognitive artifacts are material media possessing the cognitive properties of generating, manipulating, or propagating *representations*" (Nersessian 2005, p 41, emphasis added).

These definitions have two elements in common: cognitive artifacts are defined as (a) human-made, physical objects that (b) are deployed by human agents for the purpose of functionally contributing to performing a cognitive task. A third and more specific element which only Donald Norman, Philip Brey, and Nancy Nersessian's definitions contain is that cognitive artifacts provide (and sometimes manipulate or process) representational information. Hutchins' definition is neutral about the specific informational properties of cognitive artifacts. Of the above three elements, only the second is truly distinctive for the category of cognitive artifacts. In other words, their most distinctive property is their function, which is to contribute to performing a cognitive task. Given the centrality of function, I demarcate the boundaries of the target domain by conceptualizing cognitive artifacts as a functional kind, i.e., a kind of artifact that is defined purely by its function (Carrara and Vermaas 2009). In Hilary Kornblith's words: "At least for the most part, it seems that what makes two artifacts members of the same kind is that they perform the same function" (Kornblith 1980, p. 112). Thus, any artifact that functions to aid it user(s) in performing a cognitive task is a member of the functional kind of cognitive artifacts.

By contrast, artifacts with pragmatic functions such as hammers, coffee cups, chairs, screwdrivers, flower pots, light switches, and trash bins are not used for cognitive purposes, at least not in their normal use, because they have no functional role in performing a cognitive task. Hammering a nail into a wall, using a screwdriver to open a can of paint, or turning on a light switch, does not straightforwardly contribute to performing a cognitive task. We interact with those artifacts to create a change in the state or location of the artifact, not because it aids our cognition, but because that goal state is desirable for some practical purpose. The purpose for which such artifacts are deployed is therefore not cognitive.

Artifacts with pragmatic functions may of course influence human cognition, because sometimes we have to think about how to interact with those artifacts. However, being the object *of* perception and cognition is necessary but not sufficient for functionally contributing *to* a cognitive task. Many things are the object of perception and cognition, but do not in any obvious sense help us to perform or complete a cognitive task. However, this is not to say that such artifacts can never have cognitive functions. For example, when I leave a rented DVD on my desk to remind myself to bring it back to the video store, the DVD (in virtue of its location) arguably functions as a mnemonic aid. So during improvisation, we can attribute cognitive functions to initially non-cognitive artifacts, thereby demonstrating that we cognitively exploit not only our natural environment, but also our artifactual environment quite opportunistically. This example shows that a cognitive artifact is neither defined by intrinsic properties of the artifact nor by the intentions of the designer, but by its function, which is established by the intentions of the user and by how it is used. Thus, the target domain that is to be categorized are artifacts with cognitive functions, which may either be intentionally designed or improvised to perform that function.

3.2 Methodology

Current categorizations focus on representational artifacts and thus neglect non-representational artifacts. The taxonomy proposed in this paper does include non-representational artifacts and is thus more inclusive and has more explanatory power. Current categorizations are, moreover, anthropocentric, i.e., they start with human cognition and then categorize artifacts in terms of the cognitive process to which they contribute. So, for example, Norman (1993) distinguishes between two types of representational artifacts, those that aid experiential cognition and those that aid reflective cognition. Likewise, Brey (2005) develops a brief taxonomy of representational artifacts on the basis of the cognitive process to which they contribute (i.e., memory, interpretation, searching, and conceptual thought). In this paper I take an artifact-centered approach in categorizing cognitive artifacts: I take as my point of departure the specific properties of cognitive artifacts and then categorize them on the basis of those properties, not on the properties or goals of the agents that design, make, or use them.

