

The cognitive integration of scientific instruments: information, situated cognition, and scientific practice

Richard Heersmink^{1,2}

Published online: 1 July 2015

© Springer Science+Business Media Dordrecht 2015

Abstract Researchers in the biological and biomedical sciences, particularly those working in laboratories, use a variety of artifacts to help them perform their cognitive tasks. This paper analyses the relationship between researchers and cognitive artifacts in terms of integration. It first distinguishes different categories of cognitive artifacts used in biological practice on the basis of their informational properties. This results in a novel classification of scientific instruments, conducive to an analysis of the cognitive interactions between researchers and artifacts. It then uses a multidimensional framework in line with complementarity-based extended and distributed cognition theory to conceptualize how deeply instruments in different informational categories are integrated into the cognitive systems of their users. The paper concludes that the degree of integration depends on various factors, including the amount of informational malleability, the intensity and kind of information flow between agent and artifact, the trustworthiness of the information, the procedural and informational transparency, and the degree of individualisation.

Keywords Cognitive artifacts · Scientific instruments · Distributed cognition · Extended mind · Scientific practice

1 Introduction

The goal of this paper is to better understand the interactions and integration between researchers and cognitive artifacts in biology laboratories. It draws on and contributes to independent debates in situated cognition theory (Clark and Chalmers 1998; Menary 2007; Kirsh 2009; Rowlands 2009; Sutton 2010) and the philosophy of scientific

✉ Richard Heersmink
richard.heersmink@gmail.com

¹ Department of Cognitive Science, Macquarie University, Sydney, Australia

² Department of Philosophy, Macquarie University, Sydney, Australia

practice (Giere 2002a, b; Nersessian 2005, 2009). By doing so, it produces mutual benefits for both situated cognition theory and the philosophy of scientific practice by bringing issues from the other area to bear. The analysis is set up in two broad stages. In the first stage, I develop a classification system of cognitive artifacts in biological practice. To this end, I begin by briefly outlining a previously developed method for classifying cognitive artifacts on the basis of their informational properties (Heersmink 2013). Next, I apply this informational approach to classify artifacts in biology laboratories (section “[Classifying cognitive artifacts in biological practice](#)”).

In the second stage, I conceptualize how researchers and artifacts interact and are, in some cases, integrated into larger distributed cognitive systems. I first survey and evaluate how a number of philosophers have used a distributed cognition approach to theorize about agent-artifact interaction in scientific practice (Giere 2002a, b; Nersessian 2005, 2009) and point out some shortcomings of their approaches (section “[Distributed cognition in scientific practice](#)”). I then present a previously developed multidimensional framework to conceptualize the degree of cognitive integration between humans and artifacts (Heersmink 2014). This framework consists of various dimensions that are all matters of degree and jointly constitute a multidimensional space in which situated cognitive systems can be located, including those that are embedded and extended/distributed (section “[Cognitive integration](#)”). This framework is then used to analyse the degree of integration between researchers and artifacts by focussing on four case studies: (1) computer models of protein folding, (2) pH-meters, (3) laboratory notebooks, and (4) organised workplaces (section “[Conceptualizing cognitive integration: case studies](#)”).

2 Classifying cognitive artifacts in biological practice

With the exception of New Experimentalism (e.g., Hacking 1983; Ackerman 1985), the role of instruments in scientific practice has not played a prominent role in traditional philosophy of science. Only recently have philosophers of science showed interest in instruments and their functional role in experiments and the creation of scientific knowledge (Baird 2003; Harré 2003; Boon 2004; Record 2010). Some of these philosophers developed classifications of scientific instruments (Baird 2003; Harré 2003; Boon 2004), classifying what Ackerman (1985) calls the “instrumentarium”, which are all instruments used by scientists, including those that do not have cognitive functions. Because current classifications do not focus on cognitive artifacts, but on instruments in general, they have overlooked some categories of cognitive artifacts. Specifically, they have not included what I below refer to as “symbolic cognitive artifacts” and “ecological cognitive artifacts”. To overcome this issue and develop a more inclusive picture, I now propose an alternative classification of cognitive artifacts used in biology labs in terms of their informational properties.

2.1 Informational properties of artifacts

Given the importance of the nature of external information in supporting cognitive tasks, (Heersmink 2013) classifies cognitive artifacts on the basis of their distinct informational properties. In the classification, I made a distinction between artifacts

exhibiting “representational” and “ecological” properties. Those with representational properties function as stand-ins for their target system and have aboutness or representational content. Examples include maps, traffic signs, thermometers, scale models, timetables, textbooks, and many other artifacts. Those with ecological or non-representational properties, by contrast, do not function as stand-ins for their target and do not have content. Examples include leaving a rented DVD on your desk as a reminder to bring it back to the video store, or reorganising letter tiles when playing Scrabble to prompt word recall.

Further distinctions within these two broad categories were then made. In case of representational artifacts, a distinction between icons, indices, and symbols was made (Peirce 1935). Icons display isomorphism between icon and target. A map, for example, is isomorphic to what it represents. Indices have a direct causal connection to their target. A thermometer is directly linked to the temperature. If the temperature changes, the reading on a thermometer changes as well. Symbols do not have isomorphism or a direct connection to their target, but obtain their meaning through convention, logical rules, and shared use. Language and mathematics are clear examples. Some artifacts have a combination of two or more of these representational properties. A map is isomorphic but often also contains symbols such as words and numbers. Likewise, a thermometer is indexical, but typically also contains symbols such as numbers. A representational artifact is thus often predominantly iconic, indexical, or symbolic.

In case of ecological artifacts, a distinction between spatial and dynamic artifacts was made. Spatial ecological artifacts get their cognitive function because we have placed these objects, such as DVDs, deliberately in a usual or unusual location such that when we see the artifact in that location it prompts a memory. Likewise, cooks organize their work-environment such that the location of the utensils and ingredients facilitate the task they are doing. Cognitive scientist David Kirsh (1995) refers to this as “the intelligent use of space”. Kirsh makes a tripartite distinction between spatial arrangements that simplify choice, perception, or internal computation. Thus the cognitive function of such spatial arrangements is not representational, but to simplify decision-making, perception and recognition, and internal cognitive processing. These examples of self-organized environments do not exhibit aboutness or content, because they neither have isomorphism or a direct causal connection to their targets, nor are they based on convention, logical rules, or shared use. Rather, by idiosyncratically putting artifacts in certain locations that are either deliberately usual or unusual, we intentionally encode task-relevant information into the artifact and its location.

