Łukasz Afeltowicz

Nicolaus Copernicus University

Witold Wachowski

University of Warsaw

HOW FAR WE CAN GO WITHOUT LOOKING UNDER THE SKIN: THE BOUNDS OF COGNITIVE SCIENCE

Abstract. The aim of this paper is to discuss the concept of distributed cognition (DCog) in the context of classic questions posed by mainstream cognitive science. We support our remarks by appealing to empirical evidence from the fields of cognitive science and ethnography. Particular attention is paid to the structure and functioning of a cognitive system, as well as its external representations. We analyze the problem of how far we can push the study of human cognition without taking into account what is underneath an individual's skin. In light of our discussion, a distinction between DCog and the extended mind becomes important.

Keywords: cognitive science, distributed cognition, externalism, internalism, representations.

Distributed Cognition: General Remarks

In cognitive science, the term “distributed cognition” (henceforth, DCog) has at least three basic meanings or uses:

- (1) as a general, blanket category covering all “new”, non-traditional theories and conceptions of cognition;
- (2) as a loose “sensitizing” metaphor used primarily to highlight the fact that tools, the environment, and material and social factors influence the cognizing agent in some (not necessarily specified) way, and
- (3) as a name of the theory of cognition formulated by Edwin Hutchins in *Cognition in the Wild* (Hutchins, 1995a; see also Hutchins, 2001).

In this paper, we generally use DCog to refer to the specific approach introduced by Hutchins. However, the history of this approach and its current status are influenced also by other researchers and authors.

DCog focuses on an analysis of cognitive systems which are wider than single agents (e.g., humans or animals). In the case of the DCog theoretical

framework, the main units of analysis are not individuals but assemblages of human, material and cultural factors. Those assemblages are called distributed cognitive systems (DCSs). Here are a few examples of DCSs: a team of air traffic controllers and their equipment (including the bay station), financial brokers and their market devices for trading and data visualization, a team of surgeons working in an operating theater. One should always remember that DCSs may consist of humans or other cognitive agents, but the properties of a wider cognitive system differ from the properties of its parts and the former cannot be reduced to the latter. Furthermore, a precise description of a particular DCS requires observation or other research activities: one cannot specify the parts of a DCS in advance (in other words, one cannot know *a priori* what plays an important cognitive part and what does not). This is only possible in hindsight.

We usually conceptualize thinking as a process happening inside someone's head: as something that an individual does and that takes place beneath her cranium. As DCog suggests, thinking processes are not necessarily abstract and totally disembodied. While remaining agnostic about what is beneath the skull, DCog shows that thinking happens also outside, in the world (or at least DCog states that it is good to conceptualize some processes taking place in the environment as thinking). To be precise, DCog defines cognition as computation realized through the creation, transformation, and propagation of representational states (Hutchins, 1995, p. 49). This is quite a classical definition of cognition, but, in DCog, the notion of computation is applied to wider cognitive systems, consisting of whole groups of people and artifacts that function in the framework of culture. The whole process of computation is coordinated by human agents. However, the bulk of cognitive processing is not performed by human agents, who are only "lightly equipped" with cognitive skills and internalize a small part of the information and processes. They should be regarded as yet another (though special) medium of the process of computation: "The thinker in this world is a very special medium that can provide coordination among many structured media." (Hutchins, 1995a, p. 316)

In standard approaches, the "stuff" that thoughts were made of was simply mental (Thagard, 2014; Lau & Deutsch, 2014). In the case of DCog, however, representational states turn out to be mainly social and material: in DCSs, the media of information include gestures, speech, sounds, sheets of paper, instrument readings, lines drawn on charts and the charts themselves. According to DCog, the functions of these things should not be reduced to delivering input for our "internal computational engine": in many situations things function as computational devices themselves.

The approach of DCOg differs from that of traditional cognitive science in a significant way. The traditional approach assumed that it is much easier to study individual cognition than collective cognition, so we should start with a single mind and, subsequently, gradually move to more complex matters, such as the influence of culture, social factors, or type of material media on individual cognition. As Hutchins states, “The early researcher in cognitive science placed a bet that the modularity of human cognition would be such that culture, context, and history could be safely ignored at the onset, and then integrated in later.” (Hutchins, 1995a, p. 354)

Integration of cultural and material factors at the outset of a study of human cognition, as Hutchins suggests, creates a methodologically favorable situation for the cognitive scientist. By switching the unit of analysis from an individual, disembodied mind to a socio-technical collective, she may now – literally – find herself at the heart of a cognitive system and study the interactions of its components by means of standard, direct observation. For example, she can follow external representations, observe how they are transformed, look at what kind of skills and structures those transformations enable, how coordination between them and different media is achieved, etc. Eventually, this chain of representations will “hide” itself inside of agents’ heads or in some other “black boxes”, ceasing to be directly observable. But, in many cases, we are able to trace representation transformations for a surprisingly long time. Hutchins illustrates this with the example of modern navigation on board of a US Navy carrier (1995a). Let us discuss it, in order to understand how a DCS functions and what it means for a wider system to have properties that are not exhibited by the agents that act as its parts.

