Marcin Miłkowski
ORCID: 0000-0001-7646-5742

Institute of Philosophy and Sociology, Polish Academy of Sciences

Cognitive Artifacts and Their Virtues in Scientific Practice

This paper proposes a novel way to understand various kinds of scientific representations in terms of cognitive artifacts. It introduces a novel functional taxonomy of cognitive artifacts prevalent in scientific practice, which covers a huge diversity of their formats, vehicles, and functions. It is argued that toolboxes, conceptual frameworks, theories, models, and individual hypotheses can be understood as supporting our cognitive performance in scientific practice. While all these entities are external representations, their function can be best understood through the conceptual lens of wide cognition. The functional approach suggests that the assessment of knowledge representation in science should be based on functions that cognitive artifacts help us perform. By providing a conceptual link between the functionality of artifacts and their virtues, this approach also recommends an empirical approach to the study of virtues. This implies that the cognitive approach to the study of science can offer some guidance in recent philosophical debates around the nature of scientific theories or models.

1. Cognitive artifacts in scientific practice: preliminaries

One of the critical issues in the philosophy of science is to understand scientific knowledge, and in particular, various kinds of representations of this knowledge. In this paper, I offer a unified cognitive perspective on scientific representations by framing them in terms of cognitive artifacts.

The notion of cognitive artifact was introduced in cognitive science to refer to entities that “maintain, display, or operate upon information in order to serve a representational function and that affect[s] human cognitive performance” (Norman, 1991, p. 11). The study of cognitive artifacts has begun under the auspices of distributed cognition (Hutchins, 1995; Norman, 1993), which takes a specific approach to cognitive processes. Distributed cognition does not constrain the research to individual properties of cognitive agents, but focuses its attention on larger cognitive systems, which comprise multiple agents and artifacts, such as the ones present in a research laboratory (Nersessian, 2016). Nonetheless, philosophers investigating cognitive artifacts analyzed them predominantly in the context of the extended mind (Clark & Chalmers, 1998; Heersmink, 2013; Vaccari, 2017). The extended mind conception states that the mind may extend into what is traditionally conceived as the mind’s environment, incorporating external resources such as tools, language, or other external systems (Clark & Chalmers, 1998). While distributed cognition and the extended mind seem close in their rejection of the individual as the sole unit of analysis (Rupert, 2013), they should not be confused (Wachowski, 2018). According to the extended mind, the individual mind remains at the center of a cognitive system; this individual-centered perspective is not endorsed by proponents of distribution. In line with the distributed approach, this paper is not targeted at the notion of mind at all, which implies that it is not committed to core claims of the extended mind. Instead, it tackles scientific practice as a distributed cognitive practice, viewing cognition as essentially “wide”, i.e., including factors that go beyond intracranial processes (Miłkowski, Clowes, et al., 2018).

By following Norman’s account of cognitive artifacts, I also assume that they are essentially representational. Their representational nature was questioned by Heersmink (2013). His argument, in essence, is that there are non-representational cognitive artifacts, dubbed “ecological”, such as car keys left on someone’s desk as a reminder to return a rented DVD. However, Fasoli (2018) undermined Heersmink’s argument, simply by pointing out that car keys are used as signs in this example. They indicate that the DVD should be returned. There is no reason to suppose that there are cognitive artifacts without representational function.

The general assumption of this paper is that scientific knowledge representations are cognitive artifacts. While the claim that scientific representations serve a representational function and affect human cognitive performance might seem fairly bland, it will be argued that the cognitive approach may lead to novel insights. These insights include normative assessment of features of cognitive artifacts. As will be shown, these insights can be derived from scientific communication, by relying on natural language processing.

The rest of the paper proceeds in the following way. In the next section, I argue that the functioning of cognitive artifacts in scientific practice should be understood in terms of computations over external representations. Section 3 provides a preliminary taxonomy of scientific representations that includes several dimensions: their format, their contribution to an aim of inquiry, their intended scope, and the degree to which they depend on data. Subsequently, it is argued that one can take a functional view on scientific representation. By empirically studying how artifacts function in scientific practice, we may discover virtues (and vices) of scientific representations. This is illustrated with the virtue of computational efficiency. In conclusion, future directions of work on cognitive artifacts in scientific practice are sketched.

2. External representations in distributed scientific practice

Cognitive artifacts in scientific practice are understudied. For cognitive scientists, mundane representational formats, such as written notes or databases, seem less worthy of inquiry than representations that were not studied extensively by philosophers of science, such as gestures. Philosophers of science, in turn, sometimes seem so perplexed by the immense variety of what scientists may mean by “theory” or “model”, that they become skeptical of whether these even exist (French, 2020). But there is a way to curb this apparently amorphous matter.

While “theory” or “model” may stand for particular instances of scientific representations, researchers often speak of them in a way that is relatively independent of their physical bearers. For example, when speaking of Newton’s theory of mechanics, physicists and philosophers rarely think of the first edition of his work that relied mostly on geometric diagrams. Instead, theories or models are typically identified in terms of their contents, and they may span multiple scientific publications, sometimes over a significant length of time (Hochstein, 2016). By following this usage, I assume that particular theories, models etc. are contents of cognitive artifacts, whereas by contents I understand sets of satisfaction conditions of these artifacts.

