Mind and Artifact: A Multidimensional Matrix for Exploring Cognition-Artifact Relations

Richard Heersmink
Department of Cognitive Science
Macquarie University, Sydney, Australia
richard.heersmink@gmail.com

Abstract. What are the possible varieties of cognition-artifact relations, and which dimensions are relevant for exploring these varieties? This question is answered in two steps. First, three levels of functional and informational integration between human agent and cognitive artifact are distinguished. These levels are based on the degree of interactivity and direction of information flow, and range from monocausal and bicausal relations to continuous reciprocal causation. Second, a multidimensional framework for exploring cognition-artifact relations is sketched. The dimensions in the framework include reliability, durability, trust, procedural and representational transparency, individualization, bandwidth, speed of information flow, distribution of computation, and cognitive and artifactual transformation. Together, these dimensions constitute a multidimensional space in which particular cognition-artifact relations can be located. The higher a cognition-artifact relation scores on these dimensions, the more integration occurs, and the more tightly coupled the overall system is. It is then better, for explanatory reasons, to see agent and artifact as one cognitive system with a distributed informational architecture.

1 Introduction

There is a great variety in both the kinds of cognitive artifacts and the cognitive profiles of the human agents that use those artifacts. Due to this variety, a multiplicity of relations is established between agents and cognitive artifacts. One way to look at these relations is through the lens of extended mind theory (EMT), according to which some of these relations ought to be seen as constitutive. EMT argues that human cognition is in certain cases constituted by an embodied human brain and cognitive artifacts [1]. Consequently, cognitive artifacts are not seen as merely external aids or scaffolds for thinking, but are sometimes proper and constitutive parts of an extended or distributed cognitive process. Cognitive processes are thus conceptualized as hybrids or amalgamations of neurological, bodily, and environmental processes [2].

John Sutton has identified two movements or waves in EMT. The first wave is mostly based on the parity principle and is advocated by Andy Clark and David Chalmers [3], Mike Wheeler [4, 5], and others. The second wave is based on what Sutton calls the complementarity principle and is advocated by Sutton [6, 7], Richard Menary [8], Clark [9], and Julian Kiverstein and Mirko Farina [10]. The parity principle reads as follows: "If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process" [3, p. 12]. Thus, according to this principle, a cognitive process is extended when a cognitive artifact (or other part of the world) functions in a *similar* way as a clearly recognized internal cognitive process. So the parity principle invites us to see similarities between internal and external states and processes as a sufficient condition for cognitive extension.

The complementarity principle, by contrast, reads as follows: "In extended cognitive systems, external states and processes need not mimic or replicate the formats, dynamics or functions of inner states and processes. Rather, different components of the overall (enduring or temporary) system can play quite different roles and have different properties while coupling in collective and complementary contributions to flexible thinking and acting" [6, p. 194]. So, in contrast to parity, complementarity argues that cognitive artifacts need not have similar functions to internal processes, but often *complement* internal processes with *different* properties and functionalities. In fact, complementing brain functions is often the point of deploying cognitive artifacts: so that they can perform functions the brain cannot do or cannot do well. Jointly, brain-plus-artifact is a much more versatile and powerful cognitive or problem-solving system than the brain alone.

The parity principle stresses functional isomorphism, and downplays differences between internal and external states and processes, implying that the nature and properties of cognitive artifacts as well as their impact on our brains and behavior do not really matter. It also downplays individual differences between humans and how they interact with cognitive artifacts. There are clear differences in how different humans rely on and deploy cognitive artifacts. Some people remember their appointments, while others rely on their agenda, some people are good at mathematics, while others need pen and paper to solve calculations, et cetera. Such differences ought to be and are conceptually and empirically studied, but the parity principle does not encourage such interdisciplinary study of the variety of relationships between human agents and cognitive artifacts.

Contrary to parity-based EMT, complementarity-based EMT *does* encourage interdisciplinary study of different kinds of interaction and coupling. So rather than providing a new metaphysics of the mind, as Sutton remarks, it encourages detailed case studies in which the differences between the contributing elements are analysed. In this paper, I build on and further develop complementarity-based EMT by sketching a multidimensional matrix for better understanding and exploring the different kinds of epistemic interactions and coupling between humans and cognitive artifacts. The dimensions in the matrix include reliability, durability, trust, procedural and representational transparency, individualization, bandwidth and speed of information flow, distribution of computation, and cognitive and artifactual transformation. There are at least two reasons why we need such a multidimensional framework. First, because it encourages and provides a conceptual toolbox for a detailed interdisciplinary study of the variety of relationships between human agents and cognitive artifacts. Second, given that a substantial part of our cognitive activity quite heavily depends on cognitive artifacts, it is important to have a framework that gives us a richer and deeper understanding of the interactions with such artifacts. If we are "natural-born cyborgs" (i.e. inherently tool-using and tool-incorporating creates), then it is important to better understand the variety of relationships that are established between us and our cognitive tools.

Importantly, although motivated by EMT, this framework is not restricted to the extended mind cases in which a minimal requirement is a two-way or reciprocal interaction. Clark and Chalmers [3] have characterized extended minds as "coupled systems" in which there is a two-way interaction between an agent and cognitive artifact. Agent and artifact both play an active causal role in an overall cognitive process. Their notion of a coupled system requires two-way interaction between agent and artifact. But, it is still important to better understand one-way or monocausal relations. Indeed, a large amount of cognitive artifacts have a monocausal influence on human thought and behavior, so for explanatory reasons (i.e., explanatory scope and completeness) it would be unwise to exclude monocausal relations from the picture even if these cases are not candidates for extended cognition. In order to develop a more inclusive picture and to better understand the causality of information flow in cognition-artifact systems, I first say a few words on the concept of a cognitive artifact and then distinguish between three levels of information flow, including monocausal and bicausal relations as well as continuous reciprocal causation between human and artifact. Thereafter, I outline the multidimensional framework in which most of the dimensions are also helpful for better understanding monocausal interactions. I end with briefly applying the framework to two concrete examples.