Dynamic ecological artifacts obtain their cognitive function in virtue of their manipulable physical structure. When playing Scrabble, for example, players sometimes reorganise the letter tiles, as to aid word recall (Kirsh 1995, 2009; Clark and Chalmers 1998). This is very difficult to do internally as our mental imagery is rather limited, so most players reorganize letter tiles to help them come up with new words. The point here is that the cognitive function of rearranging letter tiles is not based on representational properties (as is the case with for example maps, thermometers, or written language). Rather, their cognitive function is based on their manipulable physical structures, and by manipulating their physical structure one automatically manipulates the information they contain as well. Note that individual letters are not representational, they become representational only when combined into words. Compared to spatial artifacts such as putting rented DVDs at some place, rearranging letter

tiles is a much more dynamic process. It is not so much about developing static spatial categories in which artifacts are placed, but more about dynamically and constantly changing information in an ongoing cognitive task.

In a sense, this is not so much a classification of *artifacts* but more of the *informational properties* they exhibit. In developing this classification system, an informational approach was chosen because the way cognitive artifacts are used depends on the information they contain. A map, for instance, is useful for the cognitive task of navigating because the predominantly iconic information it exhibits aids decision-making processes in relation to navigating. There are at least two reasons why this taxonomy is helpful in better understanding scientific practice. First, it allows us to categorize scientific instruments in terms of their informational properties, resulting in a novel way of classifying scientific instruments, conducive to an analysis of the cognitive interactions between researcher and artifact. Second, as we will see below, it helps identify overlooked categories of cognitive artifacts in scientific practice.

2.2 Representational artifacts

Having presented a general method for classifying cognitive artifacts, I now use the above categories to group artifacts used in biology labs. Icons are characterized by exhibiting isomorphism with their targets and there are many of those in biological practice. The most prominent of these are photographs, diagrams, and scientific models. Of all the cognitive artifacts used in science, diagrams and models have received the most attention from philosophers (e.g., Perini 2005; Knuuttila 2005, 2011; Toon 2012; Charbonneau 2013; Sheredos et al. 2014). Diagrams and models are important because they aid the cognitive processes of their users. When theorizing about the interactions between agents and models Tarja Knuuttila says: “From this perspective models function as external tools for thinking, the construction and manipulation of which are crucial to their epistemic functioning” (Knuuttila 2011, p 262). There are countless models used in biology that are isomorphic to their targets, for instance two-dimensional and three-dimensional models of the structure of biomolecules. An historically important example is James Watson and Francis Crick’s physical scale-model of the structure of DNA. There also models that represent processes and mechanisms like simulations of cell division, gene expression, metabolism, endocytosis, or protein folding. However, it is important to note that models need not be isomorphic (Suarez 2003). Many models in biology (for example in genetics) are statistical or mathematical in nature, showing no similarity to their target system. Some models may thus also be classified as symbolic in nature.

Indices are characterized by a direct causal connection between the representational state of the index and its target system. Biologists use a variety of indices, including graduated cylinders to measure the amount of liquid, thermometers to measure temperature, hemocytometers to count to number of cells in a solution, pH meters to measure the acidity of a solution, scales to measure the weight of some substance, pycnometers to measure fluid density, etc. Indices are essential to many scientific practices, including those in biology. Science investigates the world empirically and the best way to do so is by measuring it with the aid of indices. Moreover, given that biology often deals with microscopic structures, which are inherently invisible to the

unaided human eye, there is only one way to empirically obtain information about such target systems and that is with indices.

Symbols obtain their meaning and content through shared use, agreement, and logical rules. Typical examples are natural and artificial languages, numbers, mathematical systems, tables, equations, and scientific formulae. Symbols and symbolic structures are ubiquitous in biology labs, as they are found in manuals for equipment, textbooks, scientific articles, patents, and laboratory notebooks. They are also found on indices like thermometers, voltmeters, and other indices that indicate the quantity of their target systems in numbers. A paradigm symbolic structure in molecular biology is the periodic table of elements, which is a table of the chemical elements organised on the basis of their atomic numbers. There is no isomorphism or direct causal connection between the table and its target, the order in which the elements are presented is based purely on logical rules and agreement. Another prominent symbolic cognitive artifact is the laboratory notebook. Researchers working in biology labs have a personalised notebook in which they document their hypotheses, experimental procedures and outcomes, and observations made during the experiment. The majority of information in laboratory notebooks is symbolic in nature (e.g., language, equations, tables, calculations). I therefore classify notebooks as a predominantly symbolic cognitive artifact, but they may also contain iconic representations such as diagrams or sketches (Yeo 2008).

2.3 Ecological cognitive artifacts

Spatial ecological artifacts obtain their cognitive function in virtue of physical-spatial structures. Specifically, by intentionally putting artifacts in usual or unusual places we encode task-relevant information into the artifact and its location. This also happens in biology labs. A researcher may, for instance, intentionally leave an alcohol bottle at an unusual location on the laboratory bench as a reminder to clean and disinfect the bench after doing an experiment. Researchers may also structure their work-environment in more complex ways. When performing an experiment, they use a variety of instruments like, for example, pipets, test tubes, beakers, flasks, racks, stopwatches, heaters, etc. Experimental procedures require that these instruments are used in a certain order. Novice researchers often use a laboratory notebook in which the experimental procedure is outlined, which helps them to remember the steps in the experiment. But more experienced researchers may have developed strategies that make the notebook partly superfluous.

Consider the following example. The main steps in a DNA isolation procedure are: (1) break cells open by adding lysozyme, (2) remove membrane lipids by adding a detergent, (3) remove proteins by adding protease, (4) remove RNA by adding RNase, and (5) purify DNA by ethanol precipitation. Given the order of these steps, it makes sense to organize the workplace such that the lysozyme is located close to the experimenter, as this is the chemical s/he needs first. Likewise, the detergent should also be located quite close as this is the second chemical s/he needs in the overall procedure, etc. How the workbench and the artifacts are organized is important for performing the cognitive task at hand. Researchers therefore prepare the laboratory bench before they begin the experiment. They will most likely also create locations on the workbench that are reserved for artifacts that have been used and are no longer needed. Thus when the first step is performed, the bottle containing lysozyme will be

put at a spot that indicates it has been used. For example, the chemicals needed for the experiment may be put on her right hand side in the same order in which they are used in the experiment, and the chemicals s/he has used may be put on her left hand side, thereby creating spatial boundaries that help the researcher in performing the experiment. I was unable to find examples of dynamic ecological artifacts in biological practice in the literature. I do, however, think it is likely that such artifacts and practices occur in biology, but this hypothesis can only be tested by empirical observations.

3 Distributed cognition in scientific practice

Before I conceptualize how some of the above classified cognitive artifacts and their users are integrated into larger cognitive systems, I briefly survey a small body of literature concerning distributed cognition in scientific practice.