Scientific practice is a paradigmatic case of a distributed cognitive process. While deserving credit and asserting their primacy of discovery motivate individual researchers in their work, the work is only rarely, or ever, attributable fully to a single individual. Given the fact that cognitive labor is distributed among lab members, or among multiple collaborating teams, there is a crucial role to play for external representations. This is because they are easily shareable and provide the linchpin for the intersubjective validity in science. They also stabilize cognitive processes, allowing one to reliably reproduce the same result over and over again. This contribution to individual cognition is underpinned by the fact that cognitive operations on external representation are less reliant on our scarce resources of working memory.

An example can be found already in the beginnings of the systematic study of geometry in ancient Greece (Hohol, 2020; Netz, 2011). The intersubjective validity of proofs relied on the combination of two specific cognitive artifacts: labeled geometric diagrams and formulaic language of proofs (Hohol & Miłkowski, 2019). Anyone familiar with such diagrams and language can reproduce the result of the cognitive process. This intersubjective reproducibility is one of the core features of reliable scientific knowledge. The role of external representations in cognitive processes is also to guide multiple agents in similar ways, which may sometimes obviate the need to account for individual differences between particular agents (Afeltowicz & Wachowski, 2015). Psychological differences among researchers do obviously exist, but they may be negligible factors in understanding the progress of reliable scientific processes such as the ones involved in Euclidean geometry.

But all this is not meant to suggest that having the right external representations is sufficient for someone to know Euclidean geometry. Actually, cognitive processes involved in geometric cognition intertwine internal and external representations. Such intertwining is a feature of distributed practice over shared external representations (Nersessian, 2016; Zhang, 1997; Zhang & Norman, 1994).[1]

Representations cannot do any cognitive work by themselves. They must be operated upon, or computed over. Otherwise, they cannot affect our cognitive processing. Nonetheless, we are still in the dark about many kinds of computational processing that underpin various types of cognitive practices. This is the case even for Euclidean geometry: we do not know how people perform computations over labeled diagrams by relying on the formulaic descriptions and what computations are even involved (Hohol & Miłkowski, 2019). The same applies to other cognitive artifacts: while we have mathematical models of verbal or logical reasoning (usually derived from formal logic), our understanding of cognitive processing involved in such reasoning is still largely lacking. In many cases, all we have are somewhat sketchy and biologically implausible neural network models (Stenning & Lambalgen, 2008). And we lack even sketchy models for such prevalent representational artifacts as communicative gestures, which also play a crucial role in scientific theorizing (Becvar, Hollan, & Hutchins, 2008).

What makes external cognitive artifacts easier to study than mental processes is that they are often (but not always) observable with the naked eye. Observability is the case for diagrams on paper, written notes on a blackboard, and gestures (Marghetis & Núñez, 2013). The procedures their users perform using these artifacts are also observable, and sometimes easily learnable this way. For example, it’s much easier to learn Euclidean geometry by observing how diagrams are drawn than by merely reading the static text.

[1] Arguably, Hutchins, in his (1995, p. 172), was skeptical of the notion of cognitive artifact because it seemed to imply that there is no role for internal processing (and applying the notion of “artifact” to such processing seems clumsy at best), and that one could ignore other cognitively significant external entities, such as stars in traditional maritime navigation, which are not artifacts either. Later work on cognitive artifacts, including the work of Hutchins, avoids such simplistic and implausible suggestions.

However, with the increasing complexity of representational systems, the transparency of processing may decrease. It is not as easy to understand logical operations simply by observing an avid reader of *Principia Mathematica*, and even harder to understand the operations of neurocomputational models by observing a human neurocognitive scientist operating a computational model. These may become obscure. In spite of being shareable, they may no longer be even reproducible (Miłkowski, Hensel, & Hohol, 2018).

There are some preliminary attempts to distinguish cognitive artifacts (as external physical entities) from procedures, which include mental procedures (Heersmink, 2013, p. 468). However, without detailed and empirically validated computational models of how we use cognitive artifacts, such a distinction remains somewhat arbitrary. This is because the same result can be achieved using various algorithms operating on various kinds of data structures, and the very distinction between a data structure and an algorithm (or computational operation) does not actually exist in some kinds of computational models, such as Turing machines. This implies that our common-sense distinction may be equally inapplicable to the computational mechanism of our cognitive processing.

3. A preliminary taxonomy of scientific cognitive artifacts

In this section, a preliminary taxonomy of scientific cognitive artifacts will be presented. In contrast to previous work (e.g., Brey, 2005; Fasoli, 2018; Heersmink, 2013, 2021), this taxonomy does not aim to cover all kinds of cognitive artifacts, but is geared towards scientific practices. Moreover, the taxonomy does not map cognitive artifacts onto individual human cognitive capacities (such as memory or perception) or onto ways cognitive artifacts may interact with individual cognition (by substituting, complementing or constituting it). There are two reasons for this theoretical choice. First, our current understanding of human cognitive capacities seems to be insufficiently empirically grounded to guide our analysis of artifacts. As argued by Poldrack et al. (2011; see also Pessoa, Medina, & Desfilis, 2021), our main categories of cognitive function are still in their infancy and likely to be revised in the future. Second, the distributed cognition perspective does not make a single individual a natural unit of analysis, so it might be difficult to identify the relation of scientific artifacts to individual cognitive processes. Take, for example, a satnav device that uses GPS data to plan a route. Cognitive processes, such as spatial orientation, of end-users might be substituted by this device (Fasoli, 2018). But in scientific practice, devices are frequently tweaked, so the same computer system might be customized for special needs of, say, archeological research by augmenting the maps. Researchers are often bricoleurs and reuse existing artifacts for new purposes or redesign them.