2 Cognitive Artifacts

In an influential paper, Donald Norman has defined a cognitive artifact as "an artificial device designed to maintain, display, or operate upon information in order to serve a representational function and that affect human cognitive performance" [11, p. 17]. A cognitive artifact is thus a device that is intentionally designed to serve a representational function that has an influence on human thought. There is a wealth of such devices in our environment, including road signs, maps, diagrams, notebooks, thermometers, agendas, textbooks, smart phones, tablet computers, PowerPoint presentations, navigation systems, software programs, laptops, and desktops.

Importantly, the representational formats of cognitive artifacts need not always be linguistic, pictorial, numerical, or diagrammatic. For example, when I leave an empty milk bottle on the kitchen dresser to remind myself that I need to buy milk, one could argue that the milk bottle is a mnemonic aid. Although the representational format is not linguistic, pictorial, numerical, or diagrammatic, it still refers to a certain task. Likewise, in order to reduce their memory load, bartenders learn to use the shapes of drink glasses and their placement on a bar as a material representation for the sequence of the ordered drinks [12]. Here a distinction can be made between a designed and improvised representational function. Milk bottles and drink glasses have not been designed as mnemonic aids, however, during improvisation we can attribute cognitive or representational functions to initially non-cognitive artifacts. An important point here is that a cognitive artifact is not defined by some intrinsic properties, but because of the way it is used.

In some cases, even non-technological objects or structures serve representational functions. For example, some seafarers can navigate with the help of celestial objects such as the sun or stars. Seafarers, in such cases, do not navigate with the help of a cognitive artifact in the sense of a human-made device with a representational function like, for instance, a compass or radar system, but navigate with the help of a naturally occurring structure. The sun or stars, then, have a representational function which is attributed to them by a human agent or group of human agents.

Thus, as Edwin Hutchins informs us, "there is no widespread consensus on how to bound the category of "cognitive artifacts." The prototypical cases seem clear, but the category is surrounded by gray areas consisting of mental and social artifacts, physical patterns that are not objects, and opportunistic practices" [13, p. 127]. Given that we attribute representational functions to initially non-representational artifacts and naturally occurring structures, and perform actions on the basis of the information those artifacts and structures represent, I would like to propose a more liberal definition of cognitive artifacts as to include any object or structure (human-made or naturally occurring) with designed or improvised representational functions that affect human cognitive performance.

3 Levels of Information Flow

Having briefly looked at what a cognitive artifact is, let us now look at how information flows between humans and cognitive artifacts. I identify three levels of information flow, which should not be seen as clearly distinguished, but as overlapping. The first level is characterized by a monocausal or one-way information flow from artifact to agent. Examples include clocks, compasses, slide rulers, road signs, maps, dictionaries, encyclopedias, cookbooks, websites, documentaries, no-smoking signs, memorials, graphs, celestial objects, diagrams, manuals, and timetables. Humans make decisions and structure their actions on the basis of the information that such cognitive artifacts provide. We depart to the train station when our watch tells us it is time to go, we take a left turn because the map says it is the shortest route to our destination, and so forth. Further, and this is essential, the agent typically does not have any influence on the content and nature of the information. Such artifacts and the information they carry are designed, installed, and distributed by other human agents including writers, designers, publishers, news agencies, companies, governmental institutions, et cetera. Thus, one can argue that such cognitive artifacts and symbol systems *mediate* information flow between a designer and user[1].

In some cases, we interact bodily with the artifact in order to obtain the information we need or want, i.e., it is an interactive process. We interact bodily with compasses, slide rulers, maps, cookbooks, and manuals to get the information we desire, in which case there are three steps involved: bodily interaction, perceptual intake, and action[2]. In other cases, we merely have to look at the artifact to extract the relevant information. No-smoking signs, road signs, and memorials, for example, need not be interacted with bodily to obtain the information we want, in which case there are two steps involved: perceptual intake and action. However, not every deployment of a cognitive artifact results in an action. Occasionally, we are inhibited from performing an action. No-smoking signs and no-parking signs, for example, do not result in an action, but in an inhibition of an action (but only for those who would have otherwise smoked or parked, of course).

The second level is characterized by a bicausal or two-way information flow, that is, from agent to artifact and then from artifact to agent. Humans frequently offload information onto their environment to relieve their memory burdens, in that way creating cognitive artifacts such as post-it notes, notebook and agenda entries, shopping-lists, to-do lists, and lists of addresses, birthdays, and telephone numbers; but also artifacts with improvised representational functions like empty milk bottles and drink glasses. Artifacts in two-way relations are often tailored for individual use and are frequently not part of publicly available artifacts or symbol systems (e.g., road signs, clocks, and textbooks), although there are exceptions such as a shared agenda. They are closed systems in the sense that the cognitive artifact is meant for one individual agent who has designed the informational content of the artifact for individual use. In one-way systems, designers outside the system have designed both the physical structure and informational content of the cognitive artifact. But in two-way systems, designers have designed the underlying physical structure of the artifact, e.g. the structure of a post-it note, notebook, or agenda that enables a user to offload information onto the artifact. However, and this is essential, the informational content and representational function is designed by the user.