3.1 Ronald Giere

Giere uses a distributed cognition framework to better understand certain aspects of scientific research. For example, Giere (2002a; see also Cheon 2014; Toon 2014; Pence et al. 2015) investigates how scientific models can be part of distributed cognitive systems. He says that models are not just aids for cognition, but are when properly manipulated, part of a distributed cognitive process. He emphasises that humans cannot store, let alone manipulate, complex scientific models in their head. They need to create and manipulate external representations that complement the shortcomings of the pattern matching abilities of embodied brains. He writes:

“Physical models provide what is probably the best case for understanding model-based reasoning as an example of distributed cognition. Here it is very clear that one need not be performing logical operations on an internal representation. It is sufficient to perform and observe appropriate physical operations” (Giere 2002a, p. 9).

Giere makes an interesting claim about the affordances of models exhibiting different kinds of representational formats (see also Vorms 2012). He claims that Watson and Crick with their physical three-dimensional model were able to discover the correct structure of DNA, whereas Rosalind Franklin, who only had two-dimensional X-ray photographs and hand-drawn diagrams, was unable to discover the structure of DNA. There may, of course, have been other factors involved in the discovery of the structure of DNA, but Giere nonetheless is onto something here. The specific informational properties and affordances of the model are important for performing the cognitive task. So physical models, in this case, are better for performing this cognitive task as compared to photographs and sketches on paper, because they complement onboard cognitive systems in a more effective way.

Giere also theorizes about larger distributed cognitive systems. When talking about the Indiana University Cyclotron Facility, he says that “It is particularly enlightening to think of the whole facility as one big cognitive system comprised, in part, of lots of

smaller cognitive systems” (Giere 2002b, p. 8). He then goes on to list some of the components of this big cognitive system, including the accelerator, detectors, computers, and the people actively working on the experiment. Likewise, he writes that the Hubble Telescope and the researchers using it can be seen as one distributed cognitive system. A single experiment may include fifty to five-hundred people and countless computing systems processing the data integrated into one distributed cognitive system.

As a criterion for thinking about when something is part of a distributed cognitive system Giere writes: “we can distinguish those features of the system that differentially influence the output of the system in scientifically relevant ways from those features that merely make it possible for the system to generate any output at all” (Giere 2002b, p. 11). Researchers interacting with instrumentation have an influence over the outcome of experiments, so both researchers and instruments are part of the system, whereas the electricity powering the instruments does not differentially influence the particular output and is thus not part of the distributed cognitive system.

3.2 Nancy Nersessian

Nersessian (2005, 2008, 2009) also uses a distributed cognition framework to analyse scientific practice. She writes: “Within this framing, to understand how problem solving is achieved requires examining the generation, manipulation, and propagation of salient representations within the system; that is, examining how representational states flow across media and accomplish cognitive work” (Nersessian 2009, p. 732). Her work is particularly relevant for this paper as she focusses on biomedical engineering (BME) labs. Such labs differ from the paradigm cases of distributed cognitive systems that Ed Hutchins (1995) has analysed. Nersessian points out that navigators on board of a ship face problems that change over time, but the artifacts and the knowledge that the agents bring to the situation remain largely stable. Such systems are thus dynamic but *synchronic*, so the functional and informational properties of the cognitive artifacts remain largely the same over time. By contrast, a BME lab can be seen as an evolving distributed cognitive system, so both the knowledge of the scientists and the functional and informational properties of the artifacts and models change over time. Thus the BME lab as a distributed cognitive system is dynamic and *diachronic*: researchers and artifacts change in relation to each other. This dynamic and diachronic relation between researcher and cognitive artifact is characterized as a “cognitive partnership”. “We use the notion of cognitive partnership to characterize the special relations that develop among researchers and between them and simulation devices (as opposed to other important artifacts such as the pipette or confocal microscope)” (Nersessian 2009, p. 741). Cognitive partnering thus does not occur with simple instruments but only with simulation devices, which are engineered models of biological target systems.

3.3 Discussion

The above approaches are perceptive and I am largely in agreement with them. There is, however, some room for improvement. Giere’s criterion for determining when something is part of a distributed cognitive system is not sufficiently clear and perhaps too liberal to be helpful. On his view, agents or artifacts have to differentially influence

the output of the system in order to be part of the distributed cognitive system. This is an interesting way to think about system membership in general, but it might be too liberal as a criterion for membership of distributed cognitive systems. This is so because it results in what critics of extended cognition have called “cognitive bloat”, which is roughly the idea that cognition extends into too many things (see e.g., Rupert 2004). If we include everything that differentially influences the output of a cognitive system, it results in cognitive systems that are distributed too widely. This criterion may be one of the reasons why some of Giere’s examples are perhaps somewhat unsupported. For instance, he argued that entire research facilities such as the Indiana University Cyclotron Facility, including all the people performing experiments, constitute one large distributed cognitive system. It is one thing to claim that entire research facilities constitute one distributed cognitive system, but it is quite another thing to explain in detail how and why hundreds of researchers interacting with instruments constitute such a system.

J. Adam Carter et al. (2014) distinguish between three varieties of active externalism: the extended mind, the extended cognition, and the distributed cognition theses. On their view, the first focuses on extended *mental states* of individuals, e.g., when Otto, a man with Alzheimer’s disease, stores his beliefs in a notebook. The second focuses on extended *cognitive processes* of individuals, e.g., when someone performs a calculation with pen and paper. The third focuses on *cognitive processes* that are distributed across members of a social group interacting with cognitive artifacts, e.g., two pilots interacting with instruments in a cockpit. A reviewer pointed out that the cognitive bloat objection plausibly applies to the extended mind thesis, but not to the extended and distributed cognition theses. Or in Carter et al.’s words: “because the existence of extended mental states—such as extended dispositional beliefs—is a claim that is more counterintuitive and thereby less easy to argue for than the claim that there are extended cognitive processes” (Carter et al. 2014, p. 11). I am sympathetic to their intuition but I am not sure whether this means that the cognitive bloat objection cannot apply to extended or distributed cognitive processes. If one’s conditions for cognitive extension or distribution are not sufficiently stringent, it may still result in a conceptualisation of processes that are extended or distributed too widely.

Nersessian’s account is elaborate, detailed, and insightful, but does not sufficiently look at the conditions for cognitive distribution. Her notion of a cognitive partnership is largely descriptive, not stipulative. So, on her view, it is difficult to determine when something is part of a distributed cognitive system. I think both Giere and Nersessian’s account can be enriched by incorporating elements from the multidimensional framework presented below. In particular, to account for the diachronic aspects of distributed cognitive systems. I will briefly come back to this point below in section “Discussion”. A second and perhaps less important shortcoming of both Giere and Nersessian’s account is that they focus on representational artifacts (diagrams, models, and certain instruments), in that way neglecting what I referred to as ecological artifacts. In order to provide a more complete picture of the variety of situated cognitive systems in biology labs (or any other laboratory or setting), we need to include ecological cognitive artifacts in the analysis.