Instead, there will be four dimensions of artifacts in our taxonomy: (1) their representational format, (2) their role vis-à-vis the general aims of inquiry, (3) their intended scope, and (4) the degree to which they are accounted for in terms of data. The taxonomy is not a complete classification, but only a first step towards this goal. This is because of the virtually infinite variety in representational practices in science.

Let me start with a representational format. Charles Peirce famously defended a taxonomy of all signs into indexes, icons, and symbols (Short, 2007). In contemporary debates, these original terms are not always used, but the spirit is still there: indexes are understood as indicators or receptors, which mean X because they correlate with X; icons are structural representations that mean X because they resemble X; and symbols are just conventional representations, linguaform or similar. The problem with the distinction is that many consider receptors as extreme cases of structural representation (Facchin, 2021;

Morgan, 2013; Nirshberg & Shapiro, 2020), which blurs the boundaries between these. Moreover, there are some distributed representations that do not seem to be fully iconic, because their component parts need not carry any representational function (Haugeland, 1998). And there are complex hybrid formats, such as labeled diagrams in geometry or maps with arbitrary symbols and linguistic descriptions, that do not fit into just one category of signs. Thus, one should not assume that there cannot be hybrid formats that add conventional elements to natural signs.

The vehicles of these signs vary considerably. Some of them rely on human perceptual modalities, especially visual, auditory, and tactile. Gestures of other researchers can be seen. Pops of an electrode can be heard when a neuron fires. And the feel of a special control button may be distinctive. Beyond these vehicles, basically any kind of intermediary relatively stable physical process can bear information, but will usually require translation to the perceptually available format for human researchers. (Sometimes no such translation may be required, e.g., if research is fully automated, cf. (Fodor, 1991)).

One may be surprised to note that there is such diversity of representational formats and vehicles. But depending on the representational format, some but not all problems can be actually solved (Larkin & Simon, 1987). Even if contentwise representations may be equivalent, they are not cognitively equivalent, and finding an appropriate format is the most important step toward solving a difficult problem.

To summarize, the format of representation may rely on only natural sign relation (resemblance or correlation), be fully conventional, or hybrid. This gives us a tripartite distinction, albeit slightly different from the Peircean one:

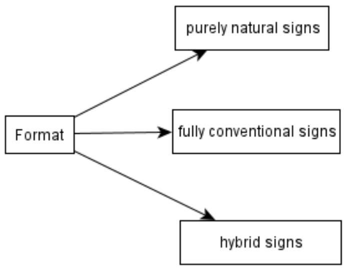

The second dimension of the taxonomy is the relationship to the general aims of inquiry, understood, following Laudan (1977), in terms of problem-solving. Traditionally, it is assumed that these aims are to describe, predict, explain, and control phenomena. But this list is a mere typology of aims. In addition, there are some more complex aims. For example, a team may wish to reproduce a certain experiment to find out whether its results were fabricated. Obviously, this would imply that some experiment must be performed, based on previous description and some predictions made, to be compared against the experimental results. Nonetheless, the specific aim of reproducing the previous research is lost in this analysis. This is because the list of aims does not include any metascientific goals, such as comparing experimental results, performing meta-analyses, or closely reproducing or conceptually replicating experiments. It also does not include relationships between the aims or inquiry-supporting aims: one of the elementary aims of many research projects is to test hypotheses, but to do that, predictions inferred from hypotheses must be compared against descriptions of phenomena (usually as measured empirically). In this case, two major aims are linked together. Moreover, testing may also rely on a specific experimental protocol that is followed closely to make results intersubjectively reproducible. Again, relying on experimental protocols is not included in any of the major aims, but it is crucial for intersubjective reliability. In lieu of a complete taxonomy of possible goals of inquiry, I will simply assume here that there is a fifth goal, which can be conceptualized as supporting reliable research (including all kinds of inferential practices).

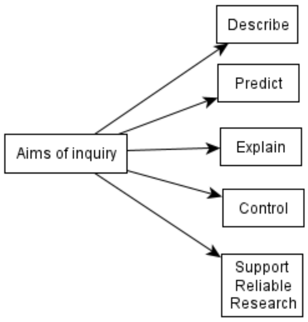

In her influential work on the research laboratory, Nersessian (2016) distinguished three kinds of cognitive artifacts: devices, instruments, and equipment. Basically, these distinctions can be easily mapped onto this dimension of the current: devices are models that serve predictive and explanatory aims; instruments are measuring tools, which allow researchers to describe phenomena; and equipment is in the service of control or auxiliary experimental labor (fifth goal).