In terms of repeatability, there are different versions of two-way relations. On one side of the spectrum there are one-offs like, for example, a post-it note with a brief reminder such as the date of a deadline. This example is a one-off, as there is only one cycle of offloading, intake, and action. On the other side of the spectrum there are often repeated interactions with a single artifact as in the case of Otto and his notebook [3]. Otto writes important information in his notebook and then consults it to act on the basis of that information. He usually does not further manipulate the existing information in the notebook, but does on occasion add new information to it when he needs to do so. The content of new entries in the notebook usually does not depend on the content of previous entries. When Otto writes down the address of MoMa, it is because he knows that in the future he might be going to MoMa and for some external reason he is triggered to write down its address. It is not because other information in the notebook triggered him to do so. With post-it notes there is in most cases one cycle of offloading, intake, and action. But with Otto's notebook there are various distinct cycles of offloading, intake, and action, which are repeated over a certain period of time. However, the informational content of each cycle does not depend on the informational content of previous ones. Hence, Otto and his notebook form a two-way system, just one that is often repeated.

[1] There are cases in which the designer of the artifact also uses the artifact, for example the writer(s) of a dictionary, encyclopaedia, or cookbook may use the books to look for information and to structure their actions in which case there is a two-way relation established, see next subsection.

[2] Although the information flow between agent and artifact is monocausal, there is strictly speaking not a monocausal relation with the artifact itself. Rather, a bicausal relation is established, because we have to causally and bodily interact with the artifact to obtain the information we need or want.

The third level is based on a reciprocal information flow. Occasionally, cognitive artifacts are integral parts of ongoing information-processing tasks. Writing an academic paper [14], making a PowerPoint presentation, solving a difficult calculation with pen and paper [15], or designing an architectural blueprint of a building often involves small incremental steps. We do not have a finished paper, presentation, calculation, or architectural blueprint in our head and then fully offload it onto the artifact. Rather, we offload small bits of information onto the artifact, and the nature and content of the offloaded information contributes to and partly determines the next step in the overall process. For example, when writing an academic paper one often starts with a rough outline, which may prompt ideas about how to fill in the details. Filling in the details may then prompt an adjustment of the outline, which may in turn prompt further details. This process may continue for a number of cycles. Each step in the overall process builds and depends on previous steps. The human agent and cognitive artifact continuously exchange information and so there is a reciprocal and cumulative information flow that constantly transforms the agent-artifact system. There is, in Clark's words, "continuous reciprocal causation" between agent and artifact [9].

Like information flow in two-way systems, reciprocal information flow often takes place in a closed system in the sense that the cognitive artifact is meant for a single human agent who has designed the informational content of the artifact for individual use. In two-way relations, there were three steps involved: offloading, intake, and action. This is roughly the same for reciprocal relations, except that each cycle depends on the outcome of the previous one. The cycles are thus interdependent. For this reason, the functional and informational integration between agent and artifact is significantly closer than in one-way and two-way systems. It is not a mere exchange of information between two entities, as in two-way systems. What is offloaded onto the artifact in a given cycle depends on what is offloaded in the previous cycle(s) and, therefore, the degree of hybridization and integration is considerably higher. In fact, this integration is so dense that it is better to conceive of agent and artifact as one cognitive and information-processing system.

To be complete, there is a fourth level of information flow (which may be called system information flow) that, for reasons of scope, has not been outlined here. Information quite often flows in systems that are comprised of more than one human agent and more than one cognitive artifact. Examples include a team of engineers working on the design of a car, researchers in a scientific laboratory, and pilots in the cockpit of an airplane. These are cases of collective cognition in the sense that there is a (more or less integrated) collective that tries to solve a particular problem by using cognitive artifacts. Extended cognition and collective cognition in the sense just explained are related but distinct phenomena. The focus in this paper is on agent-artifact relations and extended cognition, so I only focus on the first three levels of information flow.

4 Dimensions for Agent-Artifact Coupling

Having identified these three levels of information flow in human-artifact relations, let us now continue with outlining the dimensions that are important for further exploring these levels. Sutton [16, 7], Clark and Robert Wilson [17], and Kim Sterelny [18] have articulated the idea of a dimensional analysis of agent-artifact relations. Clark and Wilson identify two dimensions: first, the nature of the non-neural resources, which may be natural, technological or socio-cultural; and second, the durability and reliability of the overall system. Sutton [7] takes the dimensions of reliability and durability as well as the dimensions of trust and glue mentioned earlier by Clark and Chalmers [3], and also briefly mentions the dimension of transparency. And finally, Sterelny discusses three dimensions, namely, those of trust, individualization, and individual versus collective use. Their dimensional frameworks are perceptive and insightful, but tend to emphasize certain dimensions while overlooking others.

In this section I aim to refine and synthesize their dimensions into a coherent and systematic multidimensional matrix, add a number of dimensions to the matrix, and briefly examine where and how some of these dimensions overlap and interact. Note that this matrix is not meant as an exhaustive list; there may be other dimensions relevant for better understanding cognition-artifact relations, and the dimensions are rather sketchy for reasons of space. Before outlining and discussing each dimension, it is helpful to distinguish a number of elements that are relevant for better understanding the underlying conceptual structure of each dimension. These elements are: (1) the cognitive profile or cognitive capacities of the human agent; (2) the representational, functional, and technical properties of the cognitive artifact; (3) the task environment and context of use; and (4) the kind of epistemic action and its epistemic purpose. Although essential for better understanding agent-artifact coupling, these elements are not dimensions, but each dimension emerges out of the interplay between two or more of these elements. These dimensions are all matters of degree and relational in the sense that they never depend on only one of those elements.