On the basis of Giere and Nersessian’s work, one may distinguish two approaches to studying distributed cognitive systems in scientific laboratories. One can start by identifying a large-scale system, such as the Indiana University Cyclotron Facility,

and then try to identify all the relevant components and see how information trajectories are established and how all the components are integrated into larger distributed systems. This may be called a top-down approach. Given that there are sometimes more than a hundred researchers involved and numerous artifacts and instruments, the information trajectories and the overall system are often so complicated that no single ethnographer can observe such a system. So if one chooses a top-down approach, it is only feasible when the system is sufficiently small to observe, which might in practice be quite difficult. This is presumably why Giere, at least in some of his examples, theorizes at a very general level of explanation without giving specific details of the composition and integration of the large-scale systems he is theorizing about. By contrast, one may also start at identifying small-scale systems, comprising of an individual researcher and the artifact(s) he or she is interacting with. If that small-scale system is part of a larger one, it may be possible to work your way up and include more components in the analysis. This may be called a bottom-up approach. In the next section, a bottom-up approach is developed and then applied to conceptualize how researchers interact with their instruments and how they are, in some cases, integrated into a larger distributed cognitive system.

4 Cognitive integration

4.1 Parity versus complementarity

In addition to Clark and Chalmers (1998) conditions of reliability, trust, accessibility, and past endorsement, much extended mind theorizing has been motivated by the parity principle:

“If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process” (Clark and Chalmers 1998, p. 8).

This principle invites us to see isomorphism between the inner (biological) and the outer (artifactual) as a condition for cognitive extension. Some have argued that there is no isomorphism between the inner and the outer: artifacts and other resources have informational properties that are quite different from those of the brain. Information stored in notebooks, for example, is not automatically integrated with other information in the notebook and does not show the primacy or recency effect. John Sutton (2010) therefore argues that it is much better to think of the functions and properties of external resources as *complementary* to those of the brain (see also Menary 2007; Kiverstein and Farina 2011). On such a view:

“Contrary to any requirement of fine-grained similarity then, what the friends of extended cognition actually expect, and study, are hybrid processes in which the inner and the outer contributions are typically highly distinct in nature, yet deeply integrated and complementary” (Wilson and Clark 2009, p. 72).

Cognitive integration of agent and resource is thus an important theme within a complementarity view. This integration, I suggest, is a matter of degree and varies along a number of dimensions.

4.2 A multidimensional approach

Building on Sutton (2006, 2010), Menary (2007, 2010), Wilson and Clark (2009), and Sterelny (2010), (Heersmink 2014) develops a multidimensional framework to conceptualize the interactions and degree of cognitive integration between agent and artifact.¹ These dimensions provide a new perspective for thinking about when an artifact or other external resource is part of an extended or distributed cognitive system: Not in terms of necessary and sufficient conditions, but in terms of dimensions that are matters of degree. This provides a more realistic view on such systems. Conceiving of situated cognitive systems as either embedded/scaffolded or extended/distributed with a clear boundary in between is not a realistic way to describe such systems, because some may be more embedded or extended than others (Sutton et al. 2010).

The nature of their embeddedness or extendedness, I suggest, depends on the degree of integration. So, rather than providing a set of necessary and sufficient conditions for cognitive extension and distributed cognition, this framework provides a toolbox for investigating the degree and nature of integration of agents and artifacts into “new systemic wholes” (Clark 2007). Situated systems that rank high on these dimensions are deeply integrated and clear candidates for extended/distributed cognitive systems. Whereas those that rank low on these dimensions are shallowly integrated and are clear cases of embedded cognitive systems. In between there is a grey area in which it may not always be clear whether a system is extended or merely embedded. Below I concisely outline the eight dimensions, but for a more detailed description of each dimension see (Heersmink 2014; for discussion see Clowes 2013).

4.2.1 Information flow

In situated cognitive systems, information flow between agent and artifact may be one-way, two-way, or reciprocal. One-way information flow goes from artifact to agent. Examples include using a map to navigate or a timetable to decide which train to take. Two-way information flow goes from agent to artifact and then from artifact to agent. Examples include writing an appointment in your diary and looking it up at some later point or deliberately leaving a rented DVD on your desk so that you will bring it back. Reciprocal information flow concern more complicated informational exchanges between agent and artifact. In such cases, the artifact is part of an ongoing cognitive task. Examples include making a difficult calculation with pen and paper, reorganizing Scrabble tiles, or writing an academic paper. These examples are often seen as the paradigm cases of extended cognitive systems. In such cases, there is not one cycle of

¹ See Duncan Pritchard (2010) for a related discussion in epistemology on the conditions of extended knowledge. In future research it may be interesting to compare these conditions of extended knowledge with my multidimensional framework.

offloading and intake of information, but various incremental and interdependent cycles of informational offloading and intake.

4.2.2 Accessibility

How accessible an artifact is, is important for the degree of integration (Clark and Chalmers 1998). If the artifact is not reliably available when needed, it cannot be used for performing some cognitive task. So the more reliable the artifact is accessible, the easier it can be integrated into the cognitive system of its user. Mobile computing devices, for example, are highly accessible and thus easily integrated, whereas library books only when you go to the library. They are much less accessible and thus more difficult to integrate into the task.

4.2.3 Durability

This dimension concerns how often we use an artifact. Following Wilson and Clark (2009), a distinction can be made between one-offs, repeated, and permanent relationships to external resources. A token shopping-list is typically a one-off, a token map is often repeatedly used, and a token laptop computer is used many times over a long period of time. Generally, the more often we use an artifact, the more we rely on it and the deeper it is integrated into our cognitive system.

4.2.4 Trust

Trust is about our cognitive attitude towards the truth-value of information. Whether we trust external information is important, because information we do not trust, we usually do not use to help us perform our cognitive tasks. Trust can be implicit or explicit. When we implicitly trust information, we have not consciously evaluated it, but assume it is true. When we explicitly trust information, we have consciously evaluated it and concluded it is true. For optimal cognitive performance, we implicitly trust information.

4.2.5 Procedural transparency

Procedural transparency concerns the effortlessness and lack of conscious attention with which a cognitive artifact is used. The easier it is to use and interact with an artifact, the more procedurally transparent it is. For instance, the first time I use a computer, it is hard for me to use it smoothly, but when after a while my perceptual-motor processes are proceduralized to such an extent that I do not consciously think about how to interact with it.

4.2.6 Informational transparency

Informational transparency concerns the effortlessness with which an agent can interpret and use information. The easier it is to interpret and understand external information, the more transparent it is. Scientific equitations, for example, may be opaque for novices but fully transparent for experts. The degree of informational transparency usually depends on training and education.

4.2.7 Individualization

Some cognitive artifacts are interchangeable whereas others are individualized (Sterelny 2010). Those that are interchangeable include timetables, maps, clocks, thermometers, and others. Those that are individualized include diaries, notebooks, shopping-lists, and bookmarks of websites. Such artifacts are tailored for individual cognitive goals, which typically increases the degree of integration between agent and artifact.