The third dimension of the taxonomy is the intended scope, which will be understood as the number of observable phenomena that are represented in a given cognitive artifact. The scale should start from 0, since purely mathematical or logical research may target no observable phenomenon at all.[2] Scientific representations need not be declarative; there could be, for example, certain techniques or methods enshrined in experimental protocols, which are not, strictly speaking, true or false. But they have some ontological commitments (for example, about certain experimental materials) and satisfaction conditions that specify correctness criteria for following a certain protocol.

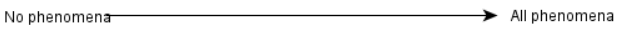

Finally, the fourth dimension specifies the degree to which the contents of the artifact can be accounted for only in terms of data. This is meant to describe the degree to which a given artifact is theoretical in its nature. However, in contrast to the intended scope, there is likely no external representation that can be entirely accounted for in terms of data because “raw data” is an oxymoron (Gitelman, 2013). There are always theoretical choices involved in structuring the artifact, for example, in choosing the representational format, which are not simply “out there” to be discovered. While it may seem difficult to operationalize this dimension, as long as we take “data” to mean “data we have already garnered” (which is sometimes called “the ground truth” in the data science community), this degree can be understood in terms of the epistemological relationship of licensing inferences (deductive or statistical). A given artifact can be less or more reliant on the data we already know, and the degree to which it can generate new outputs to be compared to some future data specifies its degree of theoretical ladenness. Purely mathematical theories are again a special case because they are never compared against any ground truth. In this respect, they are infinitely theoretical, while having zero intended scope.

[2] Alternatively, because mathematical truths hold in all possible worlds, one could claim that their scope is maximal; but they never hold of anything in particular, so it seems more apt to leave the zero as their intended scope.

No reliance → Almost "raw" data

To make this proposal a bit more concrete, let us see how it could be applied to scientific practice in more detail. To the best of my knowledge, there is no systematic study of all kinds of external representations in scientific practice. But some initial insights can be found in communicative practices of researchers. For example, by looking at a large corpus of English academic writing, one could inquire into objects of verbs that are specifically used to speak of defending or proposing certain representations, such as “propose”. Figure 1 shows the result of analyzing the objects of the verb “propose” in a large 2-billion token corpus of English academic writing, produced from open access journals indexed in the Directory of Open Access Journals (DOAJ). This corpus is available in the state-of-the-art corpus analysis software SketchEngine, which allows its users to analyze characteristic formulaic expressions (called “collocations”) and semantic relations. As can be easily seen, multiple kinds of representations are cited (with some noise, note the presence of “antenna”).

