

The metaphysics of cognitive artefacts

Richard Heersmink*

Centre for Cognition and its Disorders, Macquarie University, Sydney, Australia

This article looks at some of the metaphysical properties of cognitive artefacts. It first identifies and demarcates the target domain by conceptualizing this class of artefacts as a functional kind. Building on the work of Beth Preston, a pluralist notion of functional kind is developed, one that includes artefacts with proper functions and system functions. Those with proper functions have a history of cultural selection, whereas those with system functions are improvised uses of initially non-cognitive artefacts. Having identified the target domain, it then briefly looks at the multiple usability of physical structures and the multiple realizability of cognitive function. Further developing insights from the “dual nature of artefacts thesis”, the article ends with conceptualizing the structure–function relations of cognitive artefacts. More specifically, it unpacks the relation between physical structure, representational structure, information, and cognitive function.

Keywords: cognitive artefacts; function; metaphysics of artefacts; situated cognition; epistemic action

1. Introduction

Embodied agents like us have limited information-storage and information-processing capacities. In order to complement these limitations, we use maps, diagrams, models, diaries, timetables, calculators, computer systems, and other artefacts to help us perform our cognitive tasks. The informational properties and functionalities of such artefacts are crucial for performing a wide range of cognitive tasks, including navigating, calculating, planning, remembering, decision-making, and reasoning. Using artefacts to perform cognitive tasks gives us clear epistemic benefits, as they make such tasks easier, faster, more reliable, or possible in the first place (Kirsh and Maglio 1994). During ontogeny, we learn how to incorporate such artefacts into our cognitive system, resulting in integrated agent–artefact systems that are much more powerful, reliable, and versatile problem-solving systems than embodied agents alone. Because of these epistemic benefits, it is important to better understand the informational and functional properties of such artefacts.

Situated cognition theory has analysed the epistemic roles of artefacts in performing cognitive tasks (Norman 1993; Hutchins 1995, 1999; Clark and Chalmers 1998; Clark 2003, 2008; Menary 2007; Kirsh 2013). However, metaphysical properties of such artefacts have been largely neglected by those theorists, perhaps because their explanatory targets are situated cognitive *systems*, not *artefacts*. While this explanatory focus is understandable, a better understanding of the artefactual element in situated cognitive systems has a trickledown effect for better understanding the overall situated system. Conversely, analytic philosophy of technology has addressed metaphysical properties of artefacts, but has neglected cognitive

*Email: richard.heersmink@gmail.com

artefacts, and focussed on technological devices in general (Meijers 2001; Bakker 2004; Houkes and Vermaas 2010; Kroes 2012; Preston 2013; Franssen et al. 2014). So the situation seems to be this: there is an important class of artefacts about which we lack essential metaphysical knowledge. The goal of this article is to obtain some of this knowledge, which is beneficial for both situated cognition theory and the metaphysics of artefacts and, moreover, stimulates dialogue between these two traditionally distinct branches in philosophy.

The article proceeds as follows. It starts by conceptualizing cognitive artefacts as a functional kind, i.e. a kind of artefact defined purely by its function. Building on the work of Preston (1998, 2013), a pluralist notion of functional kind is developed, one that includes artefacts with proper (selected) functions and system (improvised) functions. The multiple usability of physical structures and the multiple realizability of cognitive function are also briefly examined (Section 2). Next, given the centrality of function for defining cognitive artefacts, it is important to understand how their functions are established. Drawing on insights from the “dual nature of artefacts thesis” (Vermaas and Houkes 2006; Kroes and Meijers 2006; Houkes et al. 2011; Kroes 2012; compare Vaesen 2011; Vaccari 2013), it will unpack the relation between physical structure, representational structure, information, and cognitive function. Better understanding this relation contains the key to better understanding cognitive artefacts and their functions (Section 3).

2. Cognitive artefacts as a functional kind

One way to think about classifying artefacts is by conceiving them as human-made functional entities and to group them in categories based on their function (Heersmink 2013). The function of cognitive artefacts, I claim, is to provide task-relevant information, thereby *complementing* internal storage and processing systems (Sutton 2010; Heersmink 2014) and making certain cognitive tasks easier, faster, more reliable, or possible at all (Kirsh and Maglio 1994). A map, for example, is a cognitive artefact because its function is to provide task-relevant information used for navigating. A chair, by contrast, does not have as its function to provide information and is not used to perform cognitive tasks, at least not in its proper use, but to support an agent in a sitting position. Functions thus define artefacts. For this reason, I demarcate the boundaries of the target domain by conceptualizing it as a functional kind, i.e. a kind of artefact that is purely defined by its function (Carrara and Vermaas 2009; Kroes 2012). In Kornblith’s (1980, 112) apt 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.”

2.1. Artefact functions

Given the centrality of function in the demarcation of cognitive artefacts from other kinds of artefacts, it is important to further articulate what a function is. The contemporary debate on function in philosophy of science originated in the works of Carl Gustav Hempel and Ernest Nagel, who were concerned with functional explanation in science, particularly in biology. Larry Wright, Robert Cummins, Ruth Garrett Millikan, Karen Neander, and others later added substantial content to the function debate in philosophy of biology, which remains a prominent topic of discussion. Philosophy of mind and philosophy of cognitive science also have given a fair share of attention to the notion of function. In functionalism, mental states are conceptualized not by their material constitution, but according to their functional role in an overall mental economy. Thus, the nature of a mental state is identified by its causal-functional role in relation to sensory input, other mental states, and behavioural output.

