

This is a preprint version of the following article:

Brey, P. (2005). 'The Epistemology and Ontology of Human-Computer Interaction,' *Minds and Machines* 15(3-4), 383-398. ISSN: 0924-6495.

The Epistemology and Ontology of Human-Computer Interaction

ABSTRACT

This paper analyzes epistemological and ontological dimensions of Human-Computer Interaction (HCI) through an analysis of the functions of computer systems in relation to their users. It is argued that the primary relation between humans and computer systems has historically been epistemic: computers are used as information-processing and problem-solving tools that extend human cognition, thereby creating hybrid cognitive systems consisting of a human processor and an artificial processor that process information in tandem. In this role, computer systems extend human cognition. Next, it is argued that in recent years, the epistemic relation between humans and computers has been supplemented by an ontic relation. Current computer systems are able to simulate virtual and social environments that extend the interactive possibilities found in the physical environment. This type of relationship is primarily ontic, and extends to objects and places that have a virtual ontology. Increasingly, computers are not just information devices, but portals to worlds that we inhabit. The aforementioned epistemic and ontic relationships are unique to information technology and distinguish human-computer relationships from other human-technology relationships.

Key words: cognitive artifacts, extended mind, human-computer interaction, social reality, virtuality

1. Introduction – What Kind of Tool is a Computer?

What kind of artifact is a computer? What is the functional relationship that we have to it? This question has always been difficult to answer because the computer is such a versatile tool. Computers have been called universal machines, machines that can execute an indefinite amount of different functions, and that can therefore function as very different tools for us at different times. So whereas the function of a screwdriver is to drive screws, and the function of a copy machine is to make copies, no straightforward account can be given of the function of a computer. It could be claimed that the function of a computer is to perform computations, to make calculations, or to process information, but this does not tell us much that is informative. Combustion engines, whose function it is to produce energy, are part of all kinds of machines. But it would not be helpful to say that the function of these machines is to produce energy. Similarly, computer hardware is teamed up with software and peripherals to create dedicated machines with specific functionalities. A better understanding of computers as human tools requires an understanding of these functionalities.

In this essay, it will be argued that contemporary computer systems perform two broad classes of functions: *epistemic functions* and *ontic functions*. Epistemic functions are what have traditionally been called information processing functions. In an epistemic function, the computer functions as a cognitive device that extends or supplements human cognitive functioning by performing information processing tasks. As long as they have existed, computers

have functioned as cognitive devices, or cognitive artifacts, and this epistemic function remains important in contemporary computers. I will argue, however, that in recent decades, the computer has acquired a new class of functions, which I term ontic. In their ontic role, computers simulate environments and tools to engage these environments. I will argue that this ontic function is not properly understood as an information function, even though it is dependent on the information processing capabilities of computers.

2. Cognitive Artifacts

Technological artifacts often serve to extend or augment existing human capacities or faculties (Brey, 2000). For example, microscopes and telescopes extend our vision, so that we can perceive objects or patterns that we could otherwise not perceive. Vehicles like bicycles and automobiles extend our locomotive abilities, so that we can move faster or with less effort. Tools like hammers and drills extend the ability of our hands to modify materials. Walls, heaters and air conditioners extend the thermoregulatory capacities of the human organism. Millions of other artifacts likewise extend perceptual, motor and regulatory functions of the human organism. Does computer technology likewise extend one or more of our faculties? According to Marshall McLuhan, it does. McLuhan claimed in his *Understanding Media* that with the advent of electric media, it is no longer just perception and motor functions of humans that are extended by technology. He argued that electric media extend the information processing functions of the central nervous system, taking over functions of information management, storage and retrieval normally performed by the central nervous system. He specifically argued that digital computers extend creative cognition and higher thought. McLuhan hence saw the digital computer as extending cognition, as opposed to perception or motor functions (McLuhan, 1966).