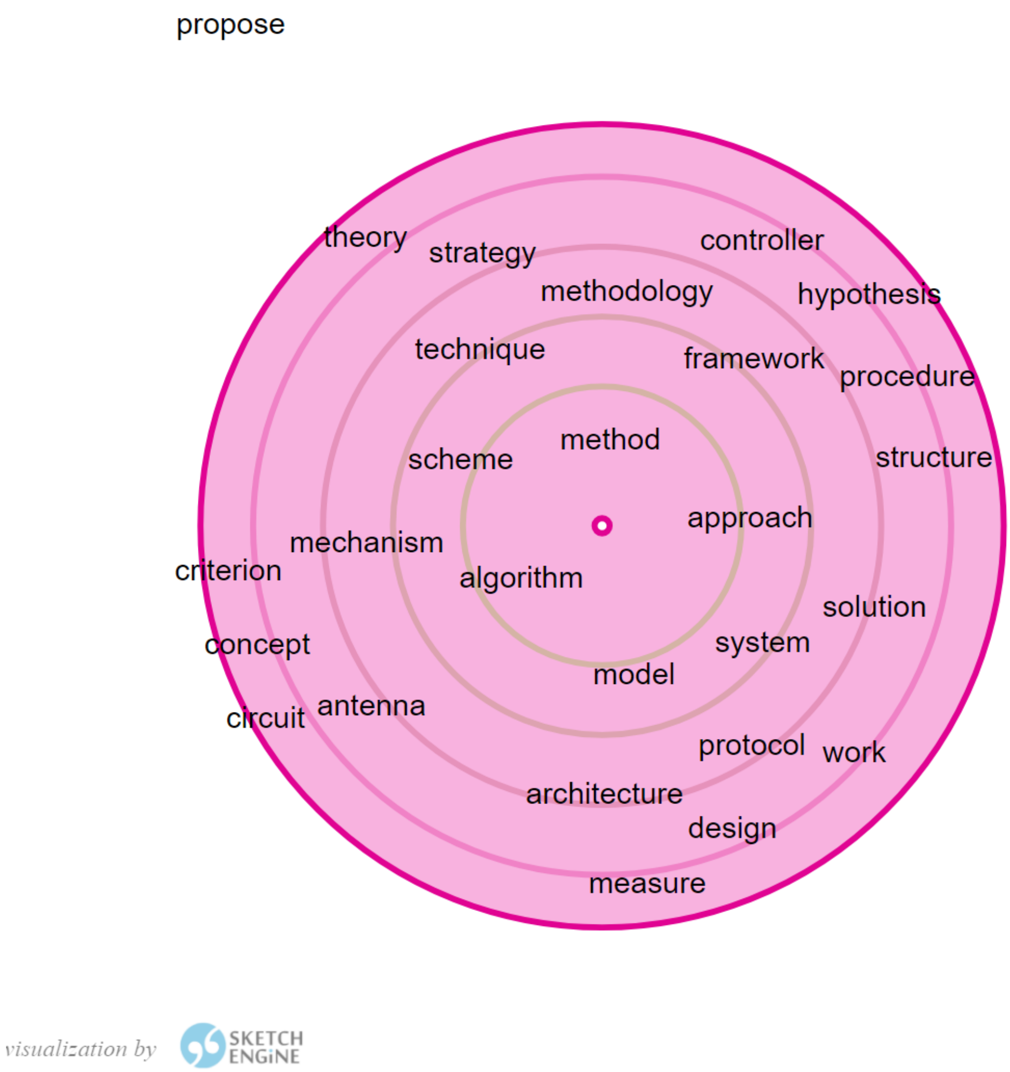

Figure 1: Objects of the verb “propose” as found in the DOAJ Corpus.

On Figure 1, one can see that the most typical kinds of entities proposed by researchers are methods, algorithms, and approaches. The exact location of a given representation in the taxonomy introduced here can be provided only for particular tokens because types are often underspecified in many respects. While algorithms are rarely represented in the format that relies on resemblance (with some exceptions, as one can actually depict a causal structure of a computational mechanism to represent a computational model), and typically specified in symbolic programming languages, it is impossible to determine a priori how any algorithm is related to the aims of the inquiry, intended scope, or its theoretical character.

In previous work, we have argued that the popular predictive processing (PP) account of cognition (Clark, 2016) is not a theory but at best a conceptual framework (Litwin & Miłkowski, 2020). We have distinguished among toolboxes, frameworks, theories, and models. The distinction can be easily understood by pointing out that mathematical toolboxes are not dependent on any ground truth (they can be freely reinterpreted, which is the case for PP work, as we have insisted), while frameworks have both a large scope and contain placeholders that allow researchers to relate mathematical details to data. In contrast, theories are much more detailed (no placeholders anymore) but retain large scope, whereas models, at least in cognitive science, focus on individual cognitive tasks: they are both available for validation against empirical data and limited in their intended scope. Similarly, the taxonomy may be used to analyze individual hypotheses stated in science. For example, the hypothesis that there are at least three pathways in vision processing (Pitcher & Ungerleider, 2021) is stated in natural language (conventional format). Its aim is to explain features of visual processing (and functional brain anatomy), whereas its scope is the whole visual system in the brain. Quite clearly, it is highly data-reliant, meaning that it should be empirically verifiable (which is consistent with the fact that it also has goals of describing, predicting, and explaining facts).

Having briefly discussed how the suggested taxonomy can be used, let us now turn to the relationship between the function of cognitive artifacts and their virtues.

4. From function to virtue, and back

This section introduces a novel method for the study of virtues of cognitive artifacts. The taxonomy provided in the previous section describes cognitive artifacts along four dimensions. One of them, namely the relationship to aims of inquiry, is not only functional but also provides a way to evaluate the performance of a given artifact. For example, descriptive artifacts have the function of describing phenomena. They function correctly if they describe them accurately (at least to a certain degree); otherwise, they are dysfunctional. This functionality can be also described in terms of virtues that have been studied by philosophers of science (Kuhn, 1977). Indeed, accuracy was one of the five virtues analyzed by Kuhn. He also insisted on consistency, scope (unification), simplicity, and fruitfulness (Kuhn, 1977, p. 332).

These virtues can be analyzed and taxonomized conceptually. In fact, most conceptual studies take their lead from Kuhn's proposal (e.g., Keas, 2018). Only recently, pioneering empirical studies of how researchers appeal to virtues have been performed (Mizrahi, 2021). Mizrahi, by performing a plain text search over JSTOR articles in multiple scientific disciplines, found that depending on the kind of scientific representation, different virtues are ascribed: models are treated differently from theories or hypotheses. The taxonomy introduced in the previous section suggests that there could be even more differences in how virtues (and vices) are assigned because various kinds of representations play different roles in scientific practice. Accuracy is obviously not a virtue of a purely formal model (with zero intended scope), for example.

Needless to say, Kuhn's catalog is not particularly systematic. Clearly, it is not a logical classification, which is exhaustive, whose categories are mutually exclusive and each is non-empty. In subsequent work, multiple other virtues were introduced. For example, Keas (2018) insists that one important virtue is applicability, which can be attributed to scientific representations aimed at control. De Regt (2017) analyzes visualizability as the virtue contributing to the intelligibility of scientific theories (at least according to some of their proponents). While these particular virtues extend the received list of virtues, a descriptively adequate account of scientific practice requires a more systematic treatment.

It should now be apparent that the previous section hints at a possible departure point: by systematically looking at various kinds of representations, we could study their functionality. However, this proposal might seem somewhat abstract without any specific suggestions regarding the methods of performing such studies. Here is one such suggestion: scientific communication is arguably at the core of scientific practice – and by focusing on how researchers describe the practice, we could at least get a "native understanding" of their conceptual world. Current language technology is particularly apt for this purpose (Lean, Rivelli, & Pence, 2021; Pence & Ramsey, 2018). Digital philosophy of science can rely on data-driven metascience, by using natural language processing tools.