4.2.8 Transformation

Interacting with cognitive artifacts transforms our onboard cognitive system. During ontogeny, we internalize publically available representations like language and number systems and learn to think in terms of those systems (Menary and Kirchhoff 2014). In general, the longer we interact with cognitive artifacts, the more they transform our cognitive system. Likewise, the informational properties of some cognitive artifacts are static (e.g., textbooks), but others are dynamic and transform their properties over time (e.g., computer simulations).

5 Conceptualizing cognitive integration: case studies

In this section, I apply these dimensions to conceptualize the interactions and degree of integration between researchers and artifacts, demonstrating the usefulness of situated cognition theory for the philosophy of scientific practice (see also Arnau et al. 2014). I focus on four case studies: computer models (icons), pH-meters (indices), laboratory notebooks (symbols), and organised work-environments (spatial ecological artifacts). These artifacts have different informational properties and are chosen to examine whether their specific informational properties are integrated into the cognitive systems of their users to different degrees.

5.1 Representational artifacts

5.1.1 Computer models of protein folding

Biologists often use computer models to gain knowledge about some target system (see e.g., Humphrey 2009; Winsberg 2010; Beisbart 2012). Here I take a closer look at computer simulations of protein folding. After a protein is synthesized by a ribosome in a cell, it starts as a long chain of amino acids referred to as its primary structure. This primary structure first coils into a number of secondary structures such as alpha-helices and beta-sheets, which in turn, fold into the eventual three-dimensional tertiary structure. A protein's tertiary structure determines its function and is therefore important to better understand. The problem molecular biologists face is how a long chain of amino acids eventually results in the tertiary structure of the protein. This folding process is quite complicated, not yet fully understood, and very difficult to study *in vivo* (i.e., within cells). For these and other reasons, molecular biologists use computer simulations of the folding process.

There are two main stages in the simulation process: building the simulation and running it. In most cases, the amino acid sequence of a protein is known and used as input for the simulation. Amino acids have certain properties that determine how they will interact with other amino acids. The software program contains all the necessary information to simulate the interactions between different amino acids such as information about bonds, angles, torsion angles, and electrostatic interactions. When the simulation has been successfully built, it is run and visualized on a computer screen. The *in vivo* folding process takes milliseconds, but the simulation can be run as slow as a researcher wants. Typically the simulation is slowed down a hundred to a thousand times, in order to scrutinize each step in the folding process.

Simulations are clear cases of cognitive artifacts that complement the shortcomings of our onboard cognitive capacities. Trying to mentally visualize a static three-dimensional structure such as a protein is already quite difficult, if possible at all, but trying to mentally visualize the folding of a protein over time, is for most people impossible. During a simulation, the angle of the overall protein and speed of the simulation can be altered, as to optimize perceptual access to salient properties of the folding process. One can also zoom-in on a particular part of the protein. Thus, in terms of information flow, a modeller first offloads information onto the computer system, then the computer creates the simulation with which a modeller interacts in an ongoing and iterative manner, thereby creating a reciprocal information flow structure.

Access to the simulation is reliable as it is run on computer systems that are typically quite reliable. Further, because the simulation can be viewed from different angles, slowed down or speeded up, and zoomed-in as to improve perceptual access to salient parts of the folding process, it is fair to say that information access is highly reliable. Because programming the simulation can take hours to days and analysing the simulation usually takes many hours as well, the durability of the coupling is long. Simulations are usually analysed a number of times, so a repeated but not permanent relationship is established.

Trust in the correctness and truth-value of the simulation is generally high, but this depends on the kind of simulation. The more that is known about a particular protein, the more trust a modeller puts into its simulation. For example, when the tertiary structure is known via crystallography but not its folding process, and when a simulation results in a similar or identical tertiary structure, then there is little reason to distrust the correctness of the simulation. But when simulating the folding of a protein whose tertiary structure is unknown, then a modeller may be more cautious in accepting the accuracy of the simulation. Note that accuracy here means degree of isomorphism between the simulation and the actual folding process. Simulations may of course be wrong in that the simulation is not isomorphic to the actual folding process. For example, in some cases, the tertiary structure is known but not its folding process. When a simulation results in a protein structure that is not similar to its known tertiary structure, then modellers have good reasons to distrust the simulation. If a high degree of trust is established, is it always after a modeller has consciously evaluated the simulation, so trust in case of simulations always concerns explicit trust.

The degree of procedural transparency depends on the degree of understanding of the software program, which may take some time and experience to develop. Using simulation software requires knowledge about both molecular biology and computer programming and the degree of transparency depends on this knowledge. For a novice,

transparency is typically low, but for an experienced modeller it is high, with various degrees of transparency in between. The informational transparency is most likely high, as simulations of protein folding are highly isomorphic to the actual folding process. Even for complete novices, the basics of a simulation are relatively easy to understand, partly because simulations are approximations, simplifications, and idealisations of their target systems, which makes them easier to understand and easier to study.

Whether simulations are individualized depends on the research goals of the modeller. In some cases, simulations are performed by individual researchers working on their specific projects, in which case the simulation is highly individualized. In other cases, simulations are performed by a small team of researchers, in which case they are less individualized. But either way, the outcomes of simulations are quite often published and are thus aimed at a larger audience and are performed to contribute to a better general understanding of molecular biology. In this sense, simulations are not individualized, but part of a publically accessible body of knowledge.

Computer simulations of protein folding and protein structure have most likely transformed the representational capacities of molecular biologists. Visualizing and perceiving how an unobservable such as a protein behaves *in vivo*, has changed how molecular biologists and others think about the structure of proteins. If a modeller spends many hours analysing simulations of protein folding, then it is likely that the modeller will have internalized at least some parts of the simulation, but it is difficult to quantify how much computer simulations have transformed the representational capacities of their users.

More generally, scientific models are representational systems that the brain is likely to absorb relatively easily. Particularly models that are simplified and idealised are easy to internalize. For example, textbook models of the solar system, the anatomy of the human body, atomic structures, plate tectonics, and many other models have changed the way we think about the physical world, not only for scientists but also for the general public, as some of these models have found their way into popular culture. This is because models make complicated processes or structures relatively easy to understand, as their simplified and idealised isomorphic format is easy to interpret and internalize.

In sum, computer simulations of protein folding rank relatively high on most dimensions. They rank medium on durability, but high on reliability and trust and after some training they also rank high on procedural and informational transparency. They may or may not be individualized, depending on each case, and it is difficult to quantify how much they have transformed the representational capacities of their users, but it is safe to say that there is some degree of transformation. But, most importantly, there is a reciprocal information flow which means there is a significant degree of integration between agent and simulation. Thus, given how computer simulations of protein folding rank on the dimensions, it populates a higher region in the multidimensional space and so integration between researcher and simulation is quite dense.