However, relatively little attention has been given to the concept of function in philosophy of technology. Only in the last decade or so, philosophers have tried to articulate and clarify a function theory for artefacts (Preston 1998, 2009, 2013; Meijers 2001; Houkes and Vermaas 2004, 2010; Kroes 2012). One of the first philosophers who tried to give a substantial account of function for the artefactual domain is Preston (1998). Building on Wright (1974), Cummins (1975), Millikan (1984), and others, she developed a pluralist theory of function for artefacts that combines the notions of “proper function” and “system function”. Some philosophers have tried to reduce one of those notions to the other, arguing that there ought to be a universal account of function. But Preston argues that both notions explain different phenomena and both are needed for a complete and coherent function theory of artefacts. Let me briefly explain these two kinds of function.

2.1.1. *Proper functions*

Millikan (1984) has argued for an etiological theory of function-ascription, which claims that in order to understand the function of a biological trait or technological artefact one has to take into account the causal-historical background of that trait or artefact. She develops the concept of proper function which is established by a causal selection history. For example, hearts pump blood and in doing so they also make a certain sound, which has nothing to do with their function. Hearts do not exist because they make a certain sound, but because their function (i.e. pumping blood) has contributed to the fitness of the organism. The sounds that hearts make are a mere epiphenomenon of the mechanical workings of the heart and are irrelevant for the successful reproduction of the organism. Natural selection does not and cannot select on the basis of epiphenomena, but on the basis of effectively performing functions that contribute to successful reproduction. Likewise, chairs exist and are re-produced because they are widely used to sit on, not because you can stand on them or use them to block your door. So if one wants to understand the function of an organ, trait, or artefact, one has to look at the selection history of that entity. For the biological domain, this is an evolutionary history of natural selection. For the artefactual domain, it is a history of cultural selection by users.

Proper functions of artefacts are thus established through a process of cultural selection. Artefacts are designed or invented, and if they are successful in performing their function, they will be re-produced. Somewhere in the past, chairs have been invented to sit on and have been quite successful in their re-production. So the proper function of a chair is to sit on, because they have been selected for this purpose by previous generations of users. Chairs can also be used for other purposes, for example, to hang your coat on, to stand on, or to block a door from opening. But these purposes are not the reason why chairs are selected by their users and are, therefore, not their proper function. Most artefacts have one proper function: chairs are for sitting, pens are for writing, cars are for transportation, calculators are for calculating, etc. Some artefacts have more than one proper function. A Swiss Army knife, for example, is for cutting, opening bottles, opening cans, sawing materials, clipping nails, and so on. Swiss Army knives are selected by their users, not for one particular function, but for many functions and are thus a multifunctional artefact with numerous proper functions. A number of proper functions can thus coexist in one artefact.

2.1.2. *System functions*

Cummins (1975) argued that functional explanation in science is not based on causal-historical selection, but on the current capacities and dispositions of a whole system in terms of

its components. By giving a number of examples, Cummins argues that causal-historical selection cannot explain the existence of certain organs or traits. Penguins, for example, have wings that are not used for flying but for swimming. Their function (i.e. to enable underwater transport) has historically not been selected by natural selection. Likewise, in the artefactual domain there are numerous examples of artefacts that are used for purposes that were not intended by their designers and are not the reason why such artefacts are selected by their users. Chairs are used to stand on, screwdrivers are used for opening cans of paint, books are used to support computer screens, etc. In such instances, the function of those artefacts has nothing to do with their history of cultural selection and has everything to do with their current capacities or dispositions in a given context. Thus, the current function of an organ, trait, or artefact is not necessarily linked to its selection history, but is in principle divergent. Preston (1998, 2013) refers to Cummins' notion of function as "system function" and I will use this terminology throughout this article.

System functions of artefacts are either improvised uses of artefacts or the functions of novel prototypes. In case of novel prototypes, the first generations only have system functions, which are over time consolidated into proper functions. In case of improvised uses of artefacts, Preston develops two conceptions of system function in analogy to exaptations in biology. An exaptation occurs when a biological trait evolves such that it loses its original function and obtains a new function. Bird feathers are a classic example, which initially evolved for insulation, but were later adapted for aerodynamic purposes during flight. Preston makes a distinction between two types of system function, namely, "standardized ongoing exaptations" and "idiosyncratic ongoing exaptations". Standardized ongoing exaptations are repeated uses of artefacts that are not their proper function. Examples include using a chair to stand on, a screwdriver to open a can of paint, or a spoon to open a cocoa tin. Such uses are not intended by their designers and are not the reason why such artefacts are selected by their users, but they are nevertheless widespread cultural practice. Such system functions are quite often ongoing additions to the artefacts' proper function. Hence, proper functions and system functions can coexist in one artefact.

Idiosyncratic ongoing exaptations are improvised uses of artefacts by individuals. They are not widespread cultural practice, but ongoing uses of artefacts for individuals or small groups of individuals. Preston put forward a number of examples, including the use of an old cast iron as a bookend, or using a shoelace to tie up a tomato plant. Such idiosyncratic uses of artefacts may over time become more established and may spread to other social groups. They could potentially even become a consolidated proper function. Preston (1998, 253) argues that the notion of system function is "crucial in understanding the history of hominid tool use, which developed from the simple exaptive use of naturally occurring objects as hammers, digging sticks, and so on, to the pervasively artifactual environment we Western industrialized humans inhabit today". System functions are thus important because they explain the development of artefacts and technology. A function theory focussing merely on proper functions would only be able to account for certain uses of artefacts and not for others. Proper and system function seem to complement each other and explain different phenomena, which is an argument for a pluralist theory of function.

