

Against cognitive artifacts: extended cognition and the problem of defining ‘artifact’

Andres Pablo Vaccari¹

© Springer Science+Business Media Dordrecht 2016

Abstract In this paper I examine the notion of ‘artifact’ and related notions (namely, ‘cognitive artifact’) in the dominant version of extended cognition theory (ECT) grounded on extended functionalism. Although the term is ubiquitous in the literature, it is far from clear what ECT means by it. How are artifacts conceptualized in ECT? Is ‘artifact’ a meaningful and useful category for ECT? If the answer to the previous question is negative, should we worry? Is it important for ECT to have a coherent theory of artifacts? And what are the demands and constraints that ECT imposes on this theory? I distinguish between two aspects of ECT, one narrow, aligned with extended functionalism (EF); and one broad or pluralistic, in which EF is combined with other theoretical resources in the context of diverse research programs. I begin by determining the problems in conceptualizing artifacts from EF. Then I address the question of why a concept of artifact may be relevant to ECT. Next, I examine the efforts of Richard Heersmink to combine ECT with dominant theories of artifacts in the philosophy of technology (ECT in its broad aspect). I argue that both approaches fail to yield a meaningful notion of artifact, let alone one of ‘cognitive’ artifact. Finally, I argue that narrow ECT places rather strong constraints on a theory of artifacts, since it locates the specificity of ‘artifact’ in material aspects of realization that are, by definition, outside its theoretical purview. I examine, then discard, the possibility that a materialist and objectivist theory of artifacts may be of help. And finally I briefly explore some ways in which a broad, pluralistic ECT may address some of these shortcomings.

Keywords Extended cognition · Cognitive artifacts · Extended functionalism · Metaphysics of artifacts

✉ Andres Pablo Vaccari
avaccari@unm.edu.ar

¹ CITECDE/UNRN, Consejo Nacional de Investigaciones Cientificas y Tecnicas, S. C. de Bariloche 8400 Río Negro, Argentina

1 Introduction

Although artifacts feature prominently in the literature on extended cognition theory (ECT), the category of ‘artifact’ itself has failed to attract attention, despite some moves in this direction (Heersmink 2013, 2014, 2015, to be discussed shortly). The notion of artifact in ECT raises a host of questions that I will address in what follows: (1) How are artifacts conceptualized in ECT? (2) Is ‘artifact’ a meaningful and useful category for ECT? (3) If the answer to the previous question is negative, should we worry? Is it important for ECT to have a coherent view on artifacts? (4) What are the demands and constraints that such a theory should address?

In this paper I argue that one particularly dominant, functionalist strand of ECT (FECT, or functionalist ECT) fails to articulate a clear notion of ‘artifact’ relevant to its purposes, and that this failure can offer some lessons for ECT as a whole. I argue that this blind spot is part of a structural feature concerning how FECT conceives of ‘resource’ and related categories. In this regard, the unintelligibility of ‘artifact’ is related to well-known problems with functionalism (in this case, extended functionalism [EF]) concerning the metaphysics of realization. FECT opens an avenue to conceptualize artifacts that places the specificity of this concept in material aspects of realization—in other words, it demands an *ontology* of artifacts. This imposes rather strict constraints on artifact theory. Further, FECT’s theoretical blindness to artifacts may be a necessary outcome of its conceptual architecture.

Why focus on FECT? Firstly because it is a widespread reading of ECT, and influential theorists like Andy Clark argue it may be the *only* reading. Secondly, the failure of FECT to think artifacts can be instructive to ECT as a whole, for reasons I will develop later. And it may also have some lessons for other approaches to technology and material culture in the human sciences. However, my focus is strictly philosophical and limited to two aspects of theories of cognition: narrow ECT, based on extended functionalism (EF) and “mandated by the existence of functionally specified cognitive systems whose boundaries are located partly outside the skin” (Wheeler 2010a, 245); and broader proposals in which ECT is admitted into pluralist research programs that incorporate other theoretical resources (such as Menary’s cognitive integrationism (2007) and Malafouris’ material engagement theory (2008, 2013); both of which will be briefly addressed in the conclusion). It is possible that this latter strategy may remedy the insufficiencies of FECT by importing extrinsic theoretical resources.

I examine one such proposal, the only one thus far to explicitly examine the question of artifacts. Heersmink (2013, 2014, 2015) has attempted to integrate FECT with dominant approaches to artifacts in the philosophy of technology. I argue that his attempt is an interesting failure that further undermines the promise of functionalism as an approach to artifacts.

In my conclusion, I briefly consider some issues around integrating ECT and artifact theory. I cast the net wider, briefly considering some alternative approaches, such as material agency and Actor-Network. As a first measure, I will leave the notion of ‘artifact’ deliberately vague, departing from the same intuitive folk conception that, I argue, both FECT and ECT as a whole rely on. ‘Artifact’ is an awfully broad term that encompasses tools, machines, systems and the whole gamut of human-made objects. I think this ‘naïve’ first approach is preferable to specifying a certain conception of artifact and then see how ECT measures up to its standards. In other words, we will

begin by diving into FECT and see how the question looks from there. Then, we will follow Heersmink in borrowing such an extrinsic conception of artifact from the philosophy of technology (that is, the branch of philosophy that makes it its business to think, precisely, about artifacts).