Case Study of DCS: Maritime Navigation

The basic problem that a navigation team on board of a surface vessel at sea must solve is “Where are we now and where are we headed?” In general, all navigational systems answer this by combining one-dimensional constraints on position. Current western navigation practices rely heavily on digital technologies. Hutchins’ study, however, took place before the advent of GPS. The author of *Cognition in the Wild* observed navigators performing such tasks as taking visual bearings of landmarks with a pelorus, measuring the vessel’s actual velocity with a dummy log and a pit sword, estimating the depth of water under the ship with a fathometer, plotting bearings on a navigation chart with a pencil, a hoey (a specialized protractor with

a movable straight-edged arm that pivots in the center of the protractor's scale), and other plotting devices.

Precise position fixing with such instruments requires combining at least three one-dimensional constraints on position. Usually, the navigation team observed by Hutchins used lines of position (LOP). A LOP is a line in a two-dimensional space (the surface of the sea or the chart representing it) crossing the actual position of the vessel. In order to draw a LOP on a chart, navigators take a visual bearing: using a pelorus, they measure the direction of the line of sight connecting the ship and some visible landmark on the shore. Knowing the angle of the line of sight one can plot it on the chart in such a way that it will cross the point representing the landmark, which was taken as a reference point. In order to do so, navigators use a hoe: they set a protractor perpendicular to the meridian in such a way that the center of its scale is at the landmark, then they set the arm at the angle which was reported and draw a line with a pencil. The position of the ship is somewhere along the drawn line. In order to fix a position, one needs another LOP: the ship's location should be the point where both lines cross on the chart. But a third line is also required as a control constraint. In practice, due to lack of precision, three visual bearings plotted on a chart always form a triangle. The ship's location is somewhere in this triangle and the size of the triangle's area represents the precision of the bearings (the smaller it is, the more precise all three bearings are). What is important, is that the information is easy to extract from the drawing even for inexperienced members of a navigation team.

A chart is probably the most important cognitive device at the navigation team's disposal. All the team's activities focus around a table with a chart. A chart is a specialized, two-dimensional model representing the geographical space surrounding the ship. The most obvious property of a chart is that points on its surface correspond with the environment, especially with landmarks. Additionally, a chart contains information on the shape of the coastline, and the depth and shape of the seabed. It gives a unique perspective: something more than a bird's-eye view. Some elements of the environment are highlighted or revealed, while others are deliberately simplified or omitted. Furthermore, a chart translates a fragment of a sphere's surface into a "flat" representation: this translation is distorting, but, simultaneously, it simplifies the task performed by the navigation team. It is one of many examples how DCSs or their elements reduce cognitive complexity.

According to Hutchins, navigation charts should be considered as something more than mere representations: they are also carefully crafted computation devices (Hutchins, 1995a, p. 61). In algebra and analytic geometry,

many calculations can be performed on graphs. For example, we can determine all points between points A and B, by simply drawing a straight line connecting A and B. A navigation chart is closer to a graphical representation of mathematical abstractions than to an ordinary map. A chart represents the geographic features in such a way so as to facilitate or even enable analog computations associated with maritime navigation. Of course, all problems that navigators solve through graphing practices could be represented in the form of equations and solved by abstract symbol processing. Chart use and bearing plotting, however, simplify the whole process.

A chart is an external representation of a problem that could be represented internally. Instead of solving navigation problems in the head, in an abstract manner, members of a navigation team use charts in order to answer the question: “Where am I?” by performing such activities as line drawing, coordinated visual bearings taking, or writing and rewriting digits, with the use of a protractor, ruler, and calipers. In a way, then, members of a navigation team think with their hands and eyes (Hutchins, 1995a, pp. 142–143).

Position fixing requires that inputs from different instruments and other sources of data be skillfully integrated. A chart integrates those inputs, but it is only one of many elements of representation processing in a DCS (see Figure 1).

Let us look more closely at the sequence of actions.