I would like to add a third type of system function, namely, "idiosyncratic exaptations". Preston's examples of standardized and idiosyncratic exaptations are *ongoing*. Standardized ongoing exaptations are often repeated uses of an artefact by a community of users and idiosyncratic ongoing exaptations are ongoing uses of artefacts for individuals or small groups of individuals. So the former are culturally well-established uses of artefacts for a community and the latter are well-established (or fairly well-established) uses of artefacts for (small

groups of) individuals. However, there are also cases of idiosyncratic uses of artefacts that are not well-established but one-offs, which are neither culturally widespread nor are they ongoing. For instance, someone may use a screwdriver as a weapon in order to defend oneself, or someone might use a hammer to break a window to get into the house because she has forgotten her key. Such uses can be called idiosyncratic exaptations, as they may occur only once or twice in a lifetime. The notion of idiosyncratic exaptations thus broadens the spectrum of system functions as to include improvised one-offs.

Finally, the above distinctions should perhaps not be seen as strict subcategories of system function, but rather as points on a continuum of system functions. Proper functions are rather constant over time, but system functions come and go. They may vary between improvised one-offs, ongoing and well-established uses of artefacts for individuals, or widespread uses of artefacts for a community of users, and everything in between. There are no clear-cut criteria to indicate when idiosyncratic exaptations become idiosyncratic ongoing exaptations. There are, likewise, no clear-cut criteria to indicate when idiosyncratic ongoing exaptations become standardized ongoing exaptations. It is therefore better to conceive of system functions as a continuum.

2.2. *A pluralist notion of functional kind*

In this subsection, I further specify the target domain by developing a pluralist notion of the functional kind of cognitive artefacts. It is pluralist in the sense that it contains cognitive artefacts with both proper and system functions. Those with proper functions have a history of cultural selection, whereas those with system functions have improvised uses of initially non-cognitive artefacts. Abacuses, for example, have a long history of cultural selection, as they have been selected by their users at least since the invention of the Salamis Tablet, an abacus-like device that was used for performing calculations, dating back to roughly 300 BC. The proper function of an abacus is, therefore, to aid its user in performing calculations, because it has been designed and selected for that purpose. Abacuses may also be used for other purposes, e.g. as a kids' toy, but this is not the reason why they are designed and not the reason why they are selected by their users and is thus not their proper function.

Likewise, maps, computers, calendars, radars, rulers, thermometers, speed dials, and countless other cognitive artefacts are designed to perform cognitive functions and have a history of cultural selection. Their proper function is thus to aid their users in performing cognitive tasks. It may, therefore, be argued that such artefacts have a *cognitive proper function*. Some of these artefacts have more than one cognitive proper function. Computers, for instance, are highly multifunctional devices and may be seen as the Swiss Army knife of the cognitive artefacts (Brey and Søraker 2009). Computers are usually not selected for one particular purpose (although in exceptional cases that may happen), but for many purposes, including web browsing, text processing, storing documents and other data, making spreadsheets, making PowerPoints, etc. Numerous cognitive proper functions may, therefore, coexist in one artefact. Furthermore, as computers may also be used for non-cognitive purposes such as playing music or online shopping, cognitive and non-cognitive functions can coexist in one artefact.

Other artefacts obtain their cognitive function through improvised uses. Sometimes such improvised uses are one-offs, e.g. when I am in a cafe and suddenly have an important idea which I quickly write down on a napkin before I forget it. Napkins have neither been designed nor selected to store information and aid memory, but to help clean or absorb liquid. However, during improvisation we can offload information onto the napkin (or

any other artefact that affords writing on it), thereby attributing a *cognitive system function* to the napkin. More specifically, as this is a one-off, it may be argued that the napkin becomes a cognitive artefact with an *idiosyncratic cognitive system function*. We may also attribute idiosyncratic cognitive system functions to a set of artefacts. Consider Norman's (1993) example of trying to explain and reconstruct how an accident happened with the aid of everyday artefacts such as pencils and paperclips as stand-ins for the objects (cars and dog) they represent. One pencil stands in for a car that was hit in the back, a second pencil stands in for a car that hit the first car in the back, a third pencil stands in for a car that hit both other cars from the side, and a paperclip stands in for a dog that ran across the street, causing the first car to hit the second. Pencils and paperclips have neither been designed nor selected to function as stand-ins for other objects, but during improvisation we may attribute cognitive functions to a set of initially non-cognitive artefacts.

In some cases we improvise cognitive artefacts with more consolidated system functions, i.e. system functions that are more entrenched than mere one-offs. Beach's (1988) study of bartenders who structure distinctively shaped drink glasses such that they correspond to the order of the drinks is a good example. Due to this improvised use of drink glasses, bartenders do not have to remember the order of the drinks, but offload it onto their work environment. Drink glasses are intended by their designers and selected by their users to contain their drink, not as mnemonic aids. However, this particular mnemonic use of drink glasses is (relatively) widespread practice for bar tenders, which are a small cultural group. This mnemonic use of drink glasses may, therefore, be seen as an *idiosyncratic ongoing cognitive system function*.