The argument is structured as follows:

- 1) I begin by locating the question of artifacts within FECT (question 1). FECT considers an artifact as a kind of ‘resource’ or ‘scaffolding’ (the difference between these two terms is far from clear). However, FECT gives us no criteria to distinguish between types of resource, and so the answer to question 2 is negative.
- 2) Despite the above, I argue that some EC theorists recognize some need for a science of scaffoldings (question 3), to which the question of artifacts may be relevant (i.e., as a type of scaffolding). But I conclude that ‘artifact’ may not be a meaningful or useful category for this science. Further, I locate some lurking need for some clear notion of ‘artifact’ in second- and third-wave ECT. These latest waves of ECT place emphasis on physical structure and the dynamics of realization, forcing us to search for some ontological specificity of artifacts at this level of analysis.
- 3) I examine the efforts of Richard Heersmink to combine ECT with available approaches to artifacts in the philosophy of technology. How are artifacts conceptualized in other areas of philosophy? Can these understandings be successfully brought in line with ECT? My answer is negative. I argue that Heersmink’s approach ends up in (interesting) failure and examine the reasons for this.
- 4) Lastly, question 4: I argue that ECT (and not just FECT) sets up rather strong demands on a theory of artifacts, since, as I have already suggested, it asks us to find ontological specificity in physical aspects of realization (or ‘materiality’). Such a theory should seek a materialist and objectivist ontology of artifacts. Even if such a metaphysic is possible (and I leave this question open), it is not clear how it would help us.

My conclusion draws some lessons from FECT’s failure for future ECT-based theory of artifacts. I also offer some observations about artifacts that are relevant to the human sciences in general, and the study of material culture in particular. These observations concern certain irreducible tensions between metaphysical intuitions about artifacts as objective entities, and their roles in human activities.

2 ECT and artifacts: the problem

The core claim of FECT is the thesis of extended functionalism (EF) (Clark 2008; Wheeler 2010a). As a version of functionalism, EF argues that what matters is the functional role that a component plays in the system of which it is part, regardless of how such a role is carried out or realized. ECT argues that the vehicles that realize cognition can be distributed across brains, external objects, and environments. Thus, things outside the skull-and-skin boundary can be legitimately assigned cognitive status, provided they make a *functional contribution* to a cognitive process (and this is FECT properly speaking).

It is important to emphasize that, for many influential theorists, ECT *is* FECT, inasmuch as EF provides its only secure foundation. As Clark says, FECT is interested in “the information-bearing states of non-biological props and aids” (2008, 40). According to this view, artifacts are a kind of prop or aid that can be functionally incorporated into cognitive processes. As such, artifacts in FECT are subsumed under the notion of ‘resource’, ‘vehicle’ or ‘scaffolding’. Anything can become a cognitive vehicle provided it satisfies certain functional criteria specified by an appropriate level of description. Vehicles may be opportunely realized by physical features of objects, but presumably not just *any* such object is an artifact. How are we to distinguish artifacts from other resources such as, say, niche engineering or the exploitation of found environmental cues?

Extended functionalism can be further unpacked in terms of a metaphysic that conceives of realization in terms of supervenience. Here, the causal contribution of material organization is filtered through levels that are consequent on the existence of a level below. Following Heersmink (2014), physical structure is the metaphysical underpinning of cognitive function, as he shows in the following diagram:

physical structure → representational structure → information → cognitive function
(9);

On his part, Menary (2007) establishes a distinction between *vehicle* and *medium* (70). The first is informational, the second physical. The term ‘medium’ denotes the *instrumental* nature of physical realization: the role (function) of material structure is exclusively to carry out a function, determined by a system output (cognitive task). A written sentence is the vehicle, distinct from the medium (pen and paper, computer screen) on which it is written. At the right level of description, both reading from a notebook and “a pattern of activation” in the brain have the same functional status (72). In this manner, all functional contributions are placed at the same ontological level, as information. Thus, the notion of vehicle leads to some metaphysical riddles, such as trying to pinpoint exactly what is playing a functional role: is it the traces of ink or the patterns triggered on the retina by the reflection of light on these traces?

Put this way, the question of artifacts runs into more general and well-known problems with functionalism. As Shapiro (2008) argues, “functional concepts must ignore some of the properties by which realizers differ and focus instead on those in virtue of which different realizers can realize the same functional property” (10). This leads to what Shapiro calls the ‘boundary problem’: there is no principled way to distinguish a realizer of a function from a causal influence on the realizer. In the case of written sentences, it is not clear whether what carries out the function are the actual written signs (or the patterns on the retina) or whether the written signs are a causal trigger for the patterns of activation in the brain—which must be considered, then, as the actual realizer. In turn, the boundary problem is related to the issue of the conditions under which realizers come to carry out their functions. Building on the work of Shoemaker (1981) and Wilson (2004), Shapiro points out that “a realizer of a property is, by itself, rarely if ever sufficient for that property” (12). To stick to the example of writing as mnemonic aid: pen-and-paper or digital word-processing technologies do not suffice for memory storage. Rather, these technologies realize mnemonic functions in the context of an embedding system in which certain conditions must hold in order for

these functions to be realized, such as an agent (say, Otto, to follow the now classic example) wanting to go to the MoMA, and Otto believing that the traces on the page have been written by him or by someone reliable. Altering the causal connections that realize this function will alter the property that the physical structure realizes (Shapiro, 12).