- 1 Navigator (N) working with a chart informs a quartermaster of the watch (Q) that the moment for taking bearing is close and suggests appropriate landmarks.
- 2a-2c Q assigns one landmark to each of three pelorus operators (PO) and – if necessary – instructs them on how to find their landmarks on the horizon; until (3) all POs constantly track assigned landmarks with their visual devices.
- 3 Q gives order “Mark!”: all POs simultaneously take bearings of all assigned landmarks and memorize their values.
- 4a-4c Q collects bearings as memorized by each PO.
- 5 Q records on a sheet of paper the name of each landmark and a three-digit number associated with it, as delivered by each PO.
- 6 Q reads the bearings from the sheet.
- 7 In the most suitable moment, Q delivers to N the landmarks’ names and associated number values.
- 8 N plots the bearing on a chart with a pencil and hoe (draws on it three LOP which form a triangle).

Figure 1. Position fixing based on visual bearing performed by a navigation team on board of a sea vessel as an example of distributed cognition. PO – team member operating one of three peloruses; N – navigator, who selects landmarks, plots bearing on a chart, and reads the solution to the problem from a chart; Q – quartermaster of the watch who passes the names of landmarks (to each PO) and bearings (to N); arrows symbolize the directions in which representational states and/or orders flow; circled Ms symbolize where short storage of data takes place.

9 N (or any other navigation team member) reads the position fix off the chart and alongside with it the feedback information on bearings' precision (he estimates the area of the triangle formed by LOPs).

10 N periodically delivers navigation reports to the Captain.

This is how a vessel's position is fixed. This, however, is not achieved by a single member of the navigation team. Not even a navigator working with a chart achieves that: he is merely reading a solution to the navigation problem from the chart. A solution is not appearing in anyone's head, but it emerges on a chart as a result of various human and non-human elements' coordination. The chart functions as a data integrator, but the quartermaster is crucial for effective work cooperation, representational states proliferation, timing, and buffering.

As can be seen, charts play multiple functions in navigation: (1) they are a crucial external representation enabling analog computations; (2) they

are a tool for integrating and comparing diverse data inputs; (3) they help to coordinate actions of other DCS parts; (4) they deliver feedback information concerning the quality of work performed. Additionally (5), the chart “embodies” the knowledge and experiences of past generations of geographers, cartographers, navigators, sailors, and discoverers. In other words, the symbolic and material structure of the chart reifies the knowledge of its users: past and present (Hutchins, 1995a, p. 96). Fortunately, a user of this “splice” of knowledge does not have to understand past innovations or the history of navigation charts: he may treat a chart as a firmly closed, easy to use, off-the-shelf tool: a black box (Latour, 1987). In fact, the western navigation tradition abounds with similar black-boxed analog computation devices; other examples include a dummy log, a wind rose, and an astrolabe (Hutchins, 1995a, pp. 99–104).

When analyzing western navigation practices, one should not overestimate the role of material artifacts, or the social or cultural factors that come into play. What is important is not the quartermaster, the chart, a pelorus, or the navigator alone; what counts is what happens between all those elements. DCog is all about the process of distribution, propagation, and transformation occurring within a DCS. The distributed nature of this system means that one is unable to identify a human mastermind who solves the problem in his head alone or something like a non-human processing unit which performs the bulk of the work. If one of these is the case (if you can reduce the whole system to one of its elements), it is not a DCS.

We should now briefly refer to other features of DCS, such as redundancy, buffers, flexibility, robustness, and learning processes. Let us discuss these issues in turn.

In every DCS, there is a trade-off between redundancy and complexity. The redundant parts, parallel processes, and inter-locking functions of human or non-human elements render a system more resilient and accident-resistant (Hutchins, 1995a, pp. 189–190). The redundant elements, such as the required third LOP, facilitate error detection. But cognitive redundancy is costly and time-consuming (see Hutchins, 1995a, pp. 180–181, 220–223, 264–266). And the more complex a DCS, the harder it is to develop and manage it.

Buffers in DCS (see Hutchins, 1995a, pp. 194–196) are decisive for distribution of representational states in a timely manner (e.g. at a rate matching the cognitive capacity of a particular part of the system) and in proper sequence (which may reduce a particular agent’s memory load). As we have seen, without the quartermaster, his stopwatch and his notes, the represen-

tational states would not circulate efficiently among the team members, and the probability of error would be higher.

Each time communication between the elements of a DCS fails and errors occur, the navigation team does not stop working in order to evaluate their work. In the case of a crisis, many DCSs are reorganized without interrupting their work. They react in a flexible and dynamic manner to changes in their environment and their own functioning, partly due to the previously mentioned buffering and redundancy.

Collective cognitive efforts create favorable conditions for organized learning processes. The tacit knowledge required for mastering a practice is transmitted mainly through the master-disciple relationship and observation of experts at work. This is precisely the way in which new members of a navigation team are trained. A novice gradually learns the capabilities and limitations of particular navigation procedures and technologies; he is familiarized with typical problems that may occur and standard ways of dealing with them. By solving problems, navigators not only learn how to handle the elements of the system, but at the same time they also take part in developing the system. Many generations of sailors have been involved in the process of historical accumulation of navigation experiences and inventions. This evolution resulted in the creation of different types of black boxes, based on which subsequent generations could develop their own innovations.