To give a brief example: the dimension of trust emerges out of a specific epistemic interaction between agent and artifact performed in a particular context and with a specific epistemic purpose in mind. Some artifacts like an authoritative textbook on some subject (say, astrophysics) are almost automatically trusted, while others like Wikipedia are trusted with much more care. So trust depends on the properties of the artifact, but also on one's cognitive attitude towards the artifact, which may differ from agent to agent. Some people may trust Wikipedia by default, while others are highly skeptical of its truth-value. Our cognitive attitude towards information also depends on the context. Libraries and universities, for example, are generally seen as contexts in which trustworthy information can be found. But the information provided by the ministry of truth in George Orwell's dystopian novel, 1984, is likely to be encountered with skepticism. The dystopian world that Orwell describes is a context in which information is distrusted because it is provided by a government that constantly gives misinformation. So whether a human agent trusts certain information depends on a number of elements, namely: the cognitive profile of the agent, properties of the artifact, context, and the purpose of the epistemic action. Let us now turn to the dimensions.

4.1 Reliable Access

Reliable access to external information is one of the key dimensions for how and how often an epistemic interaction unfolds [3, 17]. Several things are important here. First, the cognitive profile of the human agent partly determines the necessity for information access. Some people have bad memory capacities and therefore rely and depend more on memory aids such as notebooks, post-it notes, agendas, and other reminders. Other people have bad mathematical skills and rely and depend more on calculators or perform calculations with pen and paper. While yet other people have bad navigation skills and rely and depend on navigation aids such as road signs, maps, and navigation systems. There are also people who have better memory, mathematical, or navigational skills and do not or rely less on external artifacts.

Second, reliability depends on the kind and properties of the artifact. Due to their technical properties, some artifacts provide better information access than others. Take agendas, for instance. As long as one does not forget to bring one's analogue agenda when needed, it provides reliable access to the information in it. In contrast, digital agendas embedded in one's smart phone, tablet, or other electronic device, in one sense, provide less reliable access, because they are inaccessible without electricity. So one not only needs to remember to bring the device when needed, but also to charge it when the battery is empty. Further, digital cognitive artifacts can potentially malfunction in more ways than analogue ones. So next to battery issues, there may be numerous software and hardware issues that prevent one from accessing one's digital agenda. Software and hardware issues are irrelevant for analogue agendas. But, conversely, digital agendas such as Google Calendar are online systems that store information in the cloud and are therefore less susceptible for theft or loss than analogue agendas. Even if one looses one's wearable computing device or if it gets stolen, the information is still available in the cloud. Analogue agendas lack these properties.

Third, the context and kind of epistemic action are relevant for reliable information access. A carpenter only brings his slide ruler when he effectively needs it, which is during work. Carpenters only need access to slide rulers when they need to perform the epistemic action of measuring the length of some object. Such epistemic actions are frequently performed during work and thus in a work-environment. Carpenters presumably do not bring slide rulers to the supermarket or dinner parties, although there may be exceptions, because there is nothing for them to measure (unless they are working in a supermarket or at dinner parties, of course). So necessity of information access depends on the kind of epistemic action and context. Certain epistemic actions are therefore only performed in particular contexts.

4.2 Durability

There are two sides to durability: first, the durability of the artifact itself, and second, the durability of the relationship with the artifact. Certain cognitive artifacts are highly durable, while others are less durable. When handled carefully - textbooks, abacuses, and slide rulers can potentially last for decades, whereas analogue agendas last for roughly a year, and shopping-lists and to-do lists often last for just a few hours. This depends on both the material quality and properties of the artifact as well as the purpose of the epistemic action. Generally, the better the material quality of the artifact, the more durable it is.

But, more importantly, the durability and repeatability of our relationship with cognitive artifacts often depends on the kind of epistemic action (and its epistemic purpose) one performs with it. A shopping-list does not need to be very durable because after having bought the needed items, it has fulfilled its epistemic purpose. A computer, in contrast, *does* need to be durable because we need it for many kinds of epistemic actions for a long period of time. Wilson and Clark [16] introduce a trichotomy between one-offs, repeated, and permanent relationships with cognitive artifacts. Shopping-lists are one-offs. Abacuses or compasses, however, are frequently re-used because they are devices that are utilized many times for the purpose of calculating or navigating. But some cognitive artifacts enter into permanent and highly durable relationships with their users. Otto and his notebook, a carpenter and his slide ruler, and an astronomer and her telescope enter into long-lasting and interdependent relationships.

4.3 Trust

In George Orwell's dystopian novel, *1984*, the ministry of truth continuously updates and changes information in entertainment, news media, and educational books with the purpose of rewriting history so that it fits the party's political doctrines. In addition to constant misinformation, people are persistently being monitored by Big Brother and have therefore no privacy. So, if they are rational, people in this fictional world distrust (or ought to distrust) information that is provided and controlled by their government and should be very careful with what they write, publish, and distribute. Fortunately, in our non-fictional (Western) world things are better for at least two reasons. First, ideally we have freedom of press, freedom of speech, and freedom of information and thus control over the informational contents of our media and books. Second, we are not constantly being monitored by our government (although sometimes we are) and can thus write down, publish, and distribute whatever we desire. Freedom of information and informational privacy are two essential conditions for trust.