If we think of an artifact as a specific kind of realizer, then it is clear that EF cannot account for this specificity, since physical structure is irrelevant to it from the very start. Functions can be said to be multiply realizable only at an analytical level coarse-grained enough to paper over any physical differences between media, and to obtain a function abstract enough to be predicated of different material configurations. As an artifact, my house matters inasmuch as physical features of it act as substrata for holding relevant informational properties in place in the context of a particular task, such as finding the kitchen. But so far FECT appears to have nothing, or at best nothing interesting, to say about the nature of my house *qua* ‘artifact’. As a result, all mention of artifacts in FECT (and ECT in general) appeal to a tacit, uncritical folk notion of ‘artifact’ that its theoretical machinery is designed to neither support nor legitimate, and for which it frankly appears to have no use for. Take, for example, the notion of ‘cognitive naturefacts’: natural objects used for some cognitive purpose (Heersmink 2013, 468; the notion of *naturefact* is due to Hilpinen 2011). Following Hutchins (1995), we can illustrate naturefacts with the case of using the stars in the night sky to plot a course through the sea. The stars are obviously not artifacts, but they enter the cognitive system that includes navigators and their tools. However, it is not clear what kind of work is the natural/artificial distinction performing in this context or how it is useful. The distinction is extrinsic and certainly arbitrary for FECT. The much more ubiquitous notion of ‘cognitive artifact’ (to which we turn in the next section) is arguably more problematic.

It follows that FECT cannot account for even basic metaphysical facts about artifacts, such as their persistence as objects. This is because artifacts do not enter cognitive systems as *wholes*, but only insofar as material aspects of them are selectively engaged during cognitive processing. In this way, according to Barrett (2013),

“... functionalism lends itself to a kind of truncated, ‘mentalized’ view of our engagement with the world: material things matter to our everyday experience only as selectively engaged by functional properties of mind. ... [And this] constitutes an arbitrary barrier to engagement, establishing a ‘ceiling of relevance’ that extends across the body-world interface” (12).

Why is this problem significant? Why not simply shrug off artifacts and ban all mention of them from ECT? Because there are strong indications that, at some point, the question of artifacts becomes relevant to ECT-based approaches. This is particularly evident in the more programmatic statements of EC theorists. For a start, the category of artifact may be clarified by, and be useful to, a science of scaffoldings.¹ Clark writes:

“The single most important task, it seems to me, is to better understand the range and variety of types of cognitive scaffolding, and the different ways in which

¹ I take the notion of scaffolding to be roughly interchangeable with that of vehicle or resource, although it does seem to place more emphasis on materiality.

non-biological scaffoldings can augment or impair performance on a task.... The Holy Grail here is a taxonomy of different types of external prop, and a systematic understanding of how they help (and hinder) human performance” (Clark, quoted in Heersmink 2013, 466).

However, it is not clear whether these taxonomical principles may yield, or have any use for, something like a concept of artifact. Causally relevant informational features are multiply realizable and bear no necessary relation to any specific form of material instantiation. There is no reason, then, why specific functional effects should be metaphysically and necessarily related to specific types of scaffolding. The realization problem again: functional properties are intimately related to physical conditions, but EF does not offer any principled way to carry out function-to-structure and structure-to-function inferences. As Adams and Aizawa (2007) point out, there is no clear ordering principle to the range of possible causal engagements with external resources:

“Tools do not constitute a natural kind; tools are, after all, artifacts. It is for this reason that, a would-be brain-tool science would have to cover more than just a multiplicity of causal processes. It would have to cover a genuine motley. A brain-tool science would not have to cover a mere disjunction of things; it would have to cover an open disjunction” (76).

A second context that makes reference to features of artifacts is second-wave ECT (and, by extension, third-wave EC). This theoretical phase of ECT stresses complementarity, integration and manipulation (Kirchhoff 2012, 290). It argues that external resources do not function in an isomorphic, analogous way to internal resources. Thinkers of this persuasion routinely make reference to the specifics of realization, including the way in which realizers physically operate in the context of causal couplings. Menary expands:

“The manipulation of external vehicles is importantly different from manipulations of internal vehicles and their integration is the unit of cognitive analysis. ... Explaining this integration of the internal and the external involves both a dynamical account of the reciprocal causal interaction between internal and external vehicles and processes and an account of how we learn to manipulate external vehicles in accordance with relevant cognitive norms. Hence, the extended mind hypothesis contributes to the wider project of integrating the internal and the external because it focuses on the dynamical account of causal coupling” (2010, 75).