One of the most important categories of DCog are representational states. In fact, DCog focuses mainly on external representations. This category will be crucial in the context of a comparison with other non-traditional DCog approaches. Let us elaborate on this category.

External Representations

External representations (xRs) are everywhere around us. Consider a calendar, a grid drawn on a piece of paper that can also function as a board for tic-tac-toe or a visualization of an engineering problem. As cognizing agents, we are not only surrounded by xRs. They enable or support many of our cognitive activities, and also our very cognition is shaped and constituted by xRs. On the one hand, xRs play an important role in the ontogenetic and phylogenetic development of humans; on the other hand, they are constantly involved in our designing artifacts – from simple everyday things to various (more or less complicated) technical devices – as well as our interactions with them, including human-computer interactions (Malafouris, 2013; Norman, 1993, 2002; Hutchins, 1995; Hollan et al., 2000).

According to the classical approach, xRs as components of our external environment can be reduced to sources or carriers of input for our internal cognitive processes (see: Horst, 2009; Lau & Deutsch, 2014). Furthermore, xR functions can be interpreted as mere cognitive scaffolding supporting or enhancing internal processing: something that makes cognitive processing easier, faster, more reliable, etc., but does not add any new quality. As we will show, xRs should be considered as something more than “data-storage” for internal computations or cognitive “busters”.

Jaijie Zhang and colleagues conducted important research on xRs in the context of DCS. Zhang offered a description of xR and iR interactions in the domain of problem solving. According to Zhang, the interactions structure the problem solving process in such a way that a part of the agent’s activity, previously devoted to internal computations, is transferred to management of external elements which remain in the agent’s proximity. One could say that the introduction of xRs changes the nature of the work performed. Instead of manipulating internal images, an agent performs actions on external objects (e.g. manual manipulations on a drawing). This transition from internal to external manipulation may lower the agent’s general cognitive workload: for example, in the mentioned Zhang’s study, game participants didn’t need to remember the goal problem states, due to the presence of diagrams in front of the participants (Zhang, 1991, pp. 954–958).

In order to better understand the meaning of xRs, their roles and the relation between xRs and iRs, let us discuss a classic example of external visual representation used in one of the oldest and most popular games: chess (Zhang, 1997, p. 179). Chess experts are capable of memorizing the configuration of pieces on the chessboard, the sequence of movement in a particular game and even play without the use of pieces and board, only using their mental models of the game (so-called blindfold chess). We may assume that chess masters memorize not the position of particular pieces on board, but rather a whole “tactical landscape”, including spatial relations between pieces (this claim is supported by eye tracking studies), moves which are afforded by particular relations or the more or less standard, overlearned tactics which led to this particular configuration of pieces. Less experienced players must resort to the chessboard and pieces as an xR of the gameplay. By using it, they relieve themselves from activities performed by the cognitive systems of expert players. First of all, the board and set of pieces are the equivalent of a short-term memory system for novices and advanced players.

As Zhang states, xRs may be transformed into iRs through the process of memorization. Frequently, this internalization process is not necessary, for example when a particular xR is easy to access and constantly present

in our surroundings. In some cases, internalization is impossible due to the high complexity of the phenomenon it represents (of course, the degree of complexity should always be judged in relative terms, taking into account the capabilities of a particular cognitive agent). Externalization (transformation of iRs into xRs) is recommended in situations when the benefits of using this particular type of xR outweigh the costs associated with the very externalization processes, the use, and maintenance of xRs (see Zhang, 1997, p. 180). Once again this cost and benefit ratio depends on the capabilities of an agent. Let us refer to chess once again. An expert does not need a chessboard in order to play. In her case, use and maintenance of xRs (preparation of starting set, moving the pieces, visual analysis of current situation) is time consuming and requires additional cognitive activities which – at least from her perspective – do not seem to be essential for playing the game. For an expert those activities are an unnecessary cognitive workload. But for novices and advanced players, all those activities are the essential elements of the game. Without them, they probably would never be able to achieve expertise in chess. But the fact remains that novices and advanced players think of chess and even look at a chessboard and the configuration of the pieces in a different way than expert chess players.