The properties I focus on are informational properties, because a cognitive function supervenes on informational properties. Cognitive artifacts are used in virtue of the information they provide. A map, for example, is used to navigate because the information it provides is helpful for navigating. In other words, cognitive functions are established only when an artifact exhibits an information-structure that is used in performing some cognitive task. If there are no information-structures, then there are no cognitive functions. Exhibiting an information-structure is thus a necessary and sufficient condition for being a cognitive artifact. I distinguish between two kinds of information-structures: representational and non-representational information-structures. Representational artifacts contain information-structures *about* the world (i.e. representational information), whereas non-representational artifacts contain information-structures *as* the world (i.e. ecological information). These distinct informational properties are essential for better understanding the range and variety of cognitive artifacts and are therefore further conceptualized below.

The proposed taxonomy is hierarchical and starts at a very general and inclusive level and gets increasingly more specific and less inclusive when one goes deeper into the taxonomy. I distinguish between three levels or taxa, those of family, genus, and species (see Fig. 1 below). The family includes all cognitive artifacts without further specifying their informational properties. The family is thus a functional kind defined purely by the cognitive function of the artifact. On the second level in the taxonomy, i.e., the taxon of genus, I distinguish between two genera: representational and ecological artifacts. On the third level, these two genera are further divided into species. In case of representational artifacts, three species are identified, those that are iconic, indexical, and symbolic. In case of ecological artifacts, two species are identified, those that are spatial and structural. Within species, token artifacts are identified.

In the above taxonomy, family, genus, and species membership are monothetic, which means that one property is the determining factor for category membership. In case of family, having a cognitive function is the determining property for

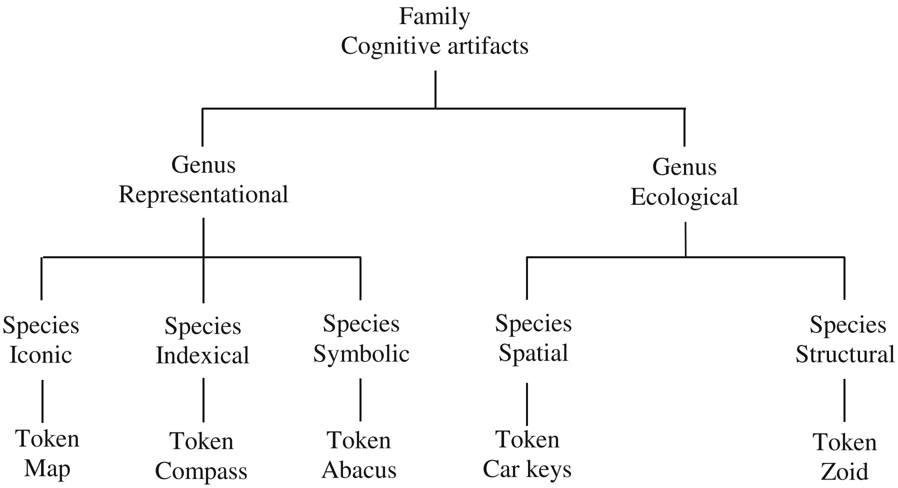

Fig. 1 A diagram of a taxonomy of cognitive artifacts, outlining three taxa: family, genus, and species

membership. In case of genus and species, exhibiting a particular informational property is the determining factor for membership. So this taxonomy combines functional and informational properties for category membership and is in that sense eclectic. Finally, upward deduction in the taxonomy is valid, but downward deduction is not. So, when membership at the species level is known, one can deduce the membership of the higher levels, but not vice versa.

4 Representational Cognitive Artifacts

Some artifacts exhibit representational properties. But what exactly are representational properties? To answer this question, we need to take a closer look at what a representation is. A useful starting point is the work of American pragmatist Charles Saunders Peirce, who has argued that a representational system involves a *triadic* relation between an *interpretant* (an understanding of the sign's object), *sign* (representation or representational vehicle), and *object* (represented world). A defining property of a representation is that it stands in for something else. In John Haugeland's words: "That which stands in for something else in this way is a *representation*; that which it stands in for is its *content*; and its standing in for that content is *representing* it" (Haugeland 1991, p. 62, original italics). A representation thus has aboutness in that its representational content is about something external to the representation.