But there are other reasons for trust. Some information we trust because we have endorsed it somewhere in the past and wrote it down because of this. This is true for Otto's notebook, agenda entries, shopping-lists, and the information on post-it notes. Other information we trust because many people rely on it for their actions. This is true for timetables of trains, dictionaries, encyclopedias, and maps, which are used and shared by many humans. Because these symbol systems are shared with many others, and many people rely on them for their actions, there is often no reason to think that they are false or incorrect [18, contrast 19]. But there are exceptions: Wikipedia, for instance, is used and shared by many people, but is in some cases still not particularly trustworthy.

In two-way and reciprocal systems, we trust the information because we have endorsed it in the past and because we offloaded it ourselves, but we also trust it because we believe the information is private and has not been tampered with. Consider a brief example: In Australia there is a TV commercial for smart phones in which a parent goes shopping with a shopping-list composed on a smart phone. The application is connected in real-time to the desktop at home where his son deliberately changes the digital shopping-list to include items he desires. So with new digital cognitive artifacts with networking abilities such as smart phones and tablets, informational privacy and security [20, 21] become increasingly important for trust in information. Privacy and security issues are less likely to emerge when using analogue shopping-lists, which are identifiable by means of one's handwriting [22]. So the nature and properties of the artifact partly determine how relevant informational privacy and security are for establishing a trust relation with the artifact and the information it carries.

4.4 Transparency

There are two types of transparency that are relevant for cognitive artifacts, namely, procedural and representational transparency. Embodied tools like bicycles, cars, hammers, and cricket bats transform the body schema. Body schemas are flexible as to incorporate tools into the sub-conscious representation of the body and its capabilities for action. Those tools, then, are not experienced as external objects with which one interacts, but one interacts with the environment through those tools [23]. When a tool is incorporated into the body schema of its user, it becomes transparent in use. We then no longer consciously need to think about how to interact with the tool, interaction goes smoothly and the tool withdraws from attention, i.e. it is transparent [24].

A similar thing happens with cognitive artifacts which I will call "procedural transparency". Procedural transparency [see also 25, 26] concerns the effortlessness and lack of conscious attention with which an agent deploys a cognitive artifact. Otto, for example, is so adapted and familiar to using his notebook that he will consult it automatically when he needs to do so. His perceptual-motor processes are proceduralized to such an extent that he does not consciously think about how to retrieve information from his notebook. So the retrieval process is not a two-step process in which Otto first believes that the address of MoMa is in his notebook and then looks up and interprets the information to form his second belief, namely, that MoMa is at 53rd street. It is a proceduralized and transparent process. In Clark's words: “the notebook has become transparent equipment for Otto, just as biological memory is for Inga" [26, p. 80]. Having a high level of procedural transparency needs substantial training and takes a considerable amount of time.

Representational transparency concerns the effortlessness with which an agent can interpret and understand external information. For example, in my neighborhood in Sydney there is a war memorial, *The El Alamein Fountain*, to remind us of the Australian soldiers that died in 1942 during the Second World War in El Alamein, Egypt. However, the memorial is a fountain and it is not immediately clear that it is meant to be a war memorial. Only after reading the commemorative plaque I understood what it is meant to represent. A fountain has very little, if any, representational isomorphism with war and casualties of war. So, for individuals who know the representational function of the fountain, it may evoke strong memories and emotions about the Second World War. Yet others who do not know its representational function, may perceive it as a mere aesthetic object and have no connection to what it represents. Thus, representational transparency is not an objective or intrinsic feature of cognitive artifacts, but partly depends on the cognitive profile and capacities of the interpreting agent. In contrast, *The Tomb of the Unknown Soldier* in Ottawa, Canada is functionally and representationally much more transparent, because it is comprised of a number of soldiers holding guns and has clearly and largely written "1914-1918" on a plaque placed under the soldiers. So for most people it is immediately clear that it is meant as a memorial for the First World War. Thus whether a memorial or other cognitive artifact fulfils its representational function partly depends on its representational transparency, which, in turn, partly depends on the degree of representational isomorphism.

4.5 Individualization & Entrenchment

Sterelny [18] has argued that certain cognitive artifacts are individualized and entrenched. For Sterelny, individualization is changing, adjusting, or fine-tuning the artifact such that its use is more effective and efficient. He argues that most of the books in his professional library are interchangeable, but some of them are massively individualized with underlining, highlighting, comments, and post-it notes. These adjustments essentially make sense to Sterelny and are less useful and valuable to others. Similarly, Otto's notebook is highly individualized and is useful only for Otto, although others may still be able to read the notebook, only Otto uses it to aid his memory and structure his actions. My tablet computer is fairly individualized: it has applications that I have downloaded and installed to fit my specific needs such as the weather forecast and train timetables for Sydney, and specific websites, documents, and books. But although it is individualized, most applications are still easily usable by others. In contrast, no-smoking signs, road signs, and library books are not individualized (and thus interchangeable) and accessible for most people. Individualization of cognitive artifacts often takes a certain period of time and highly individualized cognitive artifacts are in close equilibrium with the cognitive profile of their user.

Entrenchment of cognitive artifacts implies a close equilibrium between agent and artifact in which *both* have been transformed in order to ensure the best possible fit between agent and artifact. Sterelny acknowledges that his individualized books are not entrenched in the sense that his professional routines and habits have not been adjusted to those books in the same way as those books have been adjusted to Sterelny. So, he has individualized his books, but his books have not individualized him, or at least not sufficiently. But, according to Sterelny, there may still be cases of entrenchment concerning books. For a Locke scholar, Locke's oeuvre may have transformed the routines of the scholar in the same way as he or she has transformed Locke's oeuvre in the sense of highly individualizing his works by underlining, highlighting, comments, and so on. A more obvious and clear example of an entrenched cognitive artifact is Otto's notebook. The information in his notebook is only meant for Otto himself and is specifically geared to his needs and desires, so it is highly individualized, and his behavioral and cognitive routines are sculpted by his notebook, so it is entrenched as well.