I here intend to develop McLuhan's idea that the computer extends human cognition by building on human cognitive capacities. My focus will be on the question *how* computers extend human cognition, which I intend to answer by analyzing the functional relation between human cognition and computer activity. I will be arguing that the computer is a special kind of *cognitive artifact* that is capable of extending a broad range of cognitive abilities of human beings. The notion of a cognitive artifact has been introduced by psychologist Donald Norman (1993). According to Norman, there is a special class of artifacts that are distinguished by their ability to represent, store, retrieve and manipulate information. Norman calls such artifacts cognitive artifacts. He defines them as artificial devices designed to maintain, display, or operate upon information in order to serve a representational function. The keywords here are "information" and "representation." They distinguish cognitive artifacts from other artifacts.

Norman's definition provides a clear criterion for distinguishing cognitive artifacts, such as thermometers, newspapers, clocks and Internet search engines, from noncognitive artifacts, such as hammers and automobiles. A thermometer has as its function is to inform us about temperatures. A newspaper has been made to store and displays information on current events. A clock has been designed to accurately represent and display the time. An Internet search engine has been designed helps us to find information on the Internet. All these functions are representational functions. A hammer, in contrast, does not normally serve a representational function, as it does not normally maintain, display or operate upon information. There are perhaps some peripheral ways in which it may still serve representational functions. For example, it may contain a symbol or language that informs who the manufacturer is. And it may be put on a coffee table at home to remind oneself about a carpentry job that needs finishing. In that case it serves a representational function by making an indexical reference to the carpentry job. But it is not designed for such a purpose and therefore these cognitive functions are peripheral to its primary functions which are to hit nails and flatten or shape materials. Hence, it

is not a cognitive artifact. Similarly, an architectural sketch that has been made to accurately represent a building is a cognitive artifact, whereas an artistic drawing of a nonexistent building is not a cognitive artifact, because it has not been designed to display information, but rather to please aesthetically.

Cognitive artifacts are properly called ‘cognitive’ because they, in quite straightforward ways, extend human cognition. They help us think, plan, solve, calculate, measure, know, categorize, identify, or remember. Various classes of cognitive artifacts may be distinguished, based on the primary cognitive capacity or capacities that they extend or aid. I will now list various basic cognitive abilities that have been recognized by cognitive psychologists, and illustrate how cognitive artifacts may extend or aid such abilities.

1. *Memory.*

Human memory is the psychological faculty by which we store information and retrieve it for later use. Cognitive artifacts that extend memory functions may be called *memory devices*. They are artifacts that help us encode, store and retrieve information. Sometimes, memory devices merely help us to locate information in our own memory. For example, some banks issue cards that help you to reconstruct the PIN-code of your ATM card based on an easier to remember verbal code. More often, memory devices serve as memory systems themselves: they store information in organized ways. If memory is a means for encoding, storing and retrieving information, then any device which has this as one of its primary functions is a memory device. So a notepad is a memory device, as its function is to store notes for ourselves or others, and pens and pencils are memory devices used for inscribing data into external memory.

Psychologist Merlin Donald (1991) has argued that one of the most important changes in the transition from Neolithic to modern culture is the emergence of a system of external memory storage, of which the storage of symbolic (linguistic) information is the most important. He claims that nowadays this external memory system contains more information than biological memories do, and that most human beings rely on it extensively. Media used for external memory storage include books, newspapers, microfilms, digital storage media, and others. For inscribing or reading them we have pens, pencils, microfiche readers, monitors and the like. Most important are paper and electronic (especially digital) storage devices. External memory devices serve in straightforward ways as extensions of human biological memory.

2. *Interpretation*

Interpretation is also a fundamental human cognitive ability. Interpretation is the ability to assign meanings to input data, through the assignment of one or more concepts or categories. For example, when one tries to recognize objects in one’s environment, one may perceive certain shapes and colors. To recognize what these shapes and colors stand for, one needs to apply concepts to them that make a ‘fit’. For example, a curved yellow shape can only be recognized as a banana when the concept of a banana is applied to it. The interpretation of perceptual data is the way in which perceptual stimuli are made useful as objects of conceptual thought, which does not range over sensory images, but over concepts.

Interpretation can be qualitative or quantitative. *Quantitative* interpretation is the assignment of a numerical value to a perceived quality. Another word for this is *measurement*. Measurement is a cognitive activity that we typically, though not invariably, perform with the aid of artifacts, *measuring devices*, like thermometers, spectrometers, clocks, yardsticks, sextants, etc. The history of science and technology, if not economics, politics and management, is to a large extent a history of measurement, along with the measuring devices that have been developed for it. Measuring devices extend our abilities to estimate the size, number or intensity of phenomena in the world, and are hence extensions of our ability to interpret the world.