In other cases we improvise cognitive artefacts with even more consolidated system functions. Most people intentionally put everyday artefacts in unusual locations such that they function as reminders. Leaving a rented DVD on your desk as a reminder to bring it back to the video store is a case in point. Such improvised mnemonic uses of everyday artefacts are quite common, e.g. leaving an empty milk bottle on the kitchen dresser as a reminder to buy milk, putting an article one has to read on top of the pile on one's desk, or tying a string around one's finger as a reminder for some action or event. In these cases, the location of the artefact is deliberately unusual such that it prompts a memory when the artefact is encountered in that location, thereby functioning as external memory. In other cases, artefacts are put in a location that is deliberately usual. Some people always put their car keys on their hall table such that it is part of their behavioural routines, ensuring that they do not forget where they have put their car keys. Such improvised uses of artefacts are widespread cultural practice and may, therefore, be referred to as *standardized ongoing cognitive system functions*.

The above examples show that we opportunistically use artefactual objects and structures for cognitive purposes, thereby improvising a variety of cognitive artefacts (see also Dahlbäck, Kristiansson, and Stjernberg 2013). Cognitive artefacts are thus neither defined by intrinsic properties of the artefact nor by the intentions of the designer, but by their function, which is established by the intentions of the user and by how it is used (see also Section 3). Due to these improvised uses of initially non-cognitive artefacts, we should conceive of the functional kind of cognitive artefacts as more inclusive than merely proper cognitive artefacts. In order to increase explanatory scope and to better understand a larger set of cognitive artefacts, we need to look at those with proper and system functions. From a user-centred perspective, it does not matter whether cognitive functions have been selected over time or improvised on the spot. What matters is that the object in question can perform a cognitive function and aid on-board cognitive capacities.

2.3. Multiple usability of structure and multiple realizability of function

Cognitive system functions vary between one-offs for individuals and widespread uses for a community of users. The examples I have used should, therefore, be located on a continuum and may, depending how often they are used for cognitive purposes, shift towards one of the extremes on the continuum. Furthermore, some of these system functions are ongoing additions to the proper function(s) of the artefact. Drink glasses, DVDs, and car keys are still used for their proper function, whereas napkins, empty milk bottles, and strings are usually thrown away after they have fulfilled their system function. Conversely, proper cognitive artefacts can also be used for other functions. I may, for example, put some money between the pages of a textbook in my bookcase so that no one can find it, I may use a rolled-up newspaper as a flyswatter, a ruler to homogenize paint in a newly opened can, an abacus as a kids' toy, or some books to support my computer screen. Therefore, like most artefacts, (proper) cognitive artefacts are multiply usable, because their physical structure affords more than one particular use.

So we can use a single (proper) cognitive artefact for a variety of functions, but we can also use a variety of (proper) cognitive artefacts for a single function. If I need to perform a difficult calculation, there are different strategies and artefacts I can deploy: I may use pen, paper, and external numerals, an abacus, a pocket calculator, a spreadsheet program, or a slide ruler. These artefacts have rather different physical and informational structures, but on a course-grained (or macro-)functional level of abstraction, they have the same function. If we were to ask: What are these devices for? Then the answer would most likely be: To perform calculations or to help us perform calculations. So it may be claimed that their macro-function is to (help us) calculate. Hence, cognitive functions are, at least on a macro-functional level of abstraction, multiply realizable, i.e. different physical and informational structures can be used for achieving the same goal.

However, on a fine-grained (or micro-)functional level of abstraction, there are differences in how these devices perform their function. When using pen, paper, and numerals, for example, most of the computation or information processing is performed by an embodied brain. When breaking down difficult calculations into easier ones, e.g. 3×7 or 5×5 , some people may just remember the outcome rather than to actually compute it, while others may perform the easier calculations in their head. Either way, information is processed (remembered or calculated) by a human agent and not by an artefact. The external numerals function as to complement working memory and to structure the task space by decomposing the task into smaller and easier ones to perform parts. When using an abacus or slide ruler, the computation is done by an agent-plus-artefact system, i.e. manipulating the beads or slides *is* computation. The artefact's function, then, is to facilitate an analogue computation performed by an agent-plus-artefact system. Finally, when using a pocket calculator or spreadsheet program, an agent merely provides the artefact with input and the computation is performed by the artefact. On a micro-level, the artefact's function is not to complement working memory, structure the task space, or to facilitate an analogue computation, but to perform a digital computation.

In these scenarios, the artefactual elements have different micro-functions. In the first scenario, there is mainly mental computation and the artefactual element is merely a medium for information storage; in the second scenario, there is agent-driven analogue computation performed by a joint agent-artefact system; and in the third scenario, there is mainly digital computation performed by a computer. Moreover, using these different artefacts also requires a different set of interactive skills and, consequently, the overall functional organization of the situated cognitive system may be quite different. But, although

different artefacts may have different micro-functions, may require different skills, and result in different overall functional architectures of agents and artefacts; ultimately, the situated cognitive system has as its function to perform calculations. Thus, on a systems level, functions can be seen as multiply realizable (see also Clark 2008).

3. Structure–function relations

In connection to structure–function relations, Kroes and Meijers (2006, 2) point out that “Technical artefacts can be said to have a dual nature: they are (i) designed physical structures, which realize (ii) functions, which refer to human intentionality.” Physical structures are often said to be mind-independent and can be described by the laws of physics, whereas artefact functions are mind-dependent (i.e. they require for their existence human intentionality) and thus require an intentional description. So, in order to properly describe artefacts, Kroes and Meijers argue, we need to somehow combine physical and intentional descriptions. Vermaas and Houkes (2006) argue that the notion of artefact function is helpful here because it is a “conceptual drawbridge” between the physical and intentional realms. Functions tie the physical and intentional realms together, i.e. functions necessarily need both physical structures and intentional human agents that design, select, improvise, and interact with those physical structures. In this sense, functions can be seen as emergent properties of intentional agents interacting with human-made physical structures that have a particular effect. Drawing on this insight, a model of the emergence of cognitive functions can be sketched.