Before introducing a small case study of the use of language technology to discover virtues, several caveats are in order. It would be naive to assume that scientific publications faithfully report the process of discovery. Scientific communication is constrained by multiple conventional factors, with each writing genre incurring its own biases. For example, typical scientific papers are structured into sections in a way that rhetorically suggests a clear-cut disconnect between results and their theoretical interpretation. Medawar (1963) went so far as to say that this "misrepresents the processes of thought that accompanied or give rise to the work that is described in the paper". Peer-reviewed publications might also avoid harsh language when criticizing previous work or, alternatively, exaggerate their novel insights by excessively criticizing others (Jurgens, Kumar, Hoover, McFarland, & Jurafsky, 2018). Negative results are difficult to publish, which also biases the view of the conceptual world of science in scientific papers, reporting results that, on average, turn out to be false (Ioannidis, 2005). Finally, some aspects of communication cannot be captured in print at all, such as gestures in a lab (Becvar et al., 2008).

While some of these objections might be somewhat exaggerated, they should be treated seriously. Their significance depends, however, on how one relies on language technology in studying artifacts. My assumption is that language technology should be treated heuristically, as a tool for systematic discovery of features inherent in patterns of scientific writing. These features may be then analyzed conceptually and become parts of epistemological argument.

In the short study presented in Section 5, the issue of possible bias inherent in the data is arguably negligible. But there is a way to avoid detrimental biases in cases where they could impact results. When using language technology to analyze scientific discourse, one should not constrain the analysis to a single writing genre or a single field. Various genres have their own conventions and biases. Textbooks and science outreach (such as blogs or podcasts) may oversimplify issues, while referee reports might emphasize nuances. Blog authors tend to write about issues that are not explicitly treated in scientific papers because these are shared knowledge. In short, one should include various genres in the analysis.

In addition, various fields also display diversity in their epistemic standards, which can also fluctuate over time. Thus, a comparison across scientific fields or various research traditions can lead to the discovery that there are some patterns of scientific practice that differ across them. As soon as reproducibility problems were discovered in medical sciences and psychology, these have become analyzed in other fields, such as computational neuroscience. These differences may then be put under closer scrutiny to understand rationales for maintaining certain epistemic standards.

To sum up, by juxtaposing diverse independent data, one may expect the plausibility of results to increase. The same epistemological point has long been appreciated by many philosophers, who coined different terms to talk of the same phenomenon: "convergence of independent evidence", "convergent validation" (Campbell & Fiske, 1959), "triangulation" (Davidson, 1991), "robustness analysis" (Wimsatt, 2007) are all instances thereof. In the machine learning community, a similar argument has been made with recourse to "error diversity": when we have multiple, comparable independent models of the same target, we can hope to offset errors that are almost always present in any model (Brown, Wyatt, Harris, & Yao, 2005).

Let me reiterate that language technology does not obviate human judgement. While currently available textual resources exceed by a wide margin what we had at our disposal only a few decades ago, computational processing remains somewhat brittle, even if exceeding human performance in some very specific tasks in limited domains. The validity of results achieved must be confirmed in one way or another by appealing to human qualitative judgements, which remain the "ground truth" for language models. Thus, this research methodology connects both quantitative and qualitative aspects.

A full systematic treatment of virtues (and vices) of various cognitive artifacts goes beyond the scope of this paper, but it will be instructive to provide one complete example of a virtue that was hitherto neglected by philosophers of science: computational efficiency of models.

5. Computational efficiency: a virtue of computational models

To discover virtues ascribed to scientific representations, one can study formulaic aspects of natural language (Wray, 2012). In particular, one can extract collocations, usually understood as a sequence of words that occur together more often than would be expected by chance alone (for a recent review, see (Gablasova, Brezina, & McEnery, 2017)). Scientific language is highly formulaic, not only because of the technical terminology but also owing to specific syntactical patterns (e.g., marking inferences explicitly, not avoiding word repetition for more clarity etc.). In short, our attention should turn to collocations of terms for referring to various kinds of representations. Take "model": arguably, the adjectives that are used to predicate qualities of models (i.e., predicate adjectives of "model", as in "model is …") can be candidate expressions that stand for virtues of models. To analyze predicate adjectives, one can rely on SketchEngine, which is a state-of-the-art corpus analysis and lexicography system (Kilgarriff et al., 2014). The specific feature of SketchEngine that is useful for this purpose is the availability of "word sketches", or collocations defined over surface grammar patterns. SketchEngine supports multiple languages and hosts several hundred language corpora, with accompanying word sketch grammars that usually contain a definition of predicate adjective.

For the purpose of this study, I will rely on the largest English corpus of scientific writing, which is based on open access journal papers, indexed in DOAJ, which was already mentioned in the previous section. Its primary advantage is the size, while the sample of papers is driven mostly by the easy availability of contents for further processing. This can lead to questions about whether it is not a biased sample. In fact, it is, but the bias can be innocuous, as my analysis aims to demonstrate. Table 1 provides a complete list of all predicate adjectives of “model”.