4.1 Icon, Index, and Symbol

According to Peirce, a representational system is irreducibly triadic, but for analytical purposes, we can decompose Peirce's triadic relation into two dyadic relations; one between the representation and its object, and one between the human agent and the representation. Let us first have a look at the representation-object relation. Peirce

(1935a, b) distinguished between three possible grounds for representation, which are iconic, indexical, and symbolic (see also Von Eckhardt 1995). An icon is relevantly similar or isomorphic to what it represents. Straightforward examples are maps, radar systems, and blueprints, which are iconic because there is a high structural isomorphism between the content of the map, radar system, or blueprint and their objects.

Isomorphism is a quite general concept and because of its generality it can be pushed rather far, as to include representations such as graphs and other diagrams. Consider a line graph representing the amount of CO_2 in the earth's atmosphere plotted against the time. The amount of CO_2 has been increasing over the last hundred and 50 years or so and hence the graph will show a line going up. There is indeed some kind of isomorphism between an increasing amount of CO_2 particles in the atmosphere over a certain period of time and a line going up, but this isomorphism is of a very different kind than, for example, the isomorphism between a map and its object which is much stronger. Isomorphism between properties of the representation and its object may thus be abstract.

Indices have a direct causal connection between the index and its object. There is, for instance, a direct causal connection between the direction of the wind and the direction of a weathervane. If the direction of the wind changes, then the direction of the weathervane automatically changes, too. So the position of the weathervane is an index for the direction of the wind. Note that it is also partly iconic, since there is an isomorphism between the direction of the wind and the direction of the weathervane. Other examples of indices are thermometers, compasses, scales, voltmeters, speed dials, barometers, and many other measuring instruments. These are indices because there is a direct causal connection between the representational state of the index and its object. In these indices, the following things are causally connected: temperature and the expansion of mercury in a thermometer, the location of the North Pole and the direction of an arrow in a compass, the mass of a particular object and the reading on a scale, the amount of volts in an electrical current and the reading on a voltmeter, the speed of a car and the reading on a cars' speed dial, and atmospheric pressure and the reading on a barometer.

The last class of representations that Peirce identified are symbols, which acquire their meaning and content through shared use, agreement, and logical rules. Typical cases are words and sentences in natural language, numerals, or symbols in mathematical and scientific formula. The form or structure of symbols is quite often (though certainly not always) arbitrary. There is, for example, nothing intrinsic in the structure of the word "buildings" that makes it represent the category of buildings, there is likewise nothing intrinsic in the structure of the numeral "4" that makes it represent 4 units,[2] and there is nothing intrinsic in the structure of the sign for wavelength "λ" that makes it represent wavelength. There is no resemblance or causal connection between symbols and their object. It is shared use, agreement, and logical rules that establish the meaning of symbols, not their representational structure.

Sometimes other kinds of representations have symbolic properties as well. Certain icons such as diagrams cannot function as icons merely by their isomorphic relations to their object. There needs to be an agreement or convention that indicates which elements of the icon are to be interpreted as being isomorphic to its object.

[2] Although this is different in the Roman numeral system, see Zhang and Norman (1995).

Measurement outcomes of certain indices are also partly symbolic, e.g., temperature may be measured in degrees Celsius or Fahrenheit, depending on one's geographical location and convention. Thus, as Peirce recognises, any particular representational vehicle may display a combination of iconic, indexical, and symbolic characteristics. Consequently, a useful way of conceptualizing Peirce's trichotomy is by seeing a token representation as *predominantly* iconic, indexical, or symbolic (Atkin 2008).

The power of Peirce's trichotomy is that it can account for all forms of representation as predominantly iconic, indexical, or symbolic. Donald Peterson, in his edited book, *Forms of Representation: An Interdisciplinary Theme for Cognitive Science,* correctly points out that we use a great variety of representational systems, including:

> "algebras, alphabets, animations, architectural drawings, choreographic notations, computer interfaces, computer programming languages, computer models and simulations, diagrams, flow charts, graphs, ideograms, knitting patterns, knowledge representation formalisms, logical formalisms, maps, mathematical formalisms, mechanical models, musical notations, numeral systems, phonetic scripts, punctuation systems, tables, and so on" (Peterson 1996, p. 7).