Figure 1 presents a simplified model of cognitive function as an emergent property of the interaction between intentional, embodied agents, and cognitive artefacts. To briefly illustrate this model, I use the example of navigating with a map. We interact bodily with the physical structure of the map, for example, orientating it such that the information it contains becomes available to our perceptual systems. Such epistemic actions have as their goal to make available task-relevant information (Kirsh and Maglio 1994; Clark and Chalmers 1998; compare Loader 2012). The representational structure of the map contains a large amount of information, but only a small part of it is used to perform some cognitive task. A map of Sydney’s central business district, for example, may contain an elaborate representation of all the streets, parks, landmarks, train stations, bus stops, and so on. But to navigate, only a relatively small part of all the available information is relevant. The information that is actually used is the task-relevant information. This information is perceived and then processed by internal systems, either to guide a pragmatic action (e.g. walking towards the Harbour Bridge) or to guide further epistemic action




Figure 1. Model of interaction between intentional agents and cognitive artefacts.

(e.g. reorientating the map). Cognitive functions, then, emerge out of intentional interactions with physical, information-bearing structures.

Because functions define artefacts and tie the physical and intentional realms together, it is important to better understand how they are established. One way to do this is by taking a closer look at the following relation:

Physical structure \rightarrow representational structure \rightarrow information \rightarrow cognitive function. (1)

The remainder of this article is concerned with unpacking this relationship. In the next three subsections, each arrow in Equation (1) will be discussed subsequently, starting with cognitive function and then working my way back towards physical structure. So first I briefly explain how cognitive functions supervene on information, then look at how information supervenes on representational structures, and end with conceptualizing the relation between representational structures and physical structures.

3.1. *Cognitive functions and information*

The interaction model described in [Figure 1](#) shows that cognitive functions are established only when an artefact exhibits information that is used in performing some cognitive task. If the artefact does not exhibit task-relevant information, then it does not have a cognitive function. So, I claim, exhibiting task-relevant information is both a necessary and sufficient condition for being a cognitive artefact. Maps, for example, have cognitive functions as the information they provide is used in navigating; timetables have cognitive functions as the information they provide is used in decision-making; diagrams and models have cognitive functions as the information they provide is used in making inferences and reasoning; abacuses have cognitive functions as the information they provide is used in calculating; and so on. Cognitive functions thus supervene on information.

3.2. *Information and representational structure*

In Section 2, it was shown that cognitive functions may either be selected or improvised, resulting in a rather heterogeneous class of artefacts. This class of artefacts can be said to be functionally homogeneous in that they all complement cognition, but informationally heterogeneous in that they exhibit different informational properties. For reasons of space, the focus in this article is on representational artefacts, but it is important to point out that artefacts can also complement cognition by exhibiting non-representational information ([Kirsh 1995](#); [Heersmink 2013](#)). Before identifying different ways in which information supervenes on representational structures, a brief explanation of the notion of representation is helpful. Representational systems have three elements: (1) human agents interpreting a (2) representation about some (3) target. A defining property of a representation is that it stands-in for something else and what it stands-in for is referred to as its target. James Watson and Francis Crick's scale model of DNA, for example, stands-in for the structure of actual DNA. So here we have Watson and Crick as interpreting agents, their scale model as a representation, and actual DNA as its target. Following [Peirce \(1935\)](#), we can distinguish between three types of representations (or signs in his terminology): icons, indices, and symbols. These are distinguished on the basis of their particular representational properties, i.e. on how their content and meaning is established.

3.2.1. *Icons*

Icons obtain their informational content through exhibiting relevant isomorphism to their target. Peirce (1935, 362) writes: “I call a sign which stands for something merely because it resembles it an icon.” Two kinds of isomorphism may be identified: structural and sequential. Maps, scale models, blueprints, and radar systems are examples of structurally isomorphic icons, because their representational structure is isomorphic to their target. When icons exhibit sequential isomorphism, they represent the sequence of steps in a process or mechanism. A diagram depicting the order of the steps in a biochemical process – for example, $\text{DNA} \rightarrow \text{mRNA} \rightarrow \text{protein}$ – is sequentially isomorphic in that it represents not some physical state but the sequence of steps in a process.

Maps, radar systems, and other icons are intended by their designers and selected by their users to function as icons. They may, therefore, be seen as proper icons, as they have a history of cultural selection. There are, however, also improvised or system icons. Consider again Norman’s (1993) example of trying to explain and reconstruct how an accident happened with the aid of pencils and paperclips as stand-ins for the objects (cars and dog) they represent. This example is both a sequentially isomorphic icon, as it represents the order of sequences in an event, and a structurally isomorphic icon, as it represents the locations of the targets. System icons can over time become consolidated into proper icons. Watson and Crick, for example, improvised a scale model of the structure of DNA with cardboard cut-outs representing the base pairs in DNA and metal clamps to represent the phosphate backbone. Their model began as a system icon, but it has been quite successful in re-production and is found in many biology textbooks. It is fair to say that it is now a proper icon.

3.2.2. *Indices*

Indices have a direct causal connection to their targets. Peirce (1935, 248) describes an index as follows: “a sign which refers to the object that it denotes by virtue of being really effected by that object”. Important for indices is that the target has to causally effect the representation. A thermometer, for example, is directly connected to the temperature. If the temperature changes, then the reading on the thermometer changes as well. Similar indexical relations are established in compasses, pH meters, weathervanes, speed dials, spectrometers, barometers, and other artefacts that are directly connected to their target. These are all proper indices, as they are intended by their designers and selected by their users to function as indices. System indices also exist, but are less common than system icons, because most indices have a rather complex physical structure which is hard, though perhaps not impossible, to improvise. It may, for example, be possible to improvise a weathervane out of simple materials.