One could say that xRs reduce cognitive workload by transforming a problem so as to make it easier to solve (the general costs of externalization and use of xRs is outweighed by cognitive benefits). But we should be aware that externalization frequently changes the nature of the task at hand. A good example is the multiplication, division, subtraction, or addition of multi-digit numbers. Just try to multiply two random two-digit numbers higher than 20 “in your head”: it is quite a good test of how good your working memory is. Of course, there are people (e.g. savants) who are capable of mentally performing even more complex arithmetic operations. But most of us (including the authors) must devote a few minutes to the task: the time required to solve the problem depends on i.e. our acquaintance with mnemonic techniques and specialized problem solving heuristics. Probably, most of us when solving such a task would prefer to do it the way we were taught in elementary school: by doing long multiplication on paper. A sheet of paper “remembers” for us the partial results. It also allows us to easily monitor on which step of the problem solving procedure we currently are. It is interesting that when we have to solve the problem without paper and a pencil, using only our working memory, we internally emulate this grade-school technique of multiplication. The result does not suddenly occur to us – we use the method of long multiplication, but without paper and a pencil.

The examples discussed above could be interpreted as demonstrating the role of xRs in memorizing. But xRs' role far exceeds the function of external memory. As we mentioned earlier, some theorists try to reduce components of our external environment to mere carriers of information for our internal processes which consist of iR manipulations. According to this view, even if cognition involves a lot of social interactions and interactions with material things, the whole cognitive processing takes place within individual minds. Mental manipulations of iRs are considered to be the only genuine cognitive processes there are. Cognition requires an agent to internalize elements of its environment, create a mental model based on input data, perform mental computations, receiving at the end a product which can be externalized thanks to some kind of decoding and – if suitable – preserved in the form of some kind of xRs. An alternative view is that our environment is something more than a source of data inputs, a storage facility for products of mental calculations performed by various minds and a space through which we are trying to send data to other individual minds using different media. This framework assumes that selected elements of the environment play an active role in problem solving, that mental models are not always necessary. Using the computational metaphor, according to this alternative framework, we should consider external things as parts of a wider system performing calculations. This view also suggests that we should always analyze particular xRs in the context of whole DCSs: the way xRs function will depend on the skills of humans and the presence of other non-human elements. Returning to the navigation example, a chart is a powerful xR, but it would be useless without the skillful hands and eyes of navigation team members, basic knowledge concerning what LOP means, or the data that the chart is fed.

Note that the adjective “external” is sometimes confusing, for xRs are not external to the cognitive system. xRs are data themselves and function inside the system. On other hand, note that components of the environment are xRs temporarily, depending on the context (DCog system).

Hutchins stresses that we should be careful with the notion of externalization and internalization of representations:

Internalization has long connoted some thing moving across some boundary. Both elements of this definition are misleading. What moves is not a thing, and the boundary across which movement takes place is a line that, if drawn too firmly, obscures our understanding of the nature of human cognition. Within this larger unit of analysis, what used to look like internalization now appears as a gradual propagation of organized functional properties across a set of malleable media. (Hutchins, 1995a, p. 312)

We have yet another problem with the metaphor of cognitive enhancement. Let us look at long multiplication to see how exactly xRs contribute to cognition. One could say that paper notes (columns of digits) are a cognitive scaffolding (Norman, 1993; Clark, 1997, p. 61): not only do they simplify cognitive problems, but they also enable the performing of more complex tasks. They enhance and support our mental structures and abstract reasoning. But we should use the metaphor of scaffolding, cognitive enhancement and even mind extension carefully. Use of xRs does not always lead to an “upgrade” of our intelligence or cognitive capacity: in many cases, xRs should not be considered in terms of prosthetics or enhancers, which allow us to see further, compute faster in our heads, memorize better, or analyze more reliably. The main advantage of the presented artifacts, visualizations, and problem solving techniques lies in the fact that they help to transform complex problems into problems that can be tackled by humans lacking in extraordinary cognitive capacities. DCog’s notion of mediation does not suggest that xRs or any other element of a DCS stand “in between” mental events or between the user of a cognitive tool and the task. XRs, cognitive tools, media etc. “stand with the user as resources used in the regulation of behavior in such a way that the propagation of representational state that implements the computation can take place” (Hutchins, 1995a, p. 154). DCog suggests that one should view xRs as one of many elements that are brought into coordination in a problem solving practice (Hutchins, 1995a, p. 290).

In order to see how tools and xRs participate in reducing the complexity of a cognitive task rather than boosting the cognitive capabilities of a single mind, let us use the example of maritime navigation once more. A plotter, while coordinating a hoey, a chart, numerical values (the bearing), and names (landmarks) delivered by the quartermaster, in fact coordinates various media. While he produces the coordination of the media, there is no need to concentrate on media as carriers of data; only once the coordination has been done may the plotter extract information from a drawing. But this last act is somewhat trivial, while the whole process of answering the navigation question “Where are we?” is not. Hutchins concludes that tools (including various xRs mentioned here) “permit the people using them to do the tasks that need to be done while doing the kinds of things the people are good at: recognizing patterns, modeling simple dynamics of the world, and manipulating objects in the environment” (Hutchins, 1995a, p. 155).