4.6 Bandwidth

Like information flow in computer networks, information flow in human-artifact systems has a certain bandwidth, which is the amount of information that is exchanged per unit of time and depends on properties of both the agent and artifact. For example, a map of a city on which a particular route is outlined, potentially has a greater bandwidth than a linguistic description of the same route, because for most people it is easier and more effective to interpret a map, than to read a linguistic description of a given route. Similarly, a graph of the amount of carbon dioxide in the earth's atmosphere plotted against the time, a pie chart of the distribution of species in a given ecosystem, and an organization chart of the departments of a university make complex relationships between several items or variables clear and easily understandable. Graphs, pie charts, diagrams, and other illustrations transform an abstract relationship or problem space into a relatively easy to understand visual format. Explaining these relations in linguistic terms would in most cases be significantly more burdensome. In fact, this is often the point of using non-linguistic representations: to effortlessly and quickly convey information that would take much more time to explain in linguistic format. Common wisdom would say that a picture is worth more than a thousand words. Bandwidth also depends on the interpretative skills of the agent. Some agents can take more information onboard in a given amount of time than others.

4.7 Speed of Information Flow

The speed with which information flows depends (again) both on the representational properties of the artifact and the cognitive profile of its user. Some people read quickly, while others do not. Some people interpret a map in one glance, while others have to study it before they know where to navigate. Humans have thus different interpretative skills, which partly determine how fast information is taken onboard and processed. The degree of representational transparency is also relevant here. Some information is easier to interpret than other. The higher the representational transparency, the easier the information is to interpret, and the higher the speed of information flow. So speed of information flow depends, on the one hand, on the cognitive and interpretation skills of the human agent and, on the other hand, on the informational and representational nature of the cognitive artifact. But contextual factors such as background noise may also influence speed of information flow, since one's concentration and thus also one's ability to interpret information is influenced by it.

Conversely, the speed with which one offloads information onto an artifact is also important. Again, this depends on properties of both the agent and artifact. Certain devices have input methods that are more efficient than others. A desktop computer has a keyboard that is geared towards quick data input, a tablet has a virtual keyboard that is much less efficient, and a smartphone has a virtual keyboard as well, but one that is much smaller and thus significantly less efficient. Some computing devices have auditory input methods which are potentially much quicker than conventional methods, because most people can speak quicker than they can type or write. But equally relevant are the interactive skills of the agent. Some people write or type considerably quicker than others, which often depends on training and education.

4.8 Distribution of Computation

The degree to which each element in a human-artifact system contributes to solving a problem depends on the distribution of computation. Compare, for example, making a graph by way of pen and paper with making a graph by way of a spreadsheet program. Let's assume that both graphs are based on the same dataset, so the cognitive output (i.e., the graph) is the same, but the distribution of computation is different. In the first scenario, most computation is performed by the human agent, whereas in the second scenario most computation is performed by the spreadsheet program. In the latter case, we delegate most of the information-processing to the artifact. The distribution of computation is relevant for the nature and coupling of the system, but of course only for artifacts that have themselves information-processing abilities. In case of analogue or static cognitive artifacts, all information-processing is done by the human component and the artifactual component then merely functions as a medium for storage with its own representational properties.

4.9 Cognitive and Artifactual Transformation

As we have seen, body schemas are flexible as to incorporate tools into the sub-conscious representation of the body and its capabilities for action. Tool-use thus transforms the body schema. Likewise, the use of cognitive artifacts and other external symbol systems transform the representational and cognitive capacities of the human brain. Helen de Cruz [27], Menary [28], Clark [9] and Michael Kirchhoff [29], amongst others, have argued that external symbol systems transform the brain's representational capacities. During ontogenetic development we interact with public representational systems such as mathematics and language. By so doing, we soak up and learn to think in those representational systems and the brain takes on the representational properties of those systems.

Language and mathematics are examples of external symbol systems with which we interact substantially for a

long period of time, both phylogenetically and ontogenetically. In ontogeny we call this period education. A considerable amount of research has been done on the transformation affect of those systems on our brain and cognition. Other cognitive artifacts and symbol systems such as road signs, maps, computers, and design programs have presumably also a transformation affect on our representational and cognitive capacities. For example, after navigating a city with a map for a certain period of time, the interaction with the map and the city has changed our internal spatial representation of parts of the city. At a certain point, we no longer need the map to navigate and we have to a certain degree internalized the map. Likewise, interacting with computers for many hours a day probably transforms our neuronal structures and cognitive capacities. Engineers, for example, spend many hours a day designing objects and structures with design programs. It is not unlikely that after a certain period of training and practice their brains take on the representational properties of the program. Such transformations seem to be a consequence of long-term interaction with cognitive artifacts over ontogenetic time.

It is, however, not only the human component of the system that transforms its representational properties and capacities. The artifactual component transforms its representational properties too: cognitive artifacts are often not static and fixed but active and dynamic. The representational properties of post-it notes, slide rulers, and textbooks, for instance, are fairly stable and fixed, but smart phones, tablets, laptops, and other computing devices are very dynamic in their representations. We can transform and adjust their representational properties to our own needs and desires, and it is frequently because we act on those artifacts and the information they carry that they have dynamic and changing representations.