3.2.3. *Symbols*

Symbols acquire their meaning and content through shared use, agreement, and logical rules. Peirce’s (1935, 249) description of a symbol is “a sign which refers to the object that it denotes by virtue of a law, usually an association of general ideas, which operates to cause the symbol to be interpreted as referring to its object”. Examples include natural and artificial languages, scientific formulae, mathematical systems, and musical notation systems. The representational structure of symbols is often arbitrary, as there is nothing intrinsic in the structure (i.e. shape) of symbols that makes them represent some target.

For example, the structure of the word “tree” has no obvious relation to an actual tree and the structure of the symbol for wavelength “ λ ” has no obvious relation to actual wavelength. There is no isomorphism or direct causal connection between symbols and their target. It is shared use, agreement, and logical rules that establish their informational content. Given these social or cultural properties for determining the meaning of symbols, they are more malleable when compared with icons and indices: we may change their meaning through social agreement. This happens with words as new dictionaries sometimes adjust the meaning of existing words.

Other representations sometimes have symbolic properties as well. For instance, the reading of a thermometer (an index) is either measured in degrees Celsius or Fahrenheit. Either way, degrees are quantified in numbers, which are typical symbols, as they acquire their meaning through shared use, agreement, and logical rules. Likewise, 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 actual target. A representation may thus display a mixture of iconic, indexical, or symbolic properties. Therefore, a useful way of conceiving Peirce’s trichotomy is by seeing a representation as predominantly iconic, indexical, or symbolic (Atkin 2008).

Most symbols have a fairly long history of cultural selection and may, therefore, be seen as proper symbols. Because the meaning of symbols is defined in virtue of shared use, agreement, and logical rules, it is relatively easy to develop improvised or system symbols. We may imagine a scientist, designer, or some other individual inventing new symbols or symbolic systems, e.g. a symbol for a newly discovered phenomenon, a new computer programming language, or a new musical notation system. Initially, these new symbols only have system functions, as they lack a history of cultural selection, but depending on how often they are used and how successful they are in spreading to other cultural groups, they may become consolidated into proper symbols. This is arguably what happened to most current proper icons, indices, and symbols. The notion of system function is thus helpful in explaining the development of representational systems that at some point must have started as improvised representational structures.

3.3. *Representational structure and physical structure*

In the previous subsection, I looked at the relation between representational structure and information. We have seen that information obtains its meaning by being isomorphic to its target (iconic), by having a causal connection to its target (indexical), by shared use, agreement, and logical rules (symbolic), or by a combination of these properties. In this subsection, I focus on the relation between representational structure and physical structure and identify three distinct ways in which representational and physical structures relate, which are additive, constitutive, or computational.

3.3.1. *Additive relations*

Representational structures can be carried by or added onto a physical structure. Icons and symbols are often carried by paper, whiteboards, screens, or some other material. A map of Sydney’s central business district, for example, is printed onto a piece of paper, language is written down or printed onto a piece of paper, lecture notes are written onto a whiteboard or projected onto a screen, and the metric scale and numerals are painted onto a ruler. In such cases, physical materials such as paper, plastic, or some other material carry a representational structure. Other materials have (historically) also been used to carry external

representations, including rock, papyrus, clay tablets, wax tablets, wood, animal skin, canvas, and even sand. What is important here is that the physical structure of the artefact needs to be such that it can sustain an external representation for a certain period of time. This depends on the cognitive purpose for which the artefact is deployed. Writing with one's finger in the sand on a beach may work for creating a simple drawing, explaining where the waves break, but not for making a detailed enduring architectural blueprint which needs a fairly enduring material such as paper. In the above mentioned examples, the relation between the physical and representational structure is one of carrying and, therefore, the structure of the artefact does not directly influence informational content, only the structure (i.e. shape) of the icons and symbols are relevant for informational content and thus cognitive function.

3.3.2. *Constitutive relations*

Representational structures can be directly constituted by physical structures. When this happens, (part of) the physical structure of the artefact is identical to its representational structure. There are two ways in which this can happen: statically or dynamically. Designers sometimes make scale models (i.e. icons) of the objects or structures they are designing, for example, to test the aesthetic value, aerodynamics, or physical strength of their design. Scale models have a physical structure that is identical to their representational structure. There are usually no additional representations added onto a physical structure, as is the case with, e.g. maps, notebooks, textbooks, lecture notes, etc. However, in certain scale models, for instance those that are made to test the aesthetic value, there is colour added onto their physical structure, in which case there is an additional representational structure added to a physical structure. But, if that does not happen, then the physical form or structure of the artefact is identical to its representational structure on which its task-relevant information and function supervene. This is statically constituted because its physical and thus representational structure does not change. After the scale model has been made, its structure usually remains unaltered, unless it concerns a virtual simulation.

It is perhaps helpful to briefly point out that the relationship between representational structure and information is multiply realizable. Compare, for example, a detailed blueprint and an accurate scale model of a token building. Let us assume that these contain exactly the same information about their target, i.e. size, form, ratios, colour, structural composition, and so on, presented in different representational formats. Although these icons may look different and from a user-centred perspective we may experience them differently, strictly speaking they do contain the same information. There are, however, differences in affordances of the different formats, which may have informational and functional consequences. Similarly, Simon (1978) argued that two different representational formats, for example, propositional (symbolic) and diagrammatic (iconic), might be informationally equivalent, but not computationally equivalent. So they can contain the same information, but afford different (kinds of) computations. The format that works best depends on the task.