Hutchins’s study of pilots’ interaction with their devices within an airplane cockpit system (Hutchins, 1995b) seems to provide a model example

of how an xR is generated, transformed, and propagated. The memory tasks in the cockpit are realized by processes distributed between human agents and external representational devices.

These properties seem to determine both human-device interactions and the internal cognitive activity of pilots, related to, among others, decision-making processes.

The studies presented by us above focus on visual xRs. Their contribution to cognition can seem trivial to some extent and thereby transparent for researchers and philosophers. Therefore, in our opinion, research in another area of representation is worth looking at. This will allow us to refer to a less exploited modality and draw partially counterintuitive conclusions.

Case Study: Thinking with the Body and Kinesthetic Representation

When we are talking about human cognitive activity – or more precisely, about human perception – we usually focus on visual perception. Visual representations have been raised to such a level that they are usually considered to be the most reliable witness, transfer, or illustration. The dominant role of visualization in the work of a researcher went through in-depth analysis on the part of ethnographers of science. But we do not identify visual representations in the modal external insulation. As mentioned earlier, we do not identify them passively. Our experience and knowledge – acquired in the course of an embodied interaction with the environment – has a history associated with the context of the acquisition process. Its effect – related to the interactive architecture of the brain itself – is to recall the information right away with the whole structure of perceptual and motor connections. It is not in acquiring knowledge about the characteristics of the bottle neutral, but learning to go after it and keeping it in a special, secure way to know its surface and the sounds made when in suitably dynamic contact with certain surfaces.

The body, which is actively involved in the cognitive system, is not only a medium in contact with xRs. It can also act as a generator or a “screen”. Moreover, it may serve as a medium of representation in the sphere of proprioception. To see this in more detail, let us look at Kirsh’s research, which shows that people can use their bodies as a device for simulating and modeling. In his study of the behavior of dancers, Kirsh focuses on the activity called marking. Marking is the execution of a simplified, schematic, sketchy

version of a particular movement or behavior. It is an outline of a dance phrase, a kind of sketch or model representing particular aspects of a ‘full-out’ dance phrase (Kirsh, 2010b).

Marking is something between performing a complete dance phrase and the purely mental act of visualizing it. Kirsh compares this intermediate, sketchy, “crippled” character of marking with the act of projection localized between perception and imagination. Marking can take various forms: from perfunctory gestures to relatively complex movements, from representing a phrase with the whole body to marking it with just the hands, in which case it is often highly conventionalized. Kirsh distinguishes three types of marking: (1) marking-for-self, which is when the dancers mark a dance phrase for themselves, (2) marking-for-others, and (3) joint-marking, when a team of dancers mark a phrase to explore or verify timing and coordination. Kirsh proposes to see this as the external part of a distributed vehicle of thought, with some benefits: marking can be treated as a method for anchoring projection and a method for priming as well. “A dancer’s thought of his or her phrase is partly shaped by what is marked. Dancers do think about their phrases without dancing them or marking them. But, by marking-for-self dancers think better about their full-out phrase. Physical movement replaces mental computation” (Kirsh, 2010b, p. 2869). Therefore, marking can be treated as a part of an internal-external structure.

Let us focus for a moment on marking-for-self. When we mark for ourselves – without using a mirror – we perceive a phrase only in a kinesthetic manner. The conclusions of Kirsh’s study are interesting and somewhat counterintuitive. It turns out that the incomplete, biased, distorted simulation of movements and gestures performed while marking improves learning more than a full mental simulation or even a complete, real-time performance of target dance phrases. Such a sketchy, uninvolved representation of a dance phrase seems to be the best medium for the content, the optimal structure of the intermediate, which does not excessively burden a dancer. Such imperfect kinesthetic models also allow dancers to exceed their own limits. As in the case of manipulating a real object or idea sketching tools, marking can be regularly successful in moving our cognitive activity in a new and unexpected direction and in learning skills. This kinesthetic-visual or only kinesthetic form of representation confirms that the externalization of representation perfectly serves our interaction with the world, both material and social – as well as with themselves, which seems to be good evidence for kinesthetic marking for itself.

Discussion

The DCog framework forces a redefinition of some traditional categories (such as that of an agent and cognitive process) and it occupies a special position among the “new” theories of cognition. In some respects it is similar to such notions as the extended mind, enaction, embodied mind, embedded cognition etc. But it is not another version or incarnation of embodied cognition. The reason seems to be obvious: many advocates of embodied cognition – starting from the authors of the classic book *The Embodied Mind* (Varela et al., 1991) – criticize the computational model of cognition, treat it as reductionist and insufficient, while DCog makes extensive use of this concept. This requires a clarification: DCog is a version of computationalism (it explicitly refers to cognition as a production, transformation, and propagation of representational states).