5.1.2 *pH-meters*

A common procedure in biology is to measure the acidity or alkalinity of a liquid, i.e., its pH. In order to obtain the pH value, a series of nested actions needs to be performed: activating the device, dipping the measuring probe in the liquid, and interpreting its reading. All this is typically done in a few seconds, resulting in a pH value that is

measured to two decimals. When measuring pH, information flow is one-way, i.e., from artifact to researcher. Access to pH-meters is highly reliable, as there are usually a number of pH-meters available in any biology lab. Moreover, pH-meters are regularly calibrated to ensure accurate readings, which increases the reliability of the information. Measuring pH is such a common procedure that the relation to pH-meters is often repeated and thus a long-term, permanent relationship is established. Trust in the accuracy of the reading is high, since pH-meters are regularly calibrated. Although there may be cases when the reading is unusual which may prompt suspicion or distrust, typically a researcher implicitly trusts the information it provides, i.e., s/he does not consciously evaluate it. For experienced researchers, the procedural and informational transparency are high, because they are trained in using pH-meters and have used them countless of times. The process is proceduralized and they do not have to consciously think about how to interact with the device or how to interpret its reading. pH-meters are not individualized, but are interchangeable and used by all members of the laboratory.

Finally, pH-meters may not have transformed the representational capacities of the researchers' brains in the same way as scientific models have, but they have transformed how we think about acidity and alkalinity. They have also transformed many cognitive practices in biology labs, because being able to accurately measure pH is important for many experimental procedures. The representational state of the device is dynamic and will change when the target system changes. However, in terms of Nersessian's notion of cognitive partnership, which is characterized by an evolving and dynamic relation between agent and artifact, pH-meters do not qualify as artifacts with which a cognitive partnership is established, because their functional properties do not change over time.

In sum, pH-meters rank high on most dimensions. They rank high on all dimensions except on individualization and transformation on which they rank low, and, importantly, the information flow is one-way. Given how it ranks on the above dimensions, this situated cognitive system populates a region somewhere in the lower middle regions of the multidimensional space. Indices do not afford two-way or reciprocal information flow, because they have a direct causal connection to their target systems. One can change their informational content only by altering the target system. So one can reason *about* the information an index provides and one may use that information to guide further action, but one cannot reason *with* indices in the same way as one can with icons and symbols. This is so because they lack representational malleability, i.e., the capacity to change their content in an ongoing cognitive task. It is therefore difficult for an index and its user to be deeply integrated into a larger system. Recall that one of the reasons for taxonomizing cognitive artifacts was to examine whether artifacts in different categories, i.e., with different informational properties, have different effects on cognition and are integrated into the user's cognitive system to different degrees. Because indices lack the capacity for two-way or reciprocal information flow, they also lack the capacity for deep integration.

5.1.3 Laboratory notebooks

pH-meters have only one cognitive function (to measure pH) and exhibit only one type of information (a numerical value of pH). Laboratory notebooks, by contrast, are much

more complicated as they serve a number of cognitive functions (mainly related to memory and reasoning) and exhibit various types of information (language, equations, tables, numbers, etc.). Lab notebooks are used for various cognitive tasks and serve a number of cognitive functions:

- When performing an experiment, it is very hard to remember all the steps as most experiments are rather complex and involve countless incremental steps. Therefore, lab notebooks always contain a section on methodology, which is typically written by the researcher, describing all the steps that need to be done in order to perform the experiment. In this role, the notebook functions as a to-do list, so that the researcher knows what to do and what has been done at any point in the experiment.
- Observations made during the experiment are written in the notebook. In this role, the notebook functions as a long-term external information-storage system, thereby complementing the shortcomings of internal storage systems.
- Experimental outcomes are written in the notebook, sometimes as linguistic descriptions but more often as tables, graphs, or other diagrams. In this way, researchers later know what they have done during an experiment and what the outcome of the experiment was. These experimental outcomes are often discussed during lab meetings where notebooks play a crucial role as memory aids.
- In addition to these predominantly memory aiding functions, notebooks are also used to solve equations, perform calculations, or draw graphs, tables, or other diagrams. In these roles, the notebook serves more as a facilitator of ongoing reasoning processes.

Depending on the task for which it is used, information flow between notebook and user is two-way or reciprocal. When an experimental procedure is first written in the notebook and then later used to help organise the experiment, information flow is two-way because there are two steps involved: offloading and intake. But when it is used to perform calculations, solve chemical equations, or draw graphs or tables, then a reciprocal system is established. This is so because performing these tasks typically involves many cycles of offloading and intake and each cycle depends on the outcome of the previous one. Each step in the overall cognitive task thus builds and depends on and the previous step and are therefore interdependent.

Access to the notebook is highly reliable, as it is standard procedure to always have it around when performing experiments. Some researchers make photocopies of their lab notebooks so that they have a backup of it, which increases reliability. The information itself is, in most cases, also highly reliable. But there may, of course, be cases in which mistakes are made when performing calculations, solve equations, and so on, in which case the information is less reliable or unreliable. Given that the notebook is used many times throughout the day, a repeated relationship is established. Most of the information in the notebook is implicitly trusted, i.e., not consciously evaluated, because it is written by the researcher, but also because some of the information is standardized. Most experimental procedures, for example, are standardized so that they need little evaluation when used to guide the experiment. But when interpreting experimental outcomes, more evaluation is needed and the data may not be

trusted by default. After conscious evaluation, a researcher may either explicitly trust the data or distrust it.

Interacting with notebooks is rather straightforward, so the procedural transparency is high. The information in it (e.g., sentences, calculations, equations) is written by the researcher, so it is easy to interpret and process. Although there may be cases in which a researcher has quickly written some observations in the notebook that are hard to interpret or understand at some later point, generally the informational transparency is high. The notebook is highly individualized and is deeply engrained in many of the cognitive practices of the individual user. The information in the notebook is meant to perform his/her specific experiments and the observations s/he writes in it result from her experiments. So it is as individualized as a cognitive artifact can be. Researchers also deeply depend on it. If the notebook got lost, a researcher would have trouble performing experiments and the experimental data would be lost.

Finally, the notebook most likely does not transform the representational capacities of their users. Only interacting with representational systems such as language, mathematics, and perhaps some notational systems in science like chemical equations or models will significantly transform the representational capacities of embodied brains. Single artifacts most likely do not transform the brain in the same way representational systems do. However, artifacts do (sometimes quite dramatically) transform our cognitive practices. Lab notebooks, for example, have significantly transformed many of the cognitive practices in laboratories (see, e.g., Tweney and Ayala 2015; Holmes et al. 2003). In sum, lab notebooks rank high on all dimensions, except transformation. They furthermore exhibit two-way and reciprocal information flow. Given this score, notebooks populate a higher region in the multidimensional space and, therefore, notebooks and their users are deeply integrated into a distributed cognitive system.