In case of indices such as thermometers, barometers, and compasses, their physical structure and therefore also their representational content is dynamically coupled to their target. For example, the physical structure of a mercury thermometer is causally and dynamically coupled to the temperature. In other words, mercury expands when temperature increases, which under normal conditions constitutes a linear relationship between temperature and the degree of expansion. In this case, the physical structure and state of the artefact

(i.e. the diameter of the column and the particular expansion properties of mercury) is partly identical to its informational content, i.e. the temperature. Thermometers also contain a static temperature scale such that one can precisely see what temperature it is. Jointly, a static representational structure (i.e. the temperature scale which is carried by some material) and a dynamic physical structure (i.e. mercury in a column) constitute the informational content of the artefact. Similar relationships between physical structure and informational content are established in barometers and compasses, where expansion properties of an alloy partly determine the informational content of a barometer (i.e. atmospheric pressure) and magnetic properties of a compass arrow partly determines informational content of a compass (i.e. the cardinal directions). If the properties of the physical material were even slightly different, then the informational content would be different as well. Thus, the target causally changes the physical structure and state of the artefact, which, together with a static representational structure, constitutes the informational content of the artefact.

3.3.3. *Computational relations*

Information can be computed or manipulated by a physical structure, which happens in both analogue and digital computational artefacts. Such artefacts have a physical structure that affords information processing or computation. A slide ruler, for example, is a mechanical analogue computer used mainly for multiplication and division and, to a lesser extent, for calculating roots, logarithms, and trigonometric functions. Slide rulers contain a set of static logarithmic scales that can be manually manipulated such that a mark on the sliding strip is aligned with a mark on the fixed strip. The relative positions of other marks on the strips are then observed. Numbers aligned with the marks give the approximate answer to the calculation. By manipulating the physical structure of the slide ruler, one automatically manipulates its representational structures and thus its informational content, in that way performing analogue computations.

Analogue computational artefacts are often designed such that they can perform one type of computation (e.g. mathematical computations) often consisting of one type of representations (e.g. numerals). Their physical structure severely limits how and what kind of information can be manipulated and processed. There is, for example, only one way a slide ruler can be manipulated, so they are not general-purpose machines. By contrast, digital computational artefacts, particularly modern computers, have a much more complicated structure–function relation.

In analogue computers, there is usually a one-to-one relation between structure and cognitive function, i.e. a particular physical structure can only perform one kind of function, e.g. mathematical computations. By contrast, digital computers exhibit a one-to-many relation between structure and cognitive function, i.e. a particular physical structure can perform many kinds of functions. Their functional malleability comes from the fact that they can be (re-)programmed and moulded to the users' needs, but also because they are a medium in which a variety of representational systems can be expressed and manipulated. The information that computers exhibit on their screens is highly dynamic and malleable and can be iconic (e.g. pictures), indexical (e.g. real-time weather radar), or symbolic (e.g. language). Computers also exhibit other types of information such as programming languages and software programs. These are typically symbolic, as they acquire their meaning from logical rules, but from a user-centred perspective, these do not really matter as they mainly happen inside the computer and are more relevant for computer programmers and software developers than for users. From a phenomenological user-centred

perspective, what matters is what happens on the screen, not inside the computer, because cognitive functions supervene on information that is only visible on the screen.

These three kinds of relations between physical and representational structure (i.e. adding, constituting, and computing) are not mutually exclusive, but quite often overlap. We have seen that a token representation may display a combination of iconic, indexical, or symbolic properties. It may, likewise, also display a combination of additive, constitutive, and computational properties. For example, a scale model may have paint added onto its physical structure (additive and constitutive), a thermometer has a static representational structure and a dynamic constitutive structure (additive and constitutive), or a computer can simulate a scale model (constitutive and computational). Consequently, a useful way of conceptualizing this trichotomy of relations between physical and representational structure is by seeing a token representational artefact as predominantly additive, constitutive, or computational.

To sum up this section, cognitive functions of artefacts supervene on task-relevant information. This information may be either selected over time or improvised on the spot and needs a representation as its supervenience base. Such representations may have a combination of iconic, indexical, or symbolic properties, all of which need physical structures to exist. There are different ways in which representations depend on physical structures for their existence: representations may be added onto a physical structure, they may be constituted by a physical structure, or computed by a physical structure.

4. Conclusion

This article has conceptualized some of the metaphysical properties of cognitive artefacts, thereby strengthening the ties between situated cognition theory and analytic philosophy of technology. It argued that cognitive artefacts are defined and demarcated by their function, which is to provide task-relevant information, in that way complementing internal information-processing systems and making cognitive tasks easier, faster, more reliable, or possible at all. It developed a pluralist notion of functional kind, which included artefacts with proper (selected) and system (improvised) functions. Building on the “dual nature of artefacts thesis”, it claimed that cognitive functions can be seen as emergent properties of intentional, embodied agents interacting with human-made physical structures, having a particular cognitive effect on the agent. Drawing on this insight, a model of the emergence of cognitive function was sketched. Finally, given the centrality of function, this article has unpacked the relation between physical structure, representational structure, information, and cognitive function. Better understanding this relation contains the key to better understanding cognitive artefacts and their functions.