DCog, however, distinguishes itself from traditional computationalism in at least two ways. Firstly, the unit of analysis is different: DCog applies the old metaphor of cognition as computation to a different entity. Secondly, DCog describes the production, propagation, and transformation of external representational states, not internal (mental) ones. DCog remains agnostic about the processes that take place in the mind of an individual and it is not interested in the mental content of thoughts, although precisely those categories were at the core of traditional computationalism.

Of course, we have to be aware of how the approaches to DCS are different, similarly as with the approaches to agency in DCS. We do not advocate any approach to mental processes here. However, in the light of these studies, we believe that there is no need to connect the theory of DCog with the theory of the extended mind. A connection like this can be found in some critical views of DCog (see Rupert, 2013; Button, 2008). It seems to be a cause of confusion. We should also remember that DCog is not the first conception which approaches cognition and mind in such a way – cf. John Dewey’s epistemology and his theory of knowledge and knowing (Dewey & Bentley, 1949).

It should be stressed that the development of research on DCog is related to the substantial contribution of not only psychologists and cognitive scientists, but also social scientists. Special attention should be paid to the contribution of the anthropology of science and technology to DCog. Anthropological studies of scientific practices have facilitated identification and description of new structures and interactions within the system, which lead to consolidation and preservation of some interdependence between the

components, and to the disappearance of others. This is reminiscent of the emergence and activity of the actors in the actor-network theory (ANT) (see Latour, 1987), which is a reference not denied by the advocates of DCog, and shares a careful and unprejudiced interest in involving inanimate elements of the environment in the cognitive system.

Here we need to make some limitations. For Hutchins, as Paul Cobb shows, we have direct access to observable xRs as elements of a sociocultural system, but sometimes the observable stream of representational states hides inside the individuals, as internal symbol processing. It is in these situations that Hutchins admits analyses of internal symbol processing. It seems that, to some extent, Hutchins' theoretical model is compatible with classical cognitive science. In this connection, Cobb draws attention to an interesting exchange between Latour reviewing *Cognition in the Wild* and Hutchins responding to reviewers (Cobb, 2011, pp. 95–96; Latour, 1996; Hutchins, 1996).

Latour criticizes Hutchins' alleged vision of a division between the world and cognitive skills. As Latour writes, “[d]istribution, in my view, does not go all the way. This leads Hutchins to make mistakes even in navigational matters” (Latour, 1996, p. 60). While Hutchins responds: “The work must be done somewhere, and some of the work will be done in regions that lie inside the bounds of persons.” (Hutchins, 1996, pp. 64–65).

In Cobb's opinion, the dispute between Hutchins and Latour is about how much of the cognitive activity of DCS is internal and how much is external. In the context of their arguments, and in the context of mainstream psychology, the boundary of the skin seems to be decisive. Acceptance of Hutchins's arguments seems to entail an approach related to intelligence partitioning rather than distributing. Acceptance of Latour's arguments seems to be “pushing cognition out beyond the skin”, and thereby “emptying the person” (Cobb, 2011, p. 96). We could say that Hutchins is methodologically agnostic about the cognitive content of individual minds: he does not suggest that “everything is outside”, only that many things that we were searching for inside individuals can actually be found in the environment. Hutchins openly admits that his approach does not allow us to follow individual cognitive processes and, clearly, there is a place and need for approaches that are complementary to DCog. On the opposite side is Latour, who spreads individuals' minds and cognition onto so many things that he eventually dissolves both of the categories.

Summary

Computers were the models of the mind for traditional cognitive science. Current cognitive scientists consistently abandon the literal analogy between mind and computer; however, computationalism as such is doing well: the majority of research in cognitive science is computational. The DCog framework shows that the notion of computation might actually offer an adequate description of cognitive processes on the level of larger sociocultural systems. DCog, with its definition of cognition as the production, propagation, and transformation of representational states, does not easily fit into standard distinctions present in cognitive science. It is not a traditional computationalism: although it makes extensive use of the notion of computation, it applies it to a new set of phenomena. DCog is often classified as a non-traditional approach similar to embodied cognition or the extended mind; however, it stands out from the various approaches subsumed under the label “situated cognition” in that it posits a specific unit of analysis (which is DCS, not an individual) and that it does not reject the notions of computation and representation. Actually, there is no need to connect the DCog framework with the extended mind theory or to attribute to a DCS as a whole any form of cognitive agency. Moreover, DCog’s notion of mediation seems to be more theoretically sophisticated than notions of enhancement and scaffolding.