This account of causal couplings requires that we pay close attention to the specifics of external vehicles. For example, John Sutton (2006) builds on a distinction first proposed by Donald (1991) concerning differences between *engrams* (biologically stored representations) and *exograms* (external public symbols). According to the complementarity thesis, Sutton thinks that relevant aspects of external resources include “the medium-dependence or translatability of the information they carry, and its transmissibility across instantiations, their capacity as symbol systems, the constraints on the ways in which information can be retrieved from them, the context-dependence

of their use and so on” (Sutton 242). These features point to system capacities of artifacts grounded on specific material organizations. Again, ‘artifact’ does not appear to be a necessary unit of analysis: what matters is the *integration* of internal and external resources in terms of causal couplings. However, second-wave ECT suggests that we look for the specificity of artifacts in the material aspects of realization: how different types of resource may carry out functions in distinct ways, in tandem with internal and embodied resources, thus allowing for certain cognitive and pragmatic possibilities but not others. For example, in the case of composing a text using pen and paper or a word processor, Menary writes:

“Stable and enduring external written sentences allow for manipulations, transformations, reorderings, comparisons, and deletions of text that are not available to neural processes. This is the upshot of second-wave cognitive-integration-style arguments: bodily manipulations of external vehicles are different from, but complementary to, internal processes” (2010, 240).

The digital medium has a specific materiality that makes an ineliminable contribution to the task of writing. A word-processor configures text in terms of discrete, movable and editable units—something a notebook or a brain cannot do. This insight can be cashed out in terms of the equally helpful and obfuscating distinction between *implementational* and *vital* materiality due to Wheeler (2010b).

Implementational materiality refers to those objects (or features of objects) that enable “enhanced performance in the storage and manipulation of information” (32). The materiality of such objects counts only insofar as it explains how “the computational functions in question are *implemented in the physical world*” (32). This is what EF is interested in, the “more abstract, and in some sense non-material (that is, disembodied) information processing story” (32). Vital materiality, on the other hand, refers to noninformational features of objects that have some causal role in the performance of a task, yet this causal efficacy is not, strictly speaking, ‘cognitive’. Let’s take the example of a map. The informational features in this case are carried out by visual cues (signs, shapes, colors, etc.). These features could be implemented on a wide range of media. However, in its ‘vital’ aspects, the medium (a portable, foldable, durable, light piece of paper that allows for easy storage) makes a contribution to the task. There is nothing in the implementational story that demands that the medium of the map display these features. We could easily have a map on a block of wood or a sheet of iron. We could even have various paper maps showing, respectively, national highways, rivers, topological features, hotels, etc. Of course, it makes much more practical and pragmatic sense to show all of these features on the same portable surface of paper—and this is why maps usually display information in this way. But EF is at a loss to explain why a map has precisely the material features it has in order to properly carry out the cognitive function it does—not to mention the historical, social and cultural context that enables these functions and give them some sense. In this sense, functionalist theories have poor predictive power, offering no basis to make reliable structure-function/function-structure inferences. As a result, the vital/implementational distinction has some limitations, since it is not always easy to distinguish what aspect of objects or the physical world are playing the informational role and which aspects are ‘vital’. Barrett (2013) locates the problem at the level of materiality:

“The functionality of any given material is always changing, depending on the interests, history, and skill of the agent who is interacting with it. For instance, the color of a piece of wood might be irrelevant in one context, and yet fully ‘implementational’ in another. Especially if we abandon the artificial constraints of specific problem-solving activities and consider the limitless variety of human interactions with materials, it is easy to see that ‘implementational materiality’ is a rather fluid and fuzzy category” (12).

As we can see, Barrett’s objections are analogous to Adams and Aizawa’s. One possibility is to look for some extrinsic resources (i.e., a theory of artifacts, and of function-structure relations in this context) and see whether they can be meaningfully integrated into ECT. Heersmink attempts precisely this, and we examine his argument next.

3 ECT and the philosophy of technology

In a series of papers (2013, 2014, 2015), Heersmink has been the first to attempt a rigorous definition of cognitive artifact by bringing together ECT and theories of artifacts in the philosophy of technology. The arguments of Heersmink unwittingly expose the insufficiencies of ECT in this department, beginning with the notion of artifact itself.

Heersmink begins his analysis by delimiting cognitive artifacts as the functional kind of artifacts that “perform cognitive tasks” (2). This type of artifact is distinct from ‘pragmatic’ artifacts. Heersmink provides the example of a map and a chair: the proper function of a map is to “provide task-relevant information used for navigating” while a chair “is not used to perform cognitive tasks” (2). To establish this point, Heersmink draws from some dominant approaches to artifacts in the philosophy of technology, grounded on the notions of *intentions* and *functions* (see Franssen et al. 2014 for an excellent overview). Whereas some authors emphasize the intentions of the designer of the artifact (Thomasson 2007; Kitcher 1993), others argue that the user is the central determinant of an artifact’s function (Neander 1991). In both cases, a constitutive link is established between intention and function: the latter is determined by the former. That is, functions are mind-dependent and anchored on the intentions of an agent.