Qualitative interpretation is the assignment to data of a qualitative concept or category.

There are many cognitive artifacts that aid in the qualitative interpretation by giving criteria, templates or examples for the application of a concept. For example, color charts aid in the correct identification of colors. A book on animal tracks, with drawings of typical animal tracks, helps one in the identification of tracks observed in the woods. Medical texts list criteria for the correct identification of diseases. Few artifacts exist, however, that do not just support qualitative interpretation but that do the interpretive work themselves. The digital computer is an artifact capable of autonomous interpretation. Most qualitative interpretation performed by computers takes symbolic inputs, such as sentences, numbers or names, and assigns categories to them. For example, a computer program may take names of animals and classify them as “reptile,” “mammal,” “bird,” “amphibian,” etc. Or it may take a sentence, and parse it by assigning grammatical roles to words and phrases. Computers are also capable, when suitably programmed, to recognize objects and scenes in pictures, although their capabilities to do this are more limited.

3. *Search*

When we interact with the world, we often actively look for things that we are trying to locate but have not observed yet. We constantly look around for people, pens, purses, stores, food, stamps, road signs, words, barcodes, and numerous other things that we need to see, locate or use. The ability to search and subsequently recognize things is one of our fundamental cognitive abilities. Searches sometimes take place with exact specifications of what you are looking for, but more often they are heuristic, and take place according to hypotheses: you assume that there is something in your vicinity that meets a set of loosely formulated criteria, and search for something that meets these criteria. Searches do not just take place in the external world; we also frequently search our own memories for information.

Search is a cognitive process, because it involves activities like mental scanning and pattern matching. It is another process that can be assisted by cognitive artifacts. Cognitive artifacts can aid search by structuring the search space in such a way that it can be more easily scanned, and by ‘flagging’ types of items that one may scan for (e.g., by marking them with colors or symbols). Examples of cognitive artifacts that aid search are labels and filing systems. A special ability of computer systems is that they can perform searches themselves. They can do so because of their ability to do pattern matching and their ability to systematically scan through a search space.

4. *Conceptual thought*

The most important cognitive ability that distinguishes human cognition from animal cognition is the ability to engage in conceptual thought, and particularly the ability to engage in abstract thought, using abstract concepts. Conceptual thought is the ability to arrive at new conceptual structures (ideas or beliefs) through the modification (analysis or synthesis) of existing ones. Conceptual thought often involves *problem solving*: it often involves cognitive goals like finding the solution to a mathematical equation, determining the best way to furnish a room, finding an adequate translation into English for a sentence in Spanish, or thinking up the most diplomatic answer to a potentially embarrassing question. Problem solving can be aided by cognitive artifacts that help to arrive at an accurate representation of the problem space or of the kinds of steps to take to find a solution, such as models and diagrams, and procedural manuals. Computer systems are, again, special in that they are capable of autonomous problem solving. When suitably programmed, computers are capable of solving equations, thinking up room designs, translating sentences from Spanish to English, or answering questions. Computer intelligence of course still has its limitations. Results are not impressive, for example, in the areas of language use and reasoning in informal domains. Nevertheless, computers are nowadays frequently used for all kinds of tasks that ordinarily require conceptual thought, whether they are performing calculations, correcting grammar, engaging in dialogue, planning distribution routes, or designing

copying machines.

A distinguishing feature of cognitive artifacts is that they do not just function as objects of cognition, like other structure in the world, but that they become integral components of the information processing task itself. Traditionally, cognitive scientists have located information processing tasks in the head; information processing, or cognition, is thought to be done by minds, and minds alone. Over the past twenty years, however, an alternative view has emerged, which holds that cognition often takes place in interaction with the environment, and is in effect distributed over minds and structures in the environment that have a function in cognitive tasks.