Peirce's trichotomy can be used to categorize these representational systems. Algebras, alphabets, computer programming languages–knowledge representation, logical, or mathematical formalisms–musical notation systems, numeral systems, phonetic scripts, punctuation systems, and tables are predominantly symbolic, as they mainly acquire their meaning through logical rules and convention. But some of these systems may also display isomorphism with their objects. For example, in the case of musical notations, the order of the symbols for the notes is isomorphic with the order of the notes in the actual piece of music.

Animations, architectural drawings, choreographic notations, computer models and simulations, diagrams, flow charts, graphs, ideograms, knitting patterns, maps, and mechanical models are predominantly iconic, as they acquire their meaning through exhibiting some kind of isomorphism with their objects. However, some of these systems, for example, architectural drawings, computer models, and diagrams may also contain symbolic representations such as words, sentences, and numerals. And, more importantly, there are rules for interpreting certain elements within those representations as isomorphic to their objects. For example, the legend on a map often contains rules and guidelines for interpreting certain elements in the map as being isomorphic to certain objects and structures in the world.

Finally, computer interfaces do not really belong in that list, as they are not a form of representation, but more a medium in which a variety of representational systems can be expressed and computed. So, to conclude this subsection, iconic, indexical, and symbolic cognitive artifacts are species of the genus of representational cognitive artifacts and can account for all forms of representation.

4.2 Interpretation and Consumption

Having looked at different species of representational cognitive artifacts and some of their informational properties, let us now have a closer look at the dyadic relation between agent and representational vehicle. To establish a genuine triadic relationship, an agent has to interpret or consume the representation. For this to happen, an agent has to understand the dyadic relationship between the representation and what it represents. An agent thus has to understand the representation *qua* representation. This means that an interpreter has to realize that icons display isomorphic relations to their objects, that indices have causal relations to their objects, and that symbols are based on rules and conventions. So when I perceive a map of the London subway system, I understand that I perceive a predominantly iconic representation. Due to the properties of the vehicle, I understand that the London subway system is an existing structure in the world and because of this understanding I am able to form beliefs about certain properties of the London subway system (e.g. the order of the stations). Thus representations are psychologically efficacious in that they cause beliefs about the content of the vehicle. In this sense, they causally and informationally mediate between agent and world.

Representations can be misinterpreted, which will hinder cognitive performance. This occurs when the vehicle accurately captures its object, but for whatever reason, the agent misinterprets the vehicle and attributes certain properties to its object that do (in fact) not exist, resulting in false beliefs about aspects of its object. In such cases, there is no genuine triadic relationship between agent, representation, and world, but a mere dyadic relationship between agent and representation. It is of course possible that an agent only misinterprets certain elements of the representation. For example, when navigating the London subway system with the aid of a map (not an easy cognitive task for a tourist), I may correctly interpret my current location on the map, but misinterpret how many stations it is from my current location to Piccadilly Circus and, consequently, form a false belief about its location. In this example, there is a mixture of a true and false belief, but due to the true belief, a genuine triadic relation is established.

5 Ecological Cognitive Artifacts

Having outlined the genus of representational cognitive artifacts, I now focus on the genus of ecological cognitive artifacts. These are characterized not by exhibiting information *about* the world, but *as* the world and thus exhibit non-representational or ecological information. To give an example: When playing Tetris, one has to rotate a geometrical shape called a "zoid" so that it fits into a specific socket in the lower regions of a geometrical template. One can either choose to mentally rotate the zoid or to rotate it by means of a button-push. Experienced players have learned to rotate the zoid by means of a button-push, not only because it is significantly faster, but also because it relieves the brain from performing mental rotation (Kirsh and Maglio 1994; see also Clark and Chalmers 1998). The artificial rotation of a zoid in playing Tetris clearly has a functional role in performing a cognitive task. This functional role, however, is not established by exhibiting representational properties. Zoids do not stand in for something else, they have no representational content, and neither causally nor informationally mediate between agent and world. They are mere abstract geometrical shapes that are interacted with directly, without them representing something outside the game. So, rather than triadic situated cognitive systems (agent-representational artifact-world), ecological artifacts concern dyadic situated cognitive systems (agent-artifact). The functional role of the artifactual elements in such dyadic systems is thus established through non-representational or ecological information. But what exactly is ecological information? How can it aid their users in performing a cognitive task? And how can it be categorized?