5 Dynamics in the Matrix

All these dimensions are matters of degree and relational in the sense that they emerge out of a specific epistemic interaction between agent and artifact performed in a particular context and with a specific epistemic purpose in mind. Importantly, they are not meant as necessary conditions for cognitive extension and thus do not provide a clear set of conditions to demarcate between cases of embedded and extended cognition. On my view, particular cognition-artifact relations merely populate a certain region in this multidimensional space; the higher a specific cognition-artifact relation scores on these dimensions, the more tightly coupled the system is and the closer it is integrated with the human cognitive system. Let me now briefly analyze two distinct cognition-artifact relations.

First, when using a map during a citytrip, information flows from artifact to agent and so a one-way system is established. Say the agent consistently brings the map on each day of the citytrip, access to the information is thus highly reliable. The durability of the relationship is as long as the duration of the citytrip. So it is not a one-off or a permanent relationship, but a repeated one. The amount of trust in the correctness of the map is high, since the map was provided by an official travel agency. The procedural and representational transparency increases during the citytrip. The more often the map is deployed, the easier it becomes to use and interpret it. Let's assume that the agent does not make notes on the map, so it is not individualized. Both the bandwidth and speed of information flow depend on the representational properties of the map and the interpretative skills of the agent and are likely to increase over time. Maps that are simple and easy to interpret have potentially a greater bandwidth and speed of information flow. The distribution of computation is such that the agent does all the information-processing, because the map is a mere medium for information storage. And finally, depending on how often the map is deployed, it will (slightly) transform the representational properties of the agent. It is not unlikely that after a couple of days of navigating the city with the map, the agent partly transformed her internal representation of the city. But the map itself is static and does not transform after use. Thus, given how it scores on the above dimensions, this cognition-artifact relation populates a region somewhere in the middle of the space.

Second, Otto and his notebook constitute a two-way system: Otto offloads information (e.g., addresses, phone numbers, notes, ideas, et cetera) onto the notebook and then retrieves it for later use. Otto heavily depends on his notebook to successfully get around in the world, he therefore always carries it with him and thus the information in it is reliable available. A permanent relation is established in the sense that he consistently uses the notebook over a long period of time (Alzheimer's disease can take over a decade). He automatically trusts the information in the notebook, because he has endorsed it somewhere in the past and wrote it down because of this, but also because it is extremely unlikely that people will deliberately tamper with the notebook. For Otto, the notebook is highly transparent, both procedurally and representationally. Otto's perceptual-motor processes are proceduralized to such an extent that he does not consciously think about how to use retrieve information from the notebook. And because the information is written and structured by Otto himself, he does not need to think about what it means. For example, the sentence "MoMa is at 53rd street" needs little, if any, conscious deliberation. The notebook is further deeply entrenched, i.e. Otto has personalized his notebook, which, in turn, has sculpted his cognitive and behavioral routines and capacities. The bandwidth is fairly high, since the offloaded language is likely (though not necessarily) to be geared towards easy intake. The offloading speed is relatively fast, because Otto writes in his notebook and writing is a fairly quick method for offloading information. The distribution of computation is such that Otto performs all the information-processing, since notebooks are mere analogue mediums for information storage. And finally, the notebook may not have deeply transformed Otto's representational capacities, but language (i.e. the representational medium in his notebook) in genera certainly has. Thus, given how it scores on the above dimensions, this cognition-artifact relation populates a higher region in the multidimensional space.

Further, existing relations can shift from one region to another. When a particular artifact is used for a longer period of time and it becomes gradually more individualized, transparent, and trustworthy, the relation between user and artifact becomes increasingly more integrated. As a result, the relation will shift to a higher region in the matrix. Highly individualized and entrenched cognitive artifacts are likely to maintain a stable relation with its user and, consequently, populate a given region in the space for a long period of time, but most relations are constantly shifting from one region to another. This is so because most cognition-artifact relations are very dynamic in nature, constantly changing their functional and representational properties, and renegotiating existing informational equilibriums.

For analytical purposes, I have discussed each dimension separately, but some of them overlap and interact. I shall now very briefly look at a number of these interactions. Reliable access and durability often result in individualization. The more often a certain cognitive artifact is used, the more likely it is that it will be individualized and perhaps in some cases even entrenched. But this need not be the case. There are often-used cognitive artifacts that are not individualized or entrenched such as clocks and speed dials. Individualization and entrenchment frequently result in cognitive transformation. Again, the more often we use a certain cognitive artifact, the more likely it is that the human brain takes on the representational properties of the artifact. This happens with language and mathematics, but also with maps, design programs, and perhaps with graphs, pie charts, diagrams and other illustrations as well. Individualization frequently causes both trust and procedural and representational transparency. Individualized cognitive artifacts, including agendas and notebooks, are designed by the user of the artifact and thus almost automatically trusted and transparent in use, as well as transparent in interpretation. We do not need to think about how to use such artifacts, and the information they carry is trusted because we wrote it down ourselves. And finally, representational transparency often results in a higher speed of information flow. The idea being that the easier information is to interpret and understand, the faster we can take it onboard and process it. There are more interactions between the dimensions, but these are the most obvious ones.

6 Conclusion

This paper first briefly discussed the concept of a cognitive artifact and then distinguished between three levels of functional and informational integration between human agents and cognitive artifacts, including monocausal and bicausal relations as well as continuous reciprocal causation. After that, a multidimensional framework for exploring cognition-artifact coupling was sketched. Collectively, the dimensions constitute a multidimensional space in which cognition-artifact relations can be located. The framework provides a toolbox for detailed studies of specific conceptual or empirical cases of the use of cognitive artifacts. The higher a cognition-artifact relationship scores on these dimensions, the higher a region in this space it will populate, in which case there is higher degree of integration.