In cognitive science, this view has been developed as the *distributed cognition* approach (Hutchins, 1995, Perry, 2003, Salomon, 1993). This is an approach to the study of cognitive processes that emphasizes the distributed nature of cognitive phenomena across individuals, artifacts and internal and external representations. It was developed by Ed Hutchins and his colleagues at the University of California, San Diego in the late 1980s as a radical alternative to the traditional view of cognition, according to which the unit of analysis for the cognitive sciences is the individual, specifically the human mind or brain and its subsystems, conceived of as a processor of peripheral input from the sense organs and an initiator of action. The distributed cognition approach entertains a broader conception of cognition, in which the unit of analysis is any system with components that functionally contribute to the realization of a cognitive task. Such a system may be an airline cockpit with pilots, the bridge of a ship, or an individual using a calculator or measuring rod.

A philosophical version of this view of cognition, the *extended mind* view, has been put forward by Andy Clark and David Chalmers (Clark and Chalmers, 1998; Clark, 2001, forthcoming). They argue that humans sometimes perform physical actions in order to perform cognitive tasks, actions like rotating objects, making measurements, and looking up information. Following Kirsh and Maglio (1994), they call such actions *epistemic actions*. They claim that if such actions were performed in the head, they would be recognized as part of the cognitive process. They then argue that there is no principled reason to hold that epistemic actions are *not* parts of cognitive process, and conclude that those parts of the world that aid in cognitive processes should be held to be part of the cognitive process. Clark and Chalmers introduce the notion of a *coupled system*, which is the linking of a human being with an external entity in a two-way interaction that includes information input from this entity and epistemic actions towards it. They argue that coupled systems can be understood as genuine cognitive systems, because the entity is made part of the information processing task. Clark and Chalmers recognize that epistemic actions frequently involve special tools and instruments, but they do not employ Norman's notion of a cognitive artifact. In more recent publications Clark refers to what he calls "cognitive technology", by which he means external cognitive aids like pens and calculators (Clark, 2001), and in a further elaboration of the extended mind perspective, Kim Sterelny has recently coined the term "epistemic artefacts" (Sterelny, forthcoming).

Although I am largely in agreement with the extended mind view of Clark and Chalmers, I believe that their claim that their notion of a coupled system is too liberal. It seems to be a consequence of their view that when I merely look at a knife on the table, no coupled system exists because I do not perform epistemic actions towards the knife, whereas when I move the knife an inch to observe it better, I create a coupled system, meaning that the knife and I now form a cognitive system that collectively performs the cognitive task of observing the knife. The knife is certainly more an active component of the cognitive task in the second scenario than in the first, but only in the passive role of object of cognition. Cognitive artifacts are different: they are not just objects of cognition, but also bearers or processors of information, and in this sense more active contributors to cognitive tasks.

I would hence want to distinguish between *weakly coupled systems*, in which objects in

the environment become mere objects of epistemic actions, and *strongly coupled systems*, in which a relation is created between a human user and a cognitive artifact that actively contributes to an information processing task by serving a particular representational function. Only strongly coupled systems, in my analysis, qualify as genuine cognitive systems, because only in such systems do external objects serve representational and information-processing functions. It should be added that ordinary objects are sometimes *made* to function as cognitive artifacts, even if they have not been *designed* for this purpose. They can be granted cognitive functions through the imposition on them of a particular representational or computational function. Thus, if I know the length of my knife to be seven inches, I can use my knife *as* a cognitive artifact to measure the length of my table. In this case, the combination of me and the knife functions as a cognitive system that performs a measurement task.

3. Computer Systems as Cognitive Artifacts

Among the many cognitive artifacts that exist, computer systems are certainly unique. As has been observed in the previous section, computers are special in that they often go beyond the role of facilitating or aiding human cognition: computers are capable of performing cognitive tasks autonomously. Computers are special because they are capable of actively manipulating representations. Most other cognitive artifacts cannot manipulate representations, because they are not capable of systematically discriminating different kinds or representations and responding to them in meaningful ways. This capability is the reason that computer systems are the most versatile and powerful cognitive artifact, that can support or perform almost any cognitive task.

The functional relation that computers, as cognitive artifacts, have to their human users is hence that they extend cognition. Specifically, they extend the memory, interpretation, search, pattern matching and higher-order cognitive abilities of human beings. There is not, however, a single way in which computer systems functionally extend human cognition. I observed that computers are capable of autonomous cognitive processes. But they may also serve as a mere facilitator of human cognitive processes, as happens for example in word processing. I will now go on to further analyze how exactly computer systems add to, augment or replace the cognitive capacities of human beings.