The work of cognitive scientist David Kirsh (1995, 2006, 2009; Hollan et al. 2000) is a useful starting point in answering these questions. In his seminal paper, Kirsh (1995) makes a tripartite distinction between spatial arrangements that simplify choice, perception, or internal computation. Although this a helpful and insightful categorization, it is (like all other currently existing categorizations) anthropocentric, i.e., it starts with a cognitive agent and then categorizes artifacts on the basis of the cognitive processes to which they contribute. In the taxonomy I am developing in this paper, I take an artifact-centered approach and taxonomize cognitive artifacts on the basis of their informational properties. With this non-anthropocentric approach in mind, I recycle and reclassify some of Kirsh's examples and focus on their informational properties, rather than on the cognitive processes to which they contribute. By doing so, I distinguish between two species of ecological artifacts, those that obtain their cognitive function in virtue of physical-spatial or virtual-spatial structures and those that obtain their cognitive function in virtue of manipulable physical or virtual structures.

5.1 Spatial Ecological Cognitive Artifacts

Human agents quite frequently make use of space for cognitive purposes. Kirsh refers to this as "the intelligent use of space", which is so commonplace, he argues, that "we should not assume that such cognitive or informational structuring is not taking place all the time" (Kirsh 1995, p. 33). The intelligent use of space enables us to encode important information into artifacts that are typically neither designed nor made for cognitive purposes and thus mainly (though not solely) concerns improvised uses of artifacts. Some straightforward examples include consistently leaving car keys on a certain spot in your apartment so that you know where they are, putting an article you have to read on top of the pile on your desk, leaving a book open and turned upside down so that you know where you have stopped reading, tying a string around your finger as a reminder, or leaving a rented DVD on your desk as a prompt to bring it back to the video store. By putting artifacts in certain locations that are either deliberately usual or deliberately unusual, we intentionally encode information into the artifact and its location, thereby creating what may be referred to as spatial ecological cognitive artifacts.

Consider another example: When doing the dishes, it is not always clear which items have been washed and which ones have not. In order to keep track of the items that have been washed, it is helpful to create spatial categories of "washed items" and "unwashed items" by putting them in certain locations. Most kitchen sinks have designated areas for items that have been recently washed. These areas are not just practical, so that residual dishwater can drip away without spilling it on the kitchen counter, but also have cognitive functions, as they simplify perception and reduce memory load in a task.

Kirsh (1995) describes a more idiosyncratic case of someone who is dismantling a bicycle and then puts certain parts on a sheet of newspaper placed on the floor. The newspaper demarcates a spatial boundary within which certain items are placed, in that way structuring the task space and making items easier to locate. The user may also place the dismantled items in such a way that, when reassembling the bicycle, the items that need to be reassembled first are located closer to the user than the items that are reassembled later in the process. “The virtue of spatially decomposing the task is that one need not consult a plan, except at the very highest level, to know what to do. Each task context affords only certain possibilities for action” (Kirsh 1995, p. 44). Kirsh here is referring to a mental plan, but it could also be an external one, e.g., an assembly guide, manual, or blueprint of some kind. There is no need for a manual (or an elaborate mental plan) if all the parts are placed such that they correspond to the correct order of actions for reassembling the bicycle. Spatial structuring of artifacts thus makes both complex internal representations and external representations superfluous. It is much more efficient to spatially structure the artifactual task environment such that it affords the most efficient and environmentally-embedded plan.