5.2 Ecological artifacts

Biologists working in laboratories organise their work-environment such that the location of the artifacts used in performing their experiments facilitates the cognitive tasks they are performing. This intelligent use of space reduces the cognitive load in perceptual and memory tasks, in that way complementing their cognitive processes. When preparing the experiment by intentionally putting artifacts in particular locations on one's workbench, task-relevant information is encoded into the artifacts and their locations. So information is first offloaded onto the environment by intentionally putting artifacts at certain locations and then taken onboard at some later point. A two-way information flow system is thus established. Kirsh (1995) makes a tripartite distinction between spatial arrangements that simplify choice, perception, or internal computation. Organised workplaces in laboratories seem to mainly simplify choice and perception. When scientific instruments are placed such that they facilitate the steps in performing an experiment, a researcher does not have to choose which steps to perform as these are embedded in the environment. Likewise, a researcher does not have to scan the work-environment to look for the relevant instrument, but knows they are all within reach. Decision-making and perception are thus made easier by structuring the environment.

Access to the information is highly reliable, as the information is always there when it is needed, i.e., when performing the experiment one is sitting at one's organised

workplace which is central in one's visual field. Performing an experiment can take thirty minutes or several days. So depending on each experiment, the duration of the coupling is medium to very long. Organised workplaces are repeatedly created and in that sense a permanent relationship is established with such ecological artifacts. However, given that each experiment may require different instruments, set-ups, and procedures - the ecological artifacts that are created for each experiment will have different structures and will have different informational properties. So there is no permanent relation established to specific ecological artifacts, unless an experiment is performed many times, but typically new ecological artifacts are created for each experiment.

The amount of trust in the correctness of the information is generally high, as the researcher herself has placed the artifacts in their correct location. When performing the experiment, a researcher most likely will not consciously evaluate whether the location of artifacts indicates the correct order of steps in the experiment, but will implicitly trust it is correct. The procedural transparency is high, as the researcher herself has put the artifacts at certain locations and s/he knows how to extract the task-relevant information. Similarly, the informational transparency is high, too, because the user has created the ecological artifacts herself. Organized workplaces are highly individualized, as it concerns the specific experiments of individual researchers. In some cases, however, experiments may be performed by two or more researchers, in which case they share an organised workspace. Finally, organised workplaces most likely do not transform the representational capacities of their creators, because they do not concern external representations.

In sum, organized workspaces rank low on transformation, the durability depends on the kind and complexity of the experiment, and they rank high on reliability, trust, procedural and informational transparency, and individualization. Given how organised workspaces rank on the dimensions, they populate an upper middle region of the multidimensional space and are thus not deeply integrated. They are therefore candidates for embedded cognitive systems.

5.3 Discussion

How does the proposed multidimensional framework relate to Giere and Nersessian's approaches? I criticized Giere for having a too liberal criterion for cognitive distribution and Nersessian for a lack of such a criterion. On my view, the degree of integration is a criterion for being part of a distributed cognitive system. The higher a situated cognitive system ranks on these dimensions, the denser the integration between agent and artifact, and the more likely it is that the artifact is part of a distributed cognitive system. The view I propose here is thus not one in terms of necessary and sufficient conditions, but one that puts membership of a distributed cognitive system on a spectrum.

The case studies demonstrate that the degree of cognitive integration between agent and artifact in biological practice can vary substantially such that different situated cognitive systems populate different regions in the multidimensional space (see Table 1). When an artifact is used for a longer period of time and the relationship to it becomes gradually more durable, individualized, transparent, and trustworthy, the overall system becomes increasingly more integrated. Generally, there is a tendency for situated systems to become integrated more deeply over time. However, most situated

Table 1 Overview of the case studies in terms of each dimension

Dimension	Iconic computer models	Indexical pH-meters	Symbolic lab notebooks	Ecological organised workplaces
Information flow	Reciprocal	One-way	Reciprocal	Two-way
Accessibility	High	High	High	High
Durability	Medium	High	High	Medium
Trust	Medium-High	High	High	High
P-transparency	High	High	High	High
I-transparency	High	High	High	High
Individualization	High	Low	High	High
Transformation	Low	Low	Low	Low
Integration	Dense	Shallow	Dense	Medium

cognitive systems are dynamic and can shift from one region of this multidimensional space to another. For example, when it turns out the pH of a liquid is 9.5, when one knows it is acidic and should thus be below 7, a researcher knows there is something wrong with the pH-meter. This will reduce the amount of trust in the reading of the meter and so it will rank lower on the dimension of trust, which will also reduce the amount of integration between agent and artifact. So, the degree and nature of integration may change over time, because there will always be (small) changes in one of the dimensions.

The framework presented above is thus able to deal with these diachronic aspects of situated cognitive systems and can contribute to a better understanding of what Nersessian has called a “cognitive partnership” between researcher and artifact. More concretely, the diachronic aspect of distributed cognitive systems can be further conceptualized by the dimensions of trust, procedural transparency, informational transparency, individualization, and transformation. These dimensions are meant to capture the diachronic nature of situated cognitive systems. Over time, researchers learn to trust the informational properties of the artifact. When they get more skilful in using and manipulating the artifact, the procedural transparency increases and they find it easier to interpret the information of the artifact and so the informational transparency increases as well. During an experiment, the artifact is individualized in terms of the epistemic goals of the experiment. The informational properties of the artifact are transformed and, depending on the amount of use, the artifact may also have transformed the onboard cognitive system of its user, i.e., the researcher learns to think in terms of the informational properties of the artifact.

6 Concluding summary

This paper first developed a classification system for cognitive artifacts used in biological practice. In doing so, I focussed on the informational properties of artifacts, because the way artifacts are used depends on the information they contain. Classifying artifacts in this way is more conducive to an analysis of the cognitive interactions

between researchers and artifacts as compared to other classifications. Having this classification system in place, I then used a multidimensional framework to conceptualize the degree of integration between researcher and various artifacts. I focussed on four case studies: computer models of protein folding, pH-meters, laboratory notebooks, and organised work-environments. Computer models and notebooks are deeply integrated because they have a high informational malleability and reciprocal information flow, whereas pH meters and all other indices lack this malleability and reciprocity. One important property of models (icons), notebooks (symbols) and also organised workspaces (spatial ecological artifacts) is that the users have intentional agency over the informational content of the artifact, which significantly increases the degree of integration. Other important factors include the trustworthiness of the information, the procedural and informational transparency, and the degree of individualisation.

Acknowledgments I wish to thank John Sutton and Richard Menary for constructive comments on an earlier draft of this paper, and the Collective Cognition Group at Macquarie University for funding to write it. I also wish to thank three anonymous reviewers. The second reviewer, in particular, gave very helpful feedback.