Acknowledgements

I would like to express my gratitude to my supervisors John Sutton and Richard Menary as well as two anonymous reviewers for valuable and perceptive comments on an earlier version of this paper.

References

[1] R. Heersmink, 'Defending Extension Theory: A Response to Kiran and Verbeek', *Philosophy and Technology*, doi: 10.1007/s13347-011-0035-6, (2011).

[2] M. Rowlands, 'Extended Cognition and the Mark of the Cognitive', *Philosophical Psychology,* **22,** 1-19, (2009).

[3] A. Clark & D. Chalmers, 'The Extended Mind', *Analysis,* **58**, 7-19, (1998).

[4] M. Wheeler, In Defense of Extended Functionalism, In: *The Extended Mind,* R. Menary (Ed.), MIT Press, Cambridge, (2010).

[5] M. Wheeler, 'In Search of Clarity about Parity', *Philosophical Studies,* **152**: 417-425, (2011).

[6] J. Sutton, Exograms and Interdisciplinarity: History, the Extended Mind, and the Civilizing Process. In: *The Extended Mind,* pp. 189-225, R. Menary (Ed.), MIT Press, Cambridge, (2010).

[7] J. Sutton, C.B. Harris, P. Keil, & A.J. Barnier, 'The Psychology of Memory, Extended Cognition, and Socially Distributed Remembering', *Phenomenology and the Cognitive Sciences,* **9,** 521-560, (2010).

[8] R. Menary, *Cognitive Integration: Mind and Cognition Unbounded,* Palgrave McMillan, (2007).

[9] A. Clark, *Being There: Putting Brain, Body and World Together Again,* MIT Press, Cambridge, (1997).

[10] J. Kiverstein & M. Farina, 'Embraining Culture: Leaky Minds and Spongy Brains', *Teorema,* **32**, 35-53, (2011).

[11] D.A. Norman, Cognitive Artifacts, In: *Designing Interaction: Psychology at the Human-Computer Interface,* pp. 17-38, J.M. Carroll (Ed.), Cambridge University Press, (1991).

[12] K. Beach, The Role of External Mnemonic Symbols in Acquiring an Occupation, In: *Practical Aspects of Memory: Current Research and Issues Vol. 1,* pp. 342–346, M.M. Gruneberg & R.N. Sykes (Eds.), Wiley, New York, (1988).

[13] E. Hutchins, Cognitive Artifacts, In: *The MIT Encyclopedia of the Cognitive Sciences,* pp. 126-128, R.A. Wilson & F.C. Keil (Eds.), MIT Press, (1999).

[14] R. Menary, 'Writing as Thinking', *Language Sciences,* **29**, 621-632, (2007).

[15] J.L. McClelland, D.E. Rumelhart & G.E. Hinton, The Appeal of Parallel Distributed Processing, In: *Parallel Distributed Processing, Volume 2,* pp. 3-44, J.L. McClelland & D.E. Rumelhart (Eds.), MIT Press, (1986).

[16] J. Sutton, 'Distributed Cognition: Domains and Dimensions', *Pragmatics and Cognition,* **14**, 235-247, (2006).

[17] R. Wilson & A. Clark, How to Situate Cognition: Letting Nature Take its Course, In: *The Cambridge Handbook of Situated Cognition,* pp. 55-77, M. Aydede and P. Robbins (Eds.), Cambridge University Press, (2009).

[18] K. Sterelny, 'Minds: Extended or Scaffolded?', *Phenomenology of the Cognitive Sciences,* **9**, 465-481, (2010).

[19] K. Sterelny, Externalism, Epistemic Artefacts and the Extended Mind, In: *The Externalist Challenge. New Studies on Cognition and Intentionality,* pp. 239-254, R. Schantz (Ed.), De Gruyter, (2004).

[20] L. Floridi, 'The Ontological Interpretation of Informational Privacy', *Ethics and Information Technology,* **7**, 185–200, (2005).

[21] L. Floridi, 'Four Challenges for a Theory of Informational Privacy', *Ethics and Information Technology,* **8**, 109-119, (2006).

[22] M. Parsell, 'The Cognitive Cost of Extending an Evolutionary Mind into the Environment', *Cognitive Processing,* **7,** 3-10, (2007).

[23] R. Heersmink, 'Embodied Tools, Cognitive Tools, and Brain-Computer Interfaces', *Neuroethics,* doi: 10.1007/s12152-011-9136-2, (2011).

[24] M. Merleau-Ponty, *Phenomenology of Perception*, Routledge, (2003).

[25] J. Sutton, The Feel of the World: Exograms, Habits, and the Confusion of Types of Memory. In: *Philosophers on Memento,* pp. 65-86, A. Kania (Ed.), Routledge, (2009).

[26] A. Clark, *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, Oxford University Press, New York, (2008).

[27] H. de Cruz, 'An Extended Mind Perspective on Natural Number Representation', *Philosophical Psychology,* **21**, 475-490, (2008).

[28] R. Menary, 'Dimensions of Mind', *Phenomenology and the Cognitive Sciences,* **9**, 561-578, (2010).

[29] M.D. Kirchhoff, 'Extended Cognition and Fixed Properties: Steps to a Third-Wave Version of Extended Cognition', *Phenomenology and the Cognitive Sciences,* doi: 10.1007/s11097-011-9237-8,(2011).