Acknowledgements

I wish to thank John Sutton, Richard Menary, and two anonymous reviewers for their helpful comments and suggestions on an earlier version of this article.

Notes on contributor

Richard Heersmink is a PhD student at the Centre for Cognition and its Disorders at Macquarie University in Sydney, working on situated cognition theory. He is particularly interested in the relation between cognition and artefacts. He tries to better understand this relation by drawing on theories and concepts from philosophy of cognitive science, philosophy of science, and philosophy of technology.

References

- Atkin, A. 2008. "Icon, Index, and Symbol." In *The Cambridge Encyclopaedia of Language Sciences*, edited by P. Hogan, 367–368. Cambridge, MA: CUP.
- Bakker, L. R. 2004. "The Ontology of Artifacts." *Philosophical Explorations* 7 (2): 99–111.
- Beach, K. 1988. "The Role of External Mnemonic Symbols in Acquiring an Occupation." In *Practical Aspects of Memory*, Vol. 1, edited by M. M. Gruneberg and R. N. Sykes, 342–346. New York: Wiley.
- Brey, P., and J. H. Søraker. 2009. "Philosophy of Computing and Information Technology." In *Philosophy of Technology and Engineering Sciences*, edited by A. Meijers, 1341–1408. Amsterdam: Elsevier.
- Carrara, M., and P. Vermaas. 2009. "The Fine-Grained Metaphysics of Artifactual and Biological Functional Kinds." *Synthese* 169 (1): 125–143.
- Clark, A. 2003. *Natural-born Cyborg: Minds, Technology, and the Future of Human Intelligence*. Oxford: OUP.
- Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: OUP.
- Clark, A., and D. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 10–23.
- Cummins, R. 1975. "Functional Analysis." *The Journal of Philosophy* 72 (20): 741–764.
- Dahlbäck, N., M. Kristiansson, and F. Stjernberg. 2013. "Distributed Remembering Through Active Structuring of Activities and Environments." *Review of Philosophy and Psychology* 4 (1): 153–165.
- Franssen, M., P. Kroes, T. A. C. Reydon, and P. E. Vermaas, eds. 2014. *Artifacts Kinds: Ontology and the Human-made World*. Dordrecht: Springer.
- Heersmink, R. 2013. "A Taxonomy of Cognitive Artifacts: Function, Information, and Categories." *Review of Philosophy and Psychology* 4 (3): 465–481.
- Heersmink, R. 2014. "Dimensions of Integration in Embedded and Extended Cognitive Systems." *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-014-9355-1
- Houkes, W., P. Kroes, A. Meijers, and P. Vermaas. 2011. "Dual-nature and Collectivist Frameworks for Technical Artefacts: A Constructive Comparison." *Studies in History and Philosophy of Science Part A* 42 (1): 198–205.
- 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*. Dordrecht: Springer.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Hutchins, E. 1999. "Cognitive Artifacts." In *The MIT Encyclopedia of the Cognitive Sciences*, edited by R. A. Wilson and F. C. Keil, 126–128. Cambridge, MA: MIT Press.
- Kirsh, D. 1995. "The Intelligent Use of Space." *Artificial Intelligence* 72: 1–52.
- Kirsh, D. 2013. "Embodied Cognition and the Magical Future of Interaction Design." *Transactions on Human-Computer Interaction* 20 (1): 1–30.
- 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. A Philosophy of Engineering Design*. Dordrecht: Springer.
- Kroes, P., and A. Meijers. 2006. "The Dual Nature of Technical Artefacts." *Studies in the History and Philosophy of Science* 37 (1): 1–4.
- Loader, P. 2012. "The Epistemic/Pragmatic Distinction." *Philosophical Explorations* 15 (2): 219–232.
- Meijers, A. 2001. "The Relational Ontology of Technical Artifacts." In *The Empirical Turn in the Philosophy of Technology*, edited by P. Kroes and A. Meijers, 81–96. London: JAI.
- Menary, R. 2007. *Cognitive Integration: Mind and Cognition Unbounded*. Basingstoke: Palgrave Macmillan.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Norman, D. 1993. *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. New York: Basic Books.
- Peirce, C. S. 1935. *The Collected Papers of Charles S. Peirce*, Vol. 3. Cambridge, MA: Harvard University Press.

- 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*, edited by A. Meijers, 213–234. Amsterdam: Elsevier.
- Preston, B. 2013. *A Philosophy of Material Culture: Action, Function, and Mind*. New York: Routledge.
- Simon, H. 1978. "On the Forms of Mental Representation." In *Perception and Cognition: Issues in the Foundations of Psychology*, edited by C. W. Savage, 3–18. Minneapolis: University of Minnesota Press.
- Sutton, J. 2010. "Exograms and Interdisciplinarity: History, the Extended Mind and the Civilizing Process." In *The Extended Mind*, edited by R. Menary, 189–225. Cambridge, MA: MIT Press.
- Vaccari, A. 2013. "Artifact Dualism, Materiality, and the Hard Problem of Ontology: Some Critical Remarks on the Dual Nature of Technical Artifacts Program." *Philosophy & Technology* 26 (1): 7–29.
- Vaesen, K. 2011. "The Functional Bias of the Dual-nature of Technical Artefacts Program." *Studies in History and Philosophy of Science* 42 (1): 190–197.
- Vermaas, P., and W. Houkes. 2006. "Technical Functions: A Drawbridge between the Intentional and Structural Natures of Technical Artefacts." *Studies in History and Philosophy of Science Part A* 37 (1), 5–18.
- Wright, L. 1974. "Functions." *Philosophical Review* 82 (2): 139–168.