The intelligent use of space, however, is not restricted to physical or actual space, as it also frequently occurs in virtual space. In connection to Human-Computer Interaction, Hollan, Hutchins, and Kirsh provide some suitable examples. Computer users, they write;

> “Leave certain portals open to remind them of potentially useful information or to keep changes nicely visualized; they shift objects in size to emphasize their relative importance; and they move collections of things in and out of their primary workspace when they want to keep certain information around but have other concerns that are more pressing” (Hollan et al. 2000, p. 190).

So the way we spatially organize the items and structures on our screen helps us in performing certain computer-related tasks. We may, for example, prioritize certain information by leaving certain portals open or by making them larger than other portals. We may also organize the items on our desktop or in our navigation menu’s such that they reflect their importance. Often used items typically inhabit a more prominent position than items that are used less often (for example, by placing them in preferred positions on one’s desktop or menu) thereby making them easier to locate. These everyday examples show that we organize artifactual elements (physical or virtual) within space (actual or virtual) for cognitive purposes, thereby encoding important information into physical-spatial structures or virtual-spatial structures.

5.2 Structural Ecological Cognitive Artifacts

Some artifacts obtain their cognitive function in virtue of their manipulable physical or virtual structure. During or after the manipulation of such artifacts, new information emerges that is important for performing a cognitive task. For example, when rearranging letter tiles in Scrabble to prompt word recall, new information emerges from their spatial configurations (Kirsh 2009). In this case, the non-representational information of the tiles (i.e. individual letters), their spatial location in relation to the other tiles, and the words and openings on the board are important for performing the task. Novel and larger information-structures emerge when the tiles are rearranged, which (ideally) prompts the recall of words with as many letters as possible and that fit into existing letter structures on the board. Kirsh points out that one of the purposes for rearranging letters tiles, is to generate and draw attention to often occurring two- and three-letter combinations that figure in words such as, for instance, “ES”, “TH”, "IN", "REA", et cetera. Note that such information (i.e. letter combinations) are often not representational in nature, although there is a symbolic element to it, because there are logical rules and conventions for creating words out of letters. So perhaps letter tiles can be seen as sub-representational, as they are the building blocks of proper representations. The point here is that the cognitive function of the artifacts supervenes on their manipulable physical structures and by manipulating their physical structure one automatically manipulates the information they contain, too.

Virtual structures, too, can be manipulated such that new information emerges from their spatial configurations. The functional role of the artificial rotation of a zoid in playing Tetris is established due to a constantly manipulable virtual structure, i.e., *because* the zoids can be manipulated, a user delegates rotation to the computing device, in that way enabling a user to decide quicker whether or not it fits into a socket in the lower regions of the task space. Like with rearranging Scrabble tiles, new task-relevant information emerges only in relation to some other structure. In Scrabble, new information emerges when two or more tiles are positioned in a certain way. In Tetris, new information emerges when the zoid is spatially orientated in relation to a template in a certain way. The spatial orientation of the zoid in itself, i.e., without taking into account the structure of the template in which it has to fit, is not sufficient for performing the task. Both zoid and template are important for generating the task-relevant information. Finally, the above artifacts are much more dynamic than cases of spatial ecological cognitive artifacts. They are not about developing static spatial categories in which artifacts are placed, but about dynamically and constantly changing information in an ongoing task.

6 Representational and Ecological Artifacts

Ecological artifacts, I think, are instances of what roboticist Rodney Brooks (1999) calls "the world is its own best model" and what Clark (1989) calls the "007 principle". The point of these notions is: Why create an expensive internal representation of the world, when you can use the world itself as a model? In a similar way we may ask: Why create an expensive external representation of the world, when you can use the world itself to interact with and use it to facilitate your cognitive tasks? One could, for example, consult a manual for how to reassemble a dismantled bicycle. But one can also structure the items such that they facilitate the reassembling process, in that way streamlining the task and making a cost-expensive external representation superfluous. In Clark's words, "evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned" (Clark 1989, p. 64).