Keyword	Gramrel	Collocate	Frequency	Score
model	**adjective predicates of X**		116774	2.370
		able	8243	9.620
		capable	3229	9.000
		suitable	2869	8.990
		good	3017	8.870
		appropriate	1633	8.480
		useful	1907	8.460
		applicable	1502	8.370
		accurate	1303	8.250
		significant	2651	8.150
		valid	1218	8.120
		simple	1160	8.020
		sensitive	1692	8.010
		consistent	2502	7.940
		similar	3674	7.770
		robust	901	7.690
		available	3728	7.590
		such	7226	7.490
		correct	700	7.440
		adequate	687	7.440

	close	1157	7.400
	unable	765	7.360
	superior	673	7.290
	different	1574	7.170
	complex	710	7.110
	equivalent	632	7.070
	difficult	839	6.980
	successful	517	6.950
	reliable	491	6.920
	dependent	922	6.900
	identical	632	6.870
	flexible	443	6.850
	necessary	1013	6.830
	relevant	627	6.830
	easy	514	6.770
	sufficient	620	6.760
	realistic	394	6.730
	small	1082	6.700
	preferred	434	6.700
	due	3180	6.650
	effective	616	6.630
	linear	410	6.630
	large	937	6.610
	comparable	602	6.580

	corresponding	744	6.520
	important	942	6.490
	stable	468	6.480
	efficient	417	6.460
	acceptable	344	6.430
	compatible	357	6.380
	independent	502	6.330
	high	1654	6.310
	essential	522	6.300
	possible	512	6.300
	subject	374	6.280
	reasonable	297	6.270
	low	1144	6.260
	satisfactory	292	6.240
	feasible	311	6.220
	likely	804	6.210
	general	274	6.170
	more	510	6.160
	nonlinear	245	6.030
	specific	394	5.990
	less	539	5.920
	equal	450	5.920
	relative	498	5.900
	predictive	221	5.840

	inadequate	211	5.800
	crucial	252	5.690
	perfect	191	5.690
	insensitive	195	5.670
	unique	221	5.620
	representative	229	5.610
	simulated	178	5.560
	analogous	185	5.550
	great	410	5.510
	m1	166	5.510
	complete	189	5.480
	expensive	180	5.480
	straightforward	171	5.460
	attractive	172	5.450
	insufficient	179	5.450
	ideal	164	5.450
	free	209	5.430
	helpful	169	5.420
	critical	217	5.370
	detailed	153	5.370
	valuable	157	5.360
	preferable	151	5.330
	popular	151	5.320
	fitting	141	5.290

	used	140	5.230
	deterministic	136	5.230
	tractable	135	5.230
	unlikely	164	5.210
	interesting	150	5.210
	unknown	197	5.190
	clear	176	5.160
	bad	140	5.150
	deficient	143	5.140

Table 1: All predicate adjectives of “model” (this is the grammatical relation, or “Gramrel”, under study). The column “Frequency” lists absolute frequency of a given collocate, while “Score” provides a numerical score, computed in terms of logDice, which is an association score of expressions used in SketchEngine.

Fairly low down the list, one can find “efficient”, with over 400 instances of usage (per 3 billion tokens), but relatively high association score.[3] Two facts should be noted: other corpora of academic English available through SketchEngine do not list “efficient” among the collocates of “model”, but they are often many magnitudes of order smaller, and DOAJ is the only gigacorpus, sporting billions of tokens. What makes the DOAJ corpus special, however, is not only its size but relatively recent provenience, as compared to other resources. Thus, it should make us both suspicious of the collocate, and eager to investigate what actual linguistic patterns are responsible for this score. In addition, Table 1 includes many collocates that can be best described as noise for the purposes of this study: “m1” is most likely an unformatted indexed variable m_1, which is not a predicate; “such” is not a term for a homogenous property of any model (in contrast to “general”), but simply a function word that introduces an enumeration or complex phrase (viz., “such as”, “such that”). Some do not carry much information about virtues either (e.g., “high”, “unknown”, or “great”), and these adjectives may stand for various attributes of models.

Let us now consider several examples of phrases that mention efficiency of models in DOAJ. Some uses seem uninformative: “The sexual model was much less efficient at

[3] There are several ways association scores are computed for collocations. Many scores used in collocation research depend on the notion of mutual information, but these scores remain valid only within a single corpus on which they were computed. In contrast, the score used by SketchEngine, logDice, remains relatively independent of a given corpus (Rychlý, 2008).

maintaining species despite the higher rate of species formation" (a paper from *PLoS Computational Biology*), "We emphasize that the model is fairly efficient" (from *Natural Hazards and Earth System Sciences*). But another pattern is easily found in the concordance as well: "Its simplifying assumptions make the model computationally efficient." (*International Journal of Rotating Machinery*), "These models are computationally efficient but contain certain inaccuracies especially in the areas of geometrical discontinuities." (*Lecture Notes in Engineering and Computer Science*), "Both models are computationally efficient and produce results that are correlated with subjective results." (*Radioengineering*). The recurring pattern shows a multi-word expression "computationally efficient", used to ascribe a particular quality to computational models.