References

- Ackerman, R. (1985). *Data, instruments, and theory: A dialectical approach to understanding science*. Princeton: Princeton University Press.
- Amau, E., Estany, A., Gonzalez del Solar, R., & Sturm, T. (2014). The extended cognition thesis: its significance for the philosophy of (cognitive) science. *Philosophical Psychology*, 27(1), 1–18.
- Baird, D. (2003). Thing knowledge: Outline of a materialist theory of knowledge. In H. Radder (Ed.), *The philosophy of scientific experimentation* (pp. 39–67). Pittsburgh: University of Pittsburgh Press.
- Beisbart, C. (2012). How can computer simulations produce new knowledge? *European Journal for Philosophy of Science*, 2, 395–434.
- Boon, M. (2004). Technological instruments in scientific experimentation. *International Studies in the Philosophy of Science*, 18(2–3), 221–230.
- Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. (2014). Varieties of externalism. *Philosophical Issues*, 24(1), 63–109.
- Charbonneau, M. (2013). The cognitive life of mechanical molecular models. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(4), 585–594.
- Cheon, H. (2014). Distributed cognition in scientific context. *Journal for General Philosophy of Science*, 45(1), 23–33.
- Clark, A. (2007). Re-inventing ourselves: the plasticity of embodiment, sensing, and mind. *Journal of Philosophy and Medicine*, 32(3), 362–382.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 10–23.
- Clowes, R. W. (2013). The cognitive integration of e-memory. *Review of Philosophy and Psychology*, 4(1), 107–133.
- Giere, R. (2002a). Models as part of distributed cognitive systems. In L. Magnani, N. Nersessian, & P. Thagard (Eds.), *Model-based reasoning in scientific discovery*. New York: Kluwer.
- Giere, R. (2002b). Scientific cognition as distributed cognition. In P. Carruthers, M. Stich, & P. Siegal (Eds.), *The cognitive basis of science* (pp. 285–299). Cambridge: Cambridge University Press.
- Hacking, I. (1983). *Representing and intervening*. Cambridge: Cambridge University Press.
- Harré, R. (2003). The materiality of instruments in a metaphysics for experiments. In H. Radder (Ed.), *The philosophy of scientific experimentation* (pp. 19–38). Pittsburgh: University of Pittsburgh Press.
- Heersmink, R. (2013). A taxonomy of cognitive artifacts: function, information, and categories. *Review of Philosophy and Psychology*, 4(3), 465–481.
- Heersmink, R. (2014). Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-014-9355-1.

- Holmes, F. L., Renn, J., & Rheinberger, H. J. (Eds.). (2003). *Reworking the bench: Research notebooks in the history of science*. Dordrecht: Springer.
- Humphrey, P. (2009). The philosophical novelty of computer simulation methods. *Synthese*, 169(3), 615–626.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT Press.
- Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, 73(1–2), 1–52.
- Kirsh, D. (2009). Problem-solving and situated cognition. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 264–306). Cambridge: Cambridge University Press.
- Kiverstein, J., & Farina, M. (2011). Embraining culture: leaky minds and spongy brains. *Teorema*, 30(2), 35–54.
- Knuuttila, T. (2005). Models, representation, and mediation. *Philosophy of Science*, 72(5), 1260–1271.
- Knuuttila, T. (2011). Modelling and representing: an artefactual approach to model-based representation. *Studies in History and Philosophy of Science*, 42(2), 262–271.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. New York: Palgrave MacMillan.
- Menary, R. (2010). Dimensions of mind. *Phenomenology and the Cognitive Sciences*, 9(4), 561–578.
- Menary, R., & Kirchhoff, M. (2014). Cognitive transformations and extended expertise. *Educational Philosophy and Theory*, 46(6), 610–623.
- Nersessian, N. J. (2005). Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In M. Gorman, R. Tweney, D. Gooding, & A. Kincannon (Eds.), *New directions in scientific and technical thinking* (pp. 17–56). Mahway: Erlbaum.
- Nersessian, N. J. (2008). *Creating scientific concepts*. Cambridge: MIT Press.
- Nersessian, N. (2009). How do engineering scientists think? Model-based simulation in biomedical engineering research laboratories. *Topics in Cognitive Science*, 1, 730–757.
- Peirce, C. S. (1935). *The collected papers of Charles S. Peirce* (Vol. 3). Cambridge: Harvard University Press.
- Pence, H. E., Williams, A. J., & Belford, R. E. (2015). New tools and challenges for chemical education: Mobile learning, augmented reality, and distributed cognition in the dawn of the social and semantic web. In J. Garcia-Martinez & E. Serrano-Torregrosa (Eds.), *Chemistry education: Best practices, opportunities and trends* (pp. 693–734). Weinheim: Wiley.
- Perini, L. (2005). Explanation in two dimensions: diagrams and biological explanation. *Biology and Philosophy*, 20(2–3), 257–269.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175(1), 133–151.
- Record, I. (2010). Scientific instruments: knowledge, practice, and culture. *Spontaneous Generations: A Journal for the History and Philosophy of Science*, 4(1), 1–7.
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1), 1–19.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Sheredos, B., Burnston, D., Abrahamsen, A., & Bechtel, W. (2014). Why do biologists use so many diagrams? *Philosophy of Science*, 80(5), 931–944.
- Sterelny, K. (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465–481.
- Suarez, M. (2003). Scientific representation: against similarity and isomorphism. *International Studies in the Philosophy of Science*, 17(3), 225–244.
- Sutton, J. (2006). Distributed cognition: domains and dimensions. *Pragmatics and Cognition*, 14(2), 235–247.
- Sutton, J. (2010). Exograms and interdisciplinarity: history, the extended mind and the civilizing process. In R. Menary (Ed.), *The extended mind* (pp. 189–225). Cambridge: MIT Press.
- Sutton, J., Harris, C. B., Keil, P., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9(4), 521–560.
- Toon, A. (2012). *Models as make-believe: Imagination, fiction and scientific representation*. New York: Palgrave MacMillan.
- Toon, A. (2014). Friends at last? Distributed cognition and the cognitive/social divide. *Philosophical Psychology*, 27(1), 112–125.
- Tweney, R., & Ayala, C. D. (2015). Memory and the construction of scientific meaning: Michael Faraday's use of diaries and records. *Memory Studies*. doi:10.1177/1750698015587149.
- Vorms, M. (2012). Formats of representation in scientific theorising. In P. Humphreys & C. Imbert (Eds.), *Models, simulations, and representations* (pp. 250–273). New York: Routledge.
- Wilson, R., & Clark, A. (2009). How to situate cognition: Letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 55–77). Cambridge: Cambridge University Press.
- Winsberg, E. (2010). *Science in the age of computer simulation*. Chicago: The University of Chicago Press.
- Yeo, R. (2008). Notebooks as memory aids: precepts and practices in early modern England. *Memory Studies*, 1(1), 115–136.