For analytical purposes, I presented representational and ecological artifacts as distinct genera. Up to this point, a reader may get the impression that a cognitive artifact either exhibits representational or ecological information. This analytical distinction was helpful in that it allowed me to emphasize and conceptualize their distinct cognition-aiding informational properties. However, in some cases, cognitive artifacts exhibit a combination of representational and ecological properties. To give an example: In a bookcase in which the books are alphabetically organized, representational and ecological structures jointly facilitate a perceptual task. The representational structures are the names and titles on the back of the books and the ecological properties are the spatial order in which they have been placed. In this example, the representational properties determine the spatial structure, which in turn, supports the representational structure. So alphabetically organized books may be seen as a predominantly representational cognitive artifactual structure. The genera of representational and ecological artifacts are thus not mutually exclusive.

7 Conclusion

This paper has developed a taxonomy in which cognitive artifacts with similar informational properties can be grouped into categories. An artifact-centered approach was taken: I took as my point of departure the specific informational properties of cognitive artifacts and then categorized them on the basis of those properties, not on the properties or goals of the agents that design, make, or use them. In the developed taxonomy, I distinguished between three taxa, those of family, genus, and species. This taxonomy is an important first step towards a better understanding of the variety and range of cognitive artifacts and their distinct cognition-aiding properties.

Acknowledgments I would like to thank my supervisors John Sutton and Richard Menary, Peter Woelert, and two anonymous reviewers for helpful comments on an earlier version of this paper. I would also like to thank Sadjad Soltanzadeh for inviting me to present this paper at a colloquium at the Centre for Applied Ethics and Public Philosophy (CAPPE) at Charles Sturt University in Canberra as well as the audience, particularly Steve Clarke.

References

Atkin, A. 2008. Icon, index, and symbol. In *The Cambridge encyclopaedia of language sciences*, ed. P. Hogan, 367–368. Cambridge University Press.

Brey, P. 2005. The epistemology and ontology of human-computer interaction. *Minds & Machines* 15(3–4): 383–398.

Brooks, R. 1999. *Cambrian intelligence: The early history of the new AI*. MIT Press.

Carrara, M., and P. Vermaas. 2009. The fine-grained metaphysics of artifactual and biological functional kinds. *Synthese* 169(1): 125–143.

Clark, A. 1989. *Microcognition: Philosophy, cognitive science and parallel distributed processing*. MIT Press.

Clark, A. 2003. *Natural born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford University Press.

Clark, A. 2004. Towards a science of the biotechnological mind. In *Cognition and technology: Co-existence, convergence, and co-evolution*, ed. B. Gorayska and J. Mey, 25–36. John Benjamins Publishing Company.

Clark, A. 2008. *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford University Press.

Clark, A., and D. Chalmers. 1998. The extended mind. *Analysis* 58: 10–23.

Dipert, R. 1993. *Artifacts, art works, and agency*. Temple University Press.

Donald, M. 1991. *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Harvard University Press.

Dourish, P. 2001. *Where the action is: The foundations of embodied interaction*. MIT Press.

Harris, C.B., P.G. Keil, J. Sutton, A.J. Barnier, and D.J.F. McIlwain. 2011. We remember, we forget: Collaborative remembering in older couples. *Discourse Processes* 48(4): 267–303.