As my point of departure, I will take a set of distinctions made in Brey (2000). In this essay, I argued that artifacts that amplify the functioning of human organs may maintain three different types of relations with these organs. An artifact may *replace* the functioning of an organ by performing the function of that organ in a way that makes the organ redundant. For example, when driving a car, one's legs are not used as a means for transportation. An artifact may also *supplement* an organ that it extends, by performing a function that the organ in question is also performing. For example, clothing adds to the protective and temperature control functions already performed by the skin. Third, an artifact may *enhance* the functional powers of the organ that it extends, not by independently performing a function that resembles the organ's function, but by cooperating with the organ in a way that enhances its activities, in this way engaging in a *symbiotic relationship* with the organ. For example, a telescope extends visual perception by teaming up with the eye to form a new functional unit consisting of telescope-plus-eye that is capable of doing things that neither the telescope nor the eye is capable of doing by itself.

The relevant faculty or 'organ' that is extended by computer systems is our faculty of cognition, located, according to neuroscience, in our brain, specifically in the neo-cortex. Is a computer system an artifact that mostly replaces, supplements or enhances human cognition? All three roles are visible in computer systems. In its early days, the computer was often called the 'electronic brain,' and a common fear was that computers would replace human brains as the primary locus of cognitive activity. The computer as a replacement of human cognition is an

autonomous information processing system that operates like a human cognitive agent, producing its own plans, solutions, and other knowledge structures without human intervention. In this role, the computer fits the early ideals of artificial intelligence research to ‘build a person,’ and the ideal of expert systems research to replace human experts.

The idea of the computer as a replacement of the human cognitive system has never been fully realized, and it is nowadays recognized that AI’s dream to ‘build a person’ still depends on significant breakthroughs in AI research that have not been realized in past decades. The idea of the computer of a supplement to human cognition, in contrast, was an idea already powerful in the early days of computer and one that still holds currency. The computer in its supplementary role does autonomous information processing, but remains limited to those tasks that are tedious, time-consuming, or error-prone when performed by humans. These are tasks like doing large calculations (‘number crunching’), database searches, and organizing and reformatting data. The implicit distribution of labor between humans and computers is then that humans perform the more intuitive and creative cognitive tasks and are responsible for the overall structure and goals of large cognitive tasks, whereas computer systems autonomously perform more tedious or time-consuming cognitive tasks that are defined as ‘subroutines’ within such larger cognitive tasks.

Since the rise of the personal computer, however, a third powerful interpretation of the role of the computer has emerged: that of a versatile tool that we handle directly and that enhances our own power to get work done. In this role, the computer is not an autonomous cognitive unit, but a cognitive aide, that enhances our own cognitive powers. It does not perform cognitive tasks by itself, but helps us to perform them. Our relation with the computer in this role is more *symbiotic*: the performance of a cognitive task depends on the information-processing abilities of both human and computer, and the exchange of information between them. When we use word processors, spreadsheets, web browsers, and other software tools, the cognitive tasks we perform, such as producing well-formatted documents, performing calculations, or navigating the Web, are performed in cooperation with the computer. When we check the spelling of a document with the aid of a spelling checker, for example, this cognitive task depends on both the ability of the spelling checker to identify possible misspellings, and our own ability to operate the spelling checker and to decide whether its proposed corrections are valid.

Even in their role as tool, however, computers still engage in autonomous information processing. The aforementioned spelling checker may not autonomously correct the spelling of a document, but it does make autonomous proposals. On the other hand, the computer in its role as a supplement to human cognition still requires a knowledgeable human operator, so its operations are not entirely autonomous. So the distinction between supplementary and enhancement roles of the computer is by no means absolute. In both cases, cognition is made into a distributed process that depends on the information-processing abilities of both humans and computers (Peschl and Stry, 1998). The mutual dependency is greatest, however, when the computer functions as an enhancement of human cognition. In these cases, the computer operates *in tandem* with the human mind, and the integration of cognitive functions becomes so great that human and computer are best regarded as a single cognitive unit, a *hybrid cognitive system* that is part human, part artificial, in which two semi-autonomous information-processing systems cooperate in performing cognitive tasks.