This take on artifacts admits a range of metaphysical commitments, ranging from anti-realism (e. g. Wiggins 1980) to realism about functional kinds (e.g. Baker 2004, 2007). The strongest metaphysical bet on the market is intentionalism, a form of realism. According to intentionalism, artifacts conform a distinct ontological kind because of their intended character (Dipert 1995; Hilpinen 2011; Baker 2004; Thomasson 2007). As Baker puts it: “Artifacts are objects intentionally made to serve a given purpose” (2004, 99). According to a vigorous intentionalist such as Thomasson, for example, an artifact “must be the product of a controlled process of making”, and this involves “imposing a number of intended features on the object ...” (2007, 58–9). The intention to make a thing of an artifactual kind K “must thus involve a substantive (and substantively correct) concept of what a K is, including an understanding of what sorts of properties are K-relevant and an intention to realize many of them in the object created” (59). The difference between an artifact and a physical or natural object is that

the former was designed with a purpose, which acts as a mediating term between intention and the artifact, or between mind and matter. The maker or user performs a kind of baptismal act in which intended function is impressed on a lump of matter, bestowing on it a distinct metaphysical nature. Note that this functional baptism has two distinct roles: it marks artifacts as a separate ontological category and also *individuates* different artifact kinds (i. e., gives unity to the artifact, or the functional kind, as one single thing).

Now, if we try to harmonize this approach to artifacts with ECT, we clearly see some frictions and incompatibilities. If we think of artifacts in terms of proper function, it is evident that the proper function of an artifact may not necessarily coincide with the function it performs within a given cognitive system. Let's take Heersmink's example of bartenders (lifted in turn from Beach 1993). Bartenders use certain features of objects (such as size, shape, and spatial location and distribution) to reduce the cognitive load of their work and help structure a sequence of actions. By placing glasses of certain shapes in strategic spots, the bartender can recall the drinks to be served and the order in which to serve them. Beach refers to these as "material mnemonic symbols" (Beach 1993, 193). Effectively, bartenders, glasses, tools (such as shakers, spoons, straws, etc.), and the spatial engineering of these, constitute a cognitive system. However, from a functional perspective, the proper function of a glass is to hold a body of liquid; while, for an EC approach, the function of a glass is to act as vehicle for certain informational properties (visual, in this case: size, shape and location) that are then integrated into a task. Each description is coherent in itself but incommensurable with the other.

It is easy to multiply examples that show that proper functions of artifacts, even when they are cognitive, do not coincide with the function realized as a result of integration into a cognitive system. Heersmink switches between these two perspectives without properly integrating them. On one hand, agents "opportunistically use artefactual objects and structures for cognitive purposes" (6). In other words, certain chosen capacities of artifacts are transformed into functional features. These system functions, activated within certain contexts, can be exploited in an either idiosyncratic or ongoing manner. But, on the other hand, cognitive artifacts are "neither defined by intrinsic properties of the artifact nor by the intentions of the designer, but by their function, which is established by the intentions of the user and how it is used" (6). Thus, according to Heersmink, the intention and manner of use, as established in a particular cognitive system, takes precedence over proper function: "it does not matter whether cognitive functions have been selected over time or improvised on the spot. What matters is that the object in question can perform a cognitive function and aid on-board capacities" (6).

This renders untenable the distinction between cognitive and pragmatic artifacts based on proper function. In fact, artifacts routinely communicate their functions (as well as opening up "action schemes" [Illies and Meijers 2009] or "scripts" [Akrich 1992]) by virtue of their formal features, thereby assuming cognitive roles. Take the example of a chef arranging her kitchen to maximize the efficiency of tasks such as accessing implements; say, by having knives laid out in ascending order of size (an example of what Heersmink calls ecological cognitive artifacts, further subdivided into spatial and structural, (2013)). In this instance, the artifacts' appearance are playing an informational role, while the artifact 'itself' carries the obvious pragmatic role. When approaching a table, a chair plays the functional role of directing my body to a certain

place in the context of a task (e.g., ‘sitting down to dinner’). In this case the chair communicates its proper function, yet the informational function certainly does not coincide with the proper-pragmatic function of the chair itself. It is this former informational-communicational function that is part of a cognitive system, and so the chair in this case should be rightly considered a cognitive artifact.

It follows that, for ECT, all artifacts are by definition cognitive artifacts, so that the concept is somewhat circular. If “exhibiting task-relevant information is both a necessary and sufficient condition for being a cognitive artifact” then anything becomes a cognitive artifact the moment it is engaged in the context of a cognitive system (Heersmink 2014, 9) and any of its physical features is exploited for informational purposes. Therefore, if we follow this route, we are back at the beginning. There are no grounds to define an ‘artifact’, as the category ends up swallowing pretty much anything. We could easily call these cognitive functions ‘resources’ without any conceptual loss.