Haugeland, J. 1991. Representational genera. In *Philosophy and connectionist theory*, ed. W. Ramsey, S. Stich, and D. Rumelhart, 61–89. Lawrence Erlbaum Associates.
Heersmink, R. 2012. Mind and artifact: A multidimensional matrix for exploring cognition-artifact relations. In *Proceedings of the 5th AISB Symposium on Computing and Philosophy*, ed. J.M. Bishop and Y.J. Erden, 54–61.
Heersmink, R. 2013. Embodied tools, cognitive tools, and brain-computer interfaces. *Neuroethics* 6(1): 207–219.
Hilpinen, R. 2011. Artifact. *Stanford Encyclopaedia of Philosophy.* Retrieved 23 February 2012, from http://plato.stanford.edu/entries/artifact/.
Hollan, J., E. Hutchins, and D. Kirsh. 2000. Distributed cognition: Toward a new foundation for human-computer interaction research. *Transactions on Computer-Human Interaction* 7(2): 174–196.
Houkes, W., and P. Vermaas. 2004. Actions versus functions: A plea for an alternative metaphysics of artifacts. *The Monist* 87(1): 52–71.
Houkes, W., and P. Vermaas. 2010. *Technical functions: On the use and design of artefacts*. Springer.
Hutchins, E. 1995. *Cognition in the wild*. MIT Press.
Hutchins, E. 1999. Cognitive artifacts. In *The MIT encyclopedia of the cognitive sciences*, ed. R.A. Wilson and F.C. Keil, 126–128. MIT Press.
Kirchhoff, M.D. 2011. Extended cognition and fixed properties: Steps to a third-wave version of extended cognition. *Phenomenology and the Cognitive Sciences* 11(2): 287–308.
Kirsh, D. 1995. The intelligent use of space. *Artificial Intelligence* 72: 31–68.
Kirsh, D. 2006. Distributed cognition: A methodological note. *Pragmatics and Cognition* 14(2): 249–262.
Kirsh, D. 2009. Problem-solving and situated cognition. In *The Cambridge handbook of situated cognition*, ed. P. Robbins and M. Aydede, 264–306. Cambridge University Press.
Kirsh, D., and P. Maglio. 1994. On distinguishing epistemic from pragmatic action. *Cognitive Science* 18: 513–549.
Kornblith, H. 1980. Referring to artifacts. *Philosophical Review* 89(1): 109–114.
Kroes, P. 2012. *Technical artefacts: Creations of mind and matter.* Springer.
Menary, R. 2007. *Cognitive integration: Mind and cognition unbounded*. Palgrave McMillan.
Nersessian, N.J. 2005. Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In *New directions in scientific and technical thinking*, ed. M. Gorman, R. Tweney, D. Gooding, and A. Kincannon, 17–56. Erlbaum.
Nersessian, N.J., E. Kurz-Milcke, W.C. Newstetter, and J. Davies. 2003. Research laboratories as evolving distributed cognitive systems. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, ed. R. Alterman and D. Kirsh, 857–862.
Norman, D. 1991. Cognitive artifacts. In *Designing interaction: Psychology at the human-computer interface*, ed. J.M. Carroll, 17–38. Cambridge University Press.
Norman, D. 1993. *Things that make us smart: Defending human attributes in the age of the machine*. Basic Books.
Peirce, C.S. 1935a. *The collected papers of Charles S. Peirce vol 2*. Harvard University Press.
Peirce, C.S. 1935b. *The collected papers of Charles S. Peirce vol 3*. Harvard University Press.
Peterson, D. 1996. Introduction. In *Forms of representation: An interdisciplinary theme for cognitive science*, ed. D. Peterson, 7–27. Intellect Books.
Preston, B. 1998. Why is a wing like a spoon? a pluralist theory of function. *The Journal of Philosophy* 95(5): 215–254.
Preston, B. 2009. Philosophical theories of artifact function. In *Philosophy of technology and engineering sciences*, ed. A. Meijers, 213–234. Elsevier.
Preston, B. 2013. *A philosophy of material culture*: *Action, function, and mind.* Routledge.
Robbins, P., and M. Aydede (eds.). 2009. *The Cambridge handbook of situated cognition*. Cambridge University Press.
Sutton, J. 2006. Distributed cognition: Domains and dimensions. *Pragmatics and Cognition* 14(2): 235–247.
Sutton, J. 2010. Exograms and interdisciplinarity: History, the extended mind and the civilizing process. In *The extended mind*, ed. R. Menary, 189–225. MIT Press.
Theiner, G. 2013. Transactive memory systems: A mechanistic analysis of emergent group memory. *Review of Philosophy and Psychology* 4(1): 65–89.
Von Eckhardt, B. 1995. *What is cognitive science?* MIT Press.
Zhang, J., and D. Norman. 1995. A representational analysis of numeration systems. *Cognition* 57: 271–295.