Although DCog is convergent with the cognitive anthropology of science, one can point to some differences between these two approaches. We should also remember that DCog avoids discussion of the substance and boundaries of the mind, or the nature of internal representation. DCog should not be considered as an attempt to “empty” the individual mind.

DCog does not dispute the existence of the iRs, it only valorizes xRs (or even establishes them as a full-fledged analytical category), indicating that they often fulfill functions usually attributed to internal components. DCog may be regarded as an invitation to follow chains of xRs and their transformations as far as possible. It is an attempt to explore how far we can push the study of human cognition without taking into account what is underneath an individual’s skin. But we should be aware that the stream of representational states eventually disappears inside the individual actors. There is a need for application of complementary approaches capable of answering the questions concerning the content of individual cognitive processes, which remain black boxes for the DCog framework. In other words, DCog should not be considered as a full-fledged theory of cognition at all. It delivers concepts useful when explaining cognition happening in the world, but

requires help from other cognitive theories when explaining what happens inside an individual's head. On the other hand, DCog is a desirable counterweight to the cognitive sciences focusing solely on internal representations, boiling down a cognitive system to the brain and based almost exclusively on neurocognitive methods. This is a good moment for us to ask ourselves the following question: How far can we push the study of human cognition without taking into account what is inside the human cranium and skin, and what happens when people interact and coordinate actions of various media? An important methodological consequence of the DCog framework is that one cannot understand cognition without taking into account how social structures and cultural factors (including material culture) work.

REFERENCES

- Button, G. (2008). Against "distributed cognition". *Theory, Culture & Society*, 25(2), 87–104.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Cobb, P. (2011). Learning from distributed theories of intelligence. In E. Yackel, K. Gravemeijer, A. Sfard, & P. Cobb, *A journey in mathematics education research: Insights from the work of Paul Cobb* (pp. 85–105). Dordrecht: Springer.
- Dewey, J. & Bentley, A. F. (1949/2014). *Knowing and the known*. Boston: Beacon Press / American Institute for Economic Research. Retrieved February 16, 2014, from <https://www.aier.org>.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2), 174–196.
- Horst, S. (2009). The computational theory of mind. In E. N. Zalta et al. (Eds.), *Stanford encyclopedia of philosophy*. Retrieved August 20, 2014, from <http://plato.stanford.edu/entries/cognitive-science>.
- Hutchins, E. (1995a). *Cognition in the wild*, Cambridge, MA: MIT Press.
- Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265–288.
- Hutchins, E. (1996). Response to Reviewers. *Mind, Culture, and Activity: An International Journal*, 3(1), 64–68.
- Hutchins, E. (2001). Distributed cognition. In N. J. Smelser & P. B. Baltes (Eds.), *The international encyclopedia of the social and behavioral sciences* (pp. 2068–2072). Cambridge: Elsevier Science.
- Kirsh, D. (2010a). Thinking with external representations. *AI and Society*, 25, 441–48.

- Kirsh, D. (2010b). Thinking with the body. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2864–2869). Austin, TX: Cognitive Science Society. Retrieved January 17, 2014, from <http://adrenaline.ucsd.edu>.
- Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Milton Keynes: Open University Press.
- Latour, B. (1996). Cogito ergo sumus! or psychology swept inside out by the fresh air of the upper deck... *Mind, Culture, and Activity: An International Journal*, 3(1), 54–63.
- Lau, J., & Deutsch, M. (2014). Externalism about mental content. In E. N. Zalta et al. (Eds.), *Stanford encyclopedia of philosophy*. Retrieved August 20, 2014, from <http://plato.stanford.edu/entries/content-externalism>.
- Norman, D. A. (1993). *Things that make us smart: Defending human attributes in the age of the machine*. Boston, MA: Addison-Wesley Longman Publ.
- Norman, D. A. (2002). *The design of everyday things*. New York: Basic Books.
- Rupert, R. D. (2013). Distributed cognition and extended-mind theory. In B. Kaldis (Ed.), *Encyclopedia of philosophy and the social sciences*. Los Angeles: SAGE Publications.
- Thagard, P. (2014). Cognitive science. Part 3. In E. N. Zalta et al. (Eds.), *Stanford encyclopedia of philosophy*. Retrieved August 20, 2014, from <http://plato.stanford.edu/entries/cognitive-science>.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Zhang, J., & Patel, V. (2006). Distributed cognition, representation, and affordance. *Cognition & Pragmatics*, 14(2), 333–341.
- Zhang, J. (1991). The interaction of internal and external representations in a problem solving task. *Proceedings of the Thirteenth Annual Conference of Cognitive Science Society* (pp. 954–958). Hillsdale. Retrieved December 11, 2013, from <http://citeseerx.ist.psu.edu>
- Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, 21(2), 179–217.