We can conclude that any attempt to illuminate the notion of ‘artifact’ by importing functionalist-intentionalist approaches from the philosophy of technology is bound to run into the difficulties described. A more natural approach would be to taxonomize dimensions of integration (Heersmink 2015; Sutton 2006) or “principles of coordination” (Kirsch 2006). Although these dimensions or principles must account for material aspects of realization, they focus on the human-artifact boundary, rather than on the artifacts themselves. But, again, it does not seem that we need a clear concept of ‘artifact’ to specify these dimensions.

4 Conclusion

ECT seems opens up a rather arduous route to defining an artifact, one that somehow passes through physical and metaphysical aspects of realization. We can establish a rough analogy with the problem of body-centrism (e.g., Noë 2004): the argument that the fine details of embodiment (neurobiological, physicochemical, etc.) make ineliminable, constitutive, non-trivial and/or special contributions to functional implementation, so that multiple realizability does not hold. The issue is discussed in Clark (2008), who characterizes this thesis (admitting a range of flavors) as “total implementation sensitivity” (50). We can illustrate this analogy with the same example of writing we have been using. Van Waes and Schellens (2003) have conducted a series of experiments that compared the cognitive behaviors of experienced writers when using, respectively, pen and paper, and a word processor. The results strongly suggest that cognitive behaviors are deeply affected by writing modes (that is, the artifacts with which the task is performed). In the words of the researchers, “a number of differences in the organization of the writing process have their origin in the characteristics of the computer itself” (838). Writers using a word processor tended to revise more at the level of the letter, lacked initial planning, spent more time on the first draft than on the final text, followed a more fragmentary sequence of revisions, started writing sentences without knowing how they will end, and started writing sooner than pen-and-paper writers. On their part, among other features, pen-and-paper writers tended to revise the text systematically from beginning to end, paused much less, and these pauses took place at sentence or paragraph boundaries. All of these differences can be traced to specific features of the technologies involved. For example, a word-processor allows

for deletions of words before they fully appear on the screen, and deletions of parts of words; while, when writing on pen-and-paper, the whole word has to be crossed out in order to maintain the legibility of the text. These differences are not just superficial, but have a deep impact on the very nature of the texts produced. The point is easier to see in the case of text messaging in mobile phones, which alter the very nature and form of written communication.

The case of writing modes suggest that material features of technology make non-trivial causal contributions to the nature of cognitive tasks, and are ineliminable in terms of their outputs. This and other examples suggest the focus should be the details of implementation and realization. However, it is not clear in which manner artifacts may realize functions differently from any other class of object; we could just focus on these specificities without any need for a concept of ‘artifact’.

One possible way to tackle this problem is to develop a materialist ontology of artifacts that grounds metaphysical specificity on objective properties of artifacts, such as internal structure, operational principles, and physical organization and dynamics. This route might help us (1) ontologically individuate ‘artifact’ as a specific, objective kind of entity, and (2) establish some internal principle of classification that may allow us to sort out different types of artifacts. In the words of Carrara and Vermaas (2009), the artifact has a “principle of activity” defined as “a chain of internal and/or external causal interactions, a process determining the persistence of the object” (2009, 129). Simondon has been the most notorious philosopher to develop such an objectivist metaphysic of artifacts (1958), but we can also mention Baker’s constitution view (2007). However, it is not certain that this route will tell us much about what distinguishes artifacts from other types of resource. The question of what artifacts *are*, and what it is about their *nature* that makes possible their integration into human cognition and action, is too broad and seems motivated by different philosophical concerns. Baker (2006) herself, in her critiques to the Dual Nature of Artifacts Program, rightly observes:

“... it is unclear to me how a general ontology of artifacts generally could help us understand specific patterns of reasoning from function to structure. Surely, the content of such reasoning is not covered by any general theory, but depends on the specific function in question” (134).

It could be possible that ECT-based artifact theory should be eschewed, or more productively combined with other approaches to artifacts in the human sciences. For instance, in material engagement theory (MET), artifacts are not informational entities but matter inasmuch as they afford possibilities of action. In this view, artifacts become equal partners in a “dance of agency” that is partly constitutive of intentionality (Malafouris 2008). The intentionalist view (according to which an artifact is a neatly defined entity with a stable, ontological core) has been challenged on a number of fronts within disciplines such as archaeology, anthropology, and science and technology studies (e.g.: Hayles 1993; Pickering 1995; Preda 1999; Latour 1999; Miller 2005; Turkle 2007; Latour 2000; Dant 2005, Malafouris 2013; Law 2008; Vannini 2009; Boivin 2010; Pickering 2010). Although these studies are inscribed in different research programs, with their own concerns and approaches, they all converge on the notion that material culture shapes human agency, rather than the other way around. Factors hitherto considered external to human action (such as the physical and informational

features of objects, materials and environments) are considered to be constitutive of human interactions with the world—a feature also shared by ECT. These situated factors have ineliminable causal roles in the physical and epistemic outcomes of action (Kirchhoff 2009); they also involve bodily and phenomenological aspects of “skillful coping” that are irreducible to the language of intentionality and representation (Dreyfus 2007). According to post-phenomenological approaches, for example, what matters is the way technological mediation shapes concrete existence and interpretations of the world (e.g., Rosenberger and Verbeek 2015).