The notion of a hybrid cognitive system is significantly stronger than Clark and Chalmers’ notion of a coupled system, and also stronger than my notion of a strongly coupled system. In most strongly coupled systems, the cognitive artifacts that are used are strongly dependent on actions by the human user for their functioning, and are not capable of autonomous information processing. Computers are different in that they have the capability to autonomously perform cognitive tasks. Moreover, they are capable of forming coupled systems with other objects. These may be peripheral devices, like printers, other computers, or, in case robotic arms are added, ordinary objects in the environment. More often, computers couple with informational

objects: they perform epistemic actions on symbols and images stored in memory or input by a user, in order to perform particular cognitive tasks. Even the user becomes the object of epistemic actions by computer systems, when the system is programmed to ask input from the user in order to perform assigned information processing tasks. In conclusion, then, the computer is a special cognitive artifact that is different from others in that it is capable of autonomously performing cognitive tasks and is able to engage in symbiotic relationships with humans to create hybrid cognitive systems.

4. Computing and World-Simulation

In its early days, roughly from the late forties to the late seventies, the computer was exclusively a cognitive tool, since the only tasks that it was designed to do were cognitive tasks, like performing complex calculations and managing large amounts of information. This has changed with the advent of computers with good graphical and multimedia abilities in the late 1970s, 1980s and 1990s. These computers, most of them personal computers, acquired new functions that were not primarily cognitive. When a computer system is used to create an artistic drawing, to play an adventure game, or to listen to music, it is not used as a cognitive artifact, because the performed functions are not information functions: artistic drawings, adventure games and music are not meant to inform, but rather to please or entertain. These activities may involve cognitive activity (almost any activity does), but their principal goals are not cognitive. The computer systems and software that supports such activities therefore do not qualify as cognitive artifacts in this usage.

It may be objected that any kind of activity performed by a computer is, in essence, information processing, and that all functions of computers are therefore really cognitive. Computers, the argument goes, may be capable of performing tasks that when performed by humans would not be identified as information-processing tasks. Making drawings and producing music are examples. However, computers can only perform such tasks (or help perform them) by *reducing* such tasks to information-processing tasks. Newell and Simon (1976) have convincingly argued that computers are *physical symbol systems*: they are systems in which symbol structures that are capable of representing objects in the real world are manipulated in intelligent ways. The manipulation and interpretation of symbol structures is properly called information processing, even if the effect of these actions is not recognized by users as the execution of an information function. Therefore, the argument would go, even functions of computers that do not appear as information processing functions are really still that.

Although this conclusion may be true from the perspective of computer *programmers*, it surely is not true from the perspective of computer *users*, which is the perspective that is at issue when assessing the functional role of computers. It is certainly true that when a computer plays digital music the CPU takes in chunks of information and performs operations on them. But these chunks are only visible at the algorithmic level (Marr, 1982) to which ordinary users have no access. They are not visible at the functional level defined in relation to users. The functions of artifacts are not determined by the inner workings of artifacts, if any, but by the purpose assigned to them by designers and users. From the user's perspective, a computer that plays music performs a music-playing function, which is noncognitive in that the purpose of music is aesthetic rather than informational: it is to entertain and please the listener rather than to inform her of something. Of course, listening to music can also be analyzed as information processing in that it involves cognitive operations in the brain. But again, here, this is a process that occurs at a level of analysis not relevant to the user. At the level of analysis relevant to the user, she is merely enjoying music.

Most of the noncognitive functions of computer systems that have evolved since the

1970s critically depend on newly acquired abilities of such systems to graphically and sonically represent, simulate or model interactive objects, structures and environments. Following Floridi (1999), these abilities depend on the convergence of two trends. The first is what Floridi calls the *extension of digital encoding*. Computers originally only encoded formal and natural languages and numeric data. The next step was the encoding of three analog phenomena: sounds, images and motion, and a final step the encoding of three-dimensional, immersive virtual environment. The second trend is what Floridi calls *visualization and manipulation*: the development of visual and manipulable analog forms of access to digital information such as graphic user interfaces and WYSIWYG (“What You See Is What You Get”) software (Floridi, 1999: 14).