There appears to be a conflict between certain intuitions about artifacts (that they are a class of physical object with causal independence, ontologically static, with an uncertain relation to human intentionality, etc.) and approaches that consider the active contribution of artifacts to a dynamic system. One is static, the other *relational*. If we can draw some lessons from these approaches, it is that artifacts are not ontologically independent lumps of organized matter. What matters in Actor-Network theory, MET and other approaches inspired in Dynamic Systems theory, is the way that components of a system make a contribution to that system, shaping its development and output. These approaches share with FECT the central notion that action and cognition are deeply relational phenomena, and that the essence of things is given by their effects rather than by metaphysical identity. In this sense, what matters is an empirical description of an artifact’s precise role and effects in a given phenomenon; this is perhaps why, as Baker points out above, strictly philosophical definitions of artifacts do not seem to translate well in the field, and do not provide a useful starting point for object ethnographies or object biographies.

We might also figure out a third way, one that combines the notion of an artifact as site of a process, a pattern of activity, and the basic objectivist thesis that an artifact is an arrangement of matter with its own principles of organization.² For example, in the case of the bartenders, a relational approach must account for the fact that properties (size, shape, color, portability and other features of glasses) are *necessarily* interrelated in virtue of a shared material structure that supports supervening phenomena, or that contributes to a dynamics of emergence.

It remains to be seen how the informational paradigm of EF (and, by association, ECT) can be combined with such an objectivist ontology, or whether this last would be of any use; but at least we have made some headway in identifying a problem that deserves attention. While we seek some clarification to the question of artifacts in ECT, it seems that we will have to remain content with papering over these holes with the same old vague folk notions.

Acknowledgments The author would like to acknowledge the invaluable comments of Jesús Vega Encabo and Diego Lawler in shaping the direction of the argument presented here, as well as the many informal discussions that helped develop the main problems examined in this paper. This paper was possible thanks to a grant from the Ministry of Economy and Competitiveness of Spain, for the project “Material Cultures, Epistemic Cultures: Standards, cognitive practices and knowledge,” at the Autonomous University of Madrid, Spain.

² This nagging objectivist feeling is well captured by certain currents within critical theory, such as Bill Brown’s Thing Theory (2004). According to a recent set of studies, Simondon’s philosophy of technology may successfully account for the dual nature of artifacts (thing and process, object and subject of human intentions) (Blanco et al. 2015).

References

- Adams, F., & Aizawa, K. (2007). Defending the bounds of cognition. In R. Menary (Ed.), *The extended mind* (pp. 67–80). Cambridge: MIT Press.
- Akrich, M. (1992). The de-scription of technical objects. In W. E. Bijker & J. Law (Eds.), *Shaping technology/ Building society* (pp. 205–224). Cambridge: MIT Press.
- Baker, L. R. (2004). The ontology of artifacts. *Philosophical Explorations*, 7, 99–111.
- Baker, L. R. (2006). On the twofold nature of artifacts. *Studies in History and Philosophy of Science*, 37, 132–136.
- Baker, L. R. (2007). *The metaphysics of everyday life: An essay in practical realism*. Cambridge: Cambridge University Press.
- Barrett, N. (2013). Cognitive science, embodiment, and materiality. Retrieved 12 January 2016, from https://www.academia.edu/5371657/Cognitive_Science_Embodiment_and_Materiality.
- Beach, K. (1993). Becoming a bartender: the role of external memory cues in a work-directed educational activity. *Applied Cognitive Psychology*, 7, 191–204.
- Blanco, J., Parente, D., Rodríguez, P., & Vaccari, A. (Eds.). (2015). *Amar a las máquinas: Cultura y técnica en Gilbert Simondon*. Buenos Aires: Prometeo Libros.
- Boivin, N. (2010). *Material cultures, material minds: The impact of things on human thought, society, and evolution*. New York: Cambridge University Press.
- Brown, B. (2004). *Things*. Chicago: Chicago University Press.
- Carrara, M., & Vermaas, P. (2009). The fine-grained metaphysics of artifactual and biological functional kinds. *Synthese*, 169, 125–143.
- Clark, A. (2008). Pressing the flesh: a tension in the study of embodied, embedded mind? *Philosophy and Phenomenological Research*, 76(1), 37–59.
- Dant, T. (2005). *Materiality and society*. Maidenhead: Open University Press.
- Dipert, R. (1995). Some issues in the theory of artifacts: defining ‘artifact’ and related notions. *The Monist*, 78(2), 119–136.
- Donald, M. (1991). *Origins of the modern mind*. Cambridge: Harvard University Press.
- Dreyfus, H. (2007). The return of the myth of the mental. *Inquiry*, 50(4), 352–365.
- Franssen, M., Kroes, P., Reydon, T. A. C., & Vermaas, P. E. (Eds.). (2014). *Artifact kinds: Ontology and the human-made world*. Dordrecht: Springer.
- Hayles, N. K. (1993). The materiality of informatics. *Configurations*, 11, 147–170.
- Heersmink, R. (2013). A taxonomy of cognitive artifacts: function, information, and categories. *Review of Philosophy and Psychology*, 4(3), 465–481.
- Heersmink, R. (2014). The metaphysics of cognitive artifacts. *Philosophical Explorations*, 1–16. doi 10.1080/13869795.2014.910310.
- Heersmink, R. (2015). Dimensions of integration in embedded and cognitive systems. *Phenomenology and the Cognitive Sciences*, 13(3), 577–598.
- Hilpinen, R. (2011). Artifact. *Stanford Encyclopedia of Philosophy*. Retrieved 1 June 2015, from <http://plato.stanford.edu/entries/artifact/>.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT Press.
- Illies, C., & Meijers, A. (2009). Artefacts without agency. *The Monist*, 92(3), 420–440.
- Kirchhoff, M. D. (2009). Material agency: a theoretical framework for ascribing agency to material culture. *Techné: Research in Philosophy and Technology*, 13(3), 1–11.
- Kirchhoff, M. D. (2012). Extended cognition and fixed properties: steps to a third-wave version of extended cognition. *Phenomenology and the Cognitive Sciences*, 11(2), 287–308.
- Kirsch, D. (2006). Distributed cognition: a methodological note. *Pragmatics and Cognition*, 14(2), 249–262.
- Kitcher, P. S. (1993). Function and design. *Midwest Studies in Philosophy*, 18(1), 379–397.
- Latour, B. (1999). *Pandora's hope: Essays on the reality of science studies*. Cambridge: Harvard University Press.
- Latour, B. (2000). When things strike back: A possible contribution of ‘Science Studies’ to the social sciences. *British Journal of Sociology*, 51, 107–123.
- Law, J. (2008). Actor network theory and material semiotics. In B. S. Turner (Ed.), *The new Blackwell companion to social theory* (pp. 141–158). London: Wiley-Blackwell.
- Malafouris, L. (2008). At the potter's wheel: An argument for material agency. In C. Knappett & L. Malafouris (Eds.), *Material agency: Towards a non-anthropocentric approach* (pp. 19–36). NY: Springer.
- Malafouris, L. (2013). *How things shape the mind: A theory of material engagement*. Cambridge: MIT Press.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basigstoke: Palgrave Macmillan.

- Miller, D. (Ed.). (2005). *Materiality*. Durham: Duke University Press.
- Neander, K. (1991). Functions as selected effects: the conceptual analyst's defense. *Philosophy of Science*, 58, 168–184.
- Noë, A. (2004). *Action in perception*. Cambridge: MIT Press.
- Pickering, A. (1995). *The mangle of practice: Time, agency and Science*. Chicago: University of Chicago Press.
- Pickering, A. (2010). Material culture and the dance of agency. In D. Hicks & M. Beaudry (Eds.), *The oxford handbook of material culture studies* (pp. 191–208). Oxford: Oxford University Press.
- Preda, A. (1999). The turn to things: arguments for a sociological theory of things. *The Sociological Quarterly*, 40(2), 347–366.
- Rosenberger, R., & Verbeek, P-P. (eds.). (2015). *Postphenomenological Investigations: Essays on Human-Technology Relations*. Lexington Books.
- Shapiro, L. A. (2008). Functionalism and mental boundaries. *Cognitive Systems Research*, 9, 5–14.
- Shoemaker, S. (1981). Some varieties of functionalism. *Philosophical Topics*, 12, 93–119.
- Simondon, G. (1958). *Du mode d'existence des objets techniques*. Paris: Aubier.
- Sutton, J. (2006). Distributed cognition: domains and dimensions. *Pragmatics & Cognition*, 14(2), 235–247.
- Thomasson, A. (2007). Artifacts and human concepts. In E. Margolis & S. Laurence (Eds.), *Creations of the mind: theories of artifacts and their representation* (pp. 52–71). Oxford: Oxford University Press.
- Turkle, S. (Ed.). (2007). *Evocative objects: Things we think with*. Cambridge: MIT Press.
- Van Waes, L., & Schellens, P. J. (2003). Writing profiles: the effect of the writing mode on pausing and revision patters of experienced writers. *Journal of Pragmatics*, 35, 829–853.
- Vannini, P. (Ed.). (2009). *Material culture and technology in everyday life: Ethnographic approaches*. NY: Meter Lang Publishing, Inc.
- Wheeler, M. (2010a). In defense of extended functionalism. In R. Menary (Ed.), *The extended mind* (pp. 245–270). Cambridge: MIT Press.
- Wheeler, M. (2010b). Minds, things, materiality. In L. Malafouris & C. Renfrew (Eds.), *The cognitive life of things: Recasting the boundaries of the mind* (McDonald Institute monographs, pp. 29–37). Cambridge: McDonald Institute for Archaeological Research Publications.
- Wiggins, D. (1980). *Sameness and substance*. Oxford: Blackwell.
- Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences, cognition*. NY: Cambridge University Press.