I will term the resulting representational and modeling abilities of computers *simulation abilities*. With the rise of high-quality graphical capabilities in computers, the computer is no longer just a *cognitive device*, it is now also a *simulation device* (cf. Turkle, 1995). To wit, the two functions of computer systems, cognition and simulation, are not mutually exclusive. In fact, many of the early efforts at graphical simulation were aimed at making the computer a better cognitive artifact. The Xerox Star was, in the late seventies, the first computer to make use of a graphical user interface, using a desktop metaphor that has since been copied by Apple (Macintosh) and by Microsoft (Windows). Desktop interfaces offer graphical user environment with documents, folders, trash cans, rulers, pencils and in and out boxes, that can be operated and manipulated in ways not unlike their physical counterparts. Their primary function, however, is to better support information-processing activities, particularly those performed in offices.

The advantage of graphical user interfaces over the older symbol-based interfaces (such as DOS and UNIX) is that they rely on our sensorimotor abilities to orient ourselves in space and to recognize and manipulate objects. Symbolical user interfaces make no good use of our sensorimotor abilities, and instead rely on our capacities for abstract thought. However, because people’s sensorimotor abilities are usually better developed than their capacity for abstract thought, it pays to treat data and programs as manipulable, visible objects, when possible. As a result, the tendency in software development has been to devise programs in which data strings, (sub)programs and procedures are translated into visual icons and actions like clicking, ‘dragging,’ and scrolling.

Around the same time that graphical user interfaces came into vogue, the first noncognitive graphical computer applications started to become popular: graphical computer games and creative software like paint and music programs. These applications are noncognitive because they do not have as their primary function to assist in the performance of information-processing tasks. Instead, they are intended to extend our means for entertainment and creative expression. They do so by simulating physical environments, objects and events. Tools are simulated with which we can interact with the world, like paint brushes, golf clubs, wrenches and guns, and the objects encountered in the graphical environment can be programmed to respond visually and aurally like their physical equivalents. Many environments can even be navigated, representing a position for us, and giving us the option to move to a different position. And in some environments, we can even interact verbally or nonverbally with computer-generated characters.

The computer in its role as (graphical) simulation device functions perhaps less as an extension of ourselves than as an extension of our world. The virtual interactive environments generated by computers offer us new structures to experience, navigate and interact with. They are hence an augmentation of the world as it existed before. Although these structures are not physically real, they are nevertheless meaningful or useful to us, sometimes as much as their physical equivalents. They can clearly be useful for performing cognitive tasks, as I argued in my discussion of graphical user interfaces. They are also useful for learning, particularly for learning sensorimotor skills, through their ability to faithfully simulate physical structures that we interact with. And they are useful for entertainment and creative activity. They hence serve a functional

role as broad and diverse as the functional roles of many of the structures encountered in the physical world.

The functional role of computer systems in their role of simulation devices may be termed *ontic*, because their role is to generate or represent objects and environments that form an addition to the physical world. Another development has taken place, however, that also adds to the ontic function of contemporary computers. This is a development different from the emergence of graphical simulation abilities in computers. Since the 1990s, many computers have become part of large-scale computer networks, specifically the Internet. When computers become networked, and they are used as means for social interaction by users, social structure emerges. Internet, specifically, has given rise to an online social reality that includes social roles and statuses, groups and organizations, institutions and social events, where most or all of this social structure is realized digitally.

As I have argued elsewhere (Brey, 2003), social reality, and in particular institutional reality, can be created in cyberspace with ease through the collective imposition of status functions by Internet users to appropriate digital structures (websites, buttons, windows, avatars, etc.). For example, Internet users can make it true that certain displacements of bits on the Internet count as buying, gambling, marrying, signing a contract, or winning a game by assigning the appropriate statuses to digital structures and events. Since the 1990s, this has occurred on a large scale, with the result that cyberspace now contains evolved social structures that collectively define new social realities for their users. These social structures need not be represented graphically, and are often represented by means of text and simple buttons or windows. But like graphical representations, they contribute to the ontic function of contemporary computers: they represent interactive objects with which users interact and that are part of their everyday ontology.

5. Conclusion

The functional analysis of computer systems presented here has identified computer systems as both cognitive devices and simulation devices. In its role of a cognitive device, the computer extends human cognitive faculties by both supplementing and enhancing them. In this latter role, collaboration between human minds and computer systems becomes so close that it results in hybrid cognitive systems that are part human, part artificial. In its role of a simulation device, the computer does not so much extend human faculties as extend the world. Computer-generated, virtual environments and Internet-generated social structures offer extensions of the “real” world that are useful and important for entertainment, creative activity, learning and social interaction.

This perspective has potentially important implications for human-computer interaction research. It entails that computer systems play substantially different functions for users, including various kinds of epistemic functions, ontic functions, and combined epistemic-ontic functions, with an important distinction within the class of epistemic functions between graphical simulation functions and social and institutional functions. An understanding of the differences between these functions is likely to be useful in user interface design. As for the wider implications of this perspective, perhaps the most important one is that an understanding of computers as (mere) information processing devices is increasingly outdated. Computers function more and more often as ontic devices that generate and sustain new virtual and social realities. Increasingly, they are to us not just information devices, but portals to worlds that we inhabit.

Note

This paper was originally presented as an invited plenary talk at E-CAP 2005 at Mälardalen University in Sweden.

References

- Brey, P. (2000). 'Theories of Technology as Extension of Human Faculties', in C. Mitcham, ed., *Metaphysics, Epistemology, and Technology. Research in Philosophy and Technology, vol 19*, London: Elsevier/JAI Press, pp. 59-78.
- Brey, P. (2003). 'The Social Ontology of Virtual Environments,' in D. Koepsell and L. Moss, eds., *John Searle's Ideas About Social Reality: Extensions, Criticisms and Reconstructions*, Oxford: Blackwell Publishing, 269-282.
- Clark, A. and Chalmers, D. (1998). 'The Extended Mind,' *Analysis* 58, pp. 7-19.
- Clark, A. (2001). 'Reasons, Robots and The Extended Mind', *Mind and Language* 16:2, pp. 121-145.
- Clark, A. (2003). *Natural-born cyborgs: minds, technologies, and the future of human intelligence*, New York: Oxford University Press.
- Clark, A. (forthcoming). 'Memento's Revenge: Objections and Replies to the Extended Mind', in R. Menary, ed., *The Extended Mind*, Ashgate.
- Donald, M. (1991). *Origins of the Modern Mind*, Cambridge and London: Harvard University Press.
- Floridi, L. (1999). *Philosophy and Computing. An Introduction*, London and New York: Routledge.
- Hutchins, E. (1995). *Cognition in the Wild*, Cambridge, MA: MIT Press.
- Kirsh, D. and Maglio, P. (1994). 'On distinguishing epistemic from pragmatic action', *Cognitive Science* 18, pp. 513-49.
- McLuhan, Marshall (1966/1964). *Understanding Media: The Extensions of Man*, New York: McGraw-Hill. Paperback edition, 1966.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*, W.H. Freeman Publ.
- Newell, A. and Simon, H. (1976). 'Computer Science as Empirical Inquiry: Symbols and Search', *Communications of the ACM*, 19, pp. 113—126.
- Norman, D. (1993). *Things that Make us Smart: Defending Human Attributes in the Age of the Machine*, Reading, MA: Addison-Wesley.
- Perry, M. (2003). 'Distributed Cognition,' in J. Carroll, ed., *HCI Models, Theories, and Frameworks: Toward a Multidisciplinary Science*. San Francisco: Morgan Kaufmann Publishers, pp. 193-224.
- Peschl, M. and Stary, C. (1998). 'The Role of Cognitive Modeling for User Interface Design Representations: An Epistemological Analysis of Knowledge Engineering in the Context of Human-Computer Interaction', *Minds and Machines* 8, pp. 203-236.
- Salomon, G., Ed. (1993). *Distributed Cognitions: Psychological and Educational Considerations*, New York: Cambridge University Press.
- Sterelny, K. (forthcoming). 'Externalism, Epistemic Artefacts and the Extended Mind', in R. Schantz, ed., *The Externalist Challenge: New Studies on Cognition and Intentionality*, Berlin and New York: De Gruyter.
- Turkle, S. (1995). *Life on the Screen: Identity in the Age of the Internet*, New York: Simon & Schuster.