Computational models are merely a specific subclass of all cognitive artifacts used in modeling. Let us first situate them in our taxonomy. Their format varies depending on the methodology used. There are symbolic computational models, which are purely conventional. Some models have been argued to feature structural representations (for example, in predictive processing, cf. (Gładziejewski, 2016)), and some might be hybrid. They may also serve various aims, from description, prediction, explanation, to control (in engineering). Some may also play supportive roles, for example, by providing combinations of results from multiple models to offset their individual errors. They can be purely formal, but also aim at an arbitrary number of phenomena. Finally, these models can be entirely speculative or be derived from data to some extent. In other words, there is a rich variety of various kinds of computational models and particular subclasses of computational models occupy almost any location in the taxonomy. Hence, an important question is whether there are any deep epistemological generalizations that can be made about them. Are there any specific virtues or vices related to computational modeling?

This question is all the more significant because their prominence in science and engineering cannot be denied, in particular in the era of deep learning. We experience a steady increase of computational power and growing use of computational modeling in various branches of inquiry. No wonder that older and smaller corpora of academic writing do not mention computational modeling and its features frequently because modeling of this kind was not as widely available in the 1980s or 1990s. They simply reflect the scientific practice of their day, which did not include deep learning methods. As a consequence, they cannot shed light on our question, in contrast to DOAJ, which seems to imply that computational efficiency is a virtue of computational models.

Let me elaborate. A mere description of a linguistic pattern – that "efficient" is a collocate of "computational model" – is insufficient to justify the normative claim that the efficiency of computational models is indeed a virtue. This claim can be, however, justified by showing that scientific problems of computational modeling can be solved better by more efficient models rather than by less efficient ones. The normative nature of the claim lies in the fact that to achieve their goals (of inquiry), scientists perform particular actions and the selection of this action remains rational as long as they indeed contribute to achieving these goals. In other words, this claim is an instrumental norm: it concerns the choice of a particular means to achieve a certain goal, rather than goals themselves. This may be summarized in the following argument schema:

1. X wants to achieve their goal g_1.
2. If X selects the means m_1 rather than m_2, X will achieve their goal g_1 with significantly greater probability rather than otherwise.

Thus, faced with the choice between m_1 and m_2, X should select m_1.

Note that X could find another means, m_3, to achieve their goal g_1, or even give up on achieving g_1. The conclusion of this argument schema relies on the fact that m_1 is sufficient to achieve g_1 (and more reliable than m_2) and that's basically all there is to this *methodological* kind of epistemological normativity (Miłkowski, 2010).

By finding that researchers mention efficiency when speaking of models, one can presuppose that, at least implicitly, they might endorse computational efficiency as the effective means for modeling. The endorsement may also be explicit. For example, this is the case for the state-of-the-art models for automated image classification in the ImageNet database: the current leader of the ImageNet Classification leaderboard (Soo Ko, 2019), EfficientNet-B7, is designed to be more computationally efficient and accurate at the same time (Tan & Le, 2020). By knowing this, we may presuppose that researchers, as rational agents, endorse the claim that efficiency is a virtue of computational models (at least for image classification). As long as being more efficient turns out to be more helpful in modeling rather than the disregard of efficiency, it is actually a virtue as much as simplicity or fruitfulness are also such virtues.

To sum up our inquiry at this point: we have found that there are some virtuous aspects of functionality of computational models. A possible objection that could be raised at this point is whether these aspects are already covered by extant taxonomies of virtues.

One could try to subsume computational efficiency under at least three virtues: simplicity, consistency, and applicability. Simplicity is one of the classical virtues in Kuhn's catalog. The problem, however, with subsuming efficiency under simplicity is that computational efficiency may require giving up on syntactic simplicity (this is known in terms of space-time tradeoff or time-memory efficiency tradeoff, cf. (Hellman, 1980). In other words, there is no guarantee the simplest possible model is the most efficient one; frequently, the relationship is inverse (depending on the details of the algorithm of the model). This is why one could believe that efficiency is a case of another virtue: consistency. After all, we know that more efficient models are better than non-efficient ones. The issue with this concept is that one could rephrase other virtues in the same way (e.g., scope: it's better to have models of larger than smaller scope, etc.). While one could treat another, related computational property of models, tractability, as a requirement with which any computational model should be externally consistent (in particular in cognitive science, see (Frixione, 2001; Van Rooij, 2008)), efficiency is not an absolute requirement, and one could arguably prefer simplicity to efficiency. Finally, efficiency might be considered to be subsumed under applicability, because more efficient models are more applicable in practice. Yet, again, this is a simplification: applicability is understood by Keas as providing guidance for successful action or enhancing technological control. Computational efficiency, however, might increase applicability only for models that are already in the service of successful action or that enhance our technological control. A purely speculative theoretical model without any application whatsoever can be as well computationally efficient.

Hopefully, this short justification suffices to state that "computational efficiency" is a specific virtue of computational models that should have been included among chief virtues. This is because the functionality of computational artifacts depends on their efficiency. Inefficiency, in other words, is usually a vice because inefficient models are dysfunctional.

Before I conclude, let me note briefly that efficiency cannot be ascribed to all kinds of models. For example, it would be a category mistake to ascribe computational efficiency to animal models in life sciences: while these models are sometimes painstakingly engineered to instantiate particular biological features, their own functioning cannot be understood in purely computational terms (of course, setting aside an extremely risky hypothesis that everything instantiates at least one computational mechanism; see (Piccinini & Anderson, 2018)). This is because animal models do not instantiate computations that are essential for modeling target phenomena: we need some additional mechanism to perform computations over representations of animal models to infer properties of phenomena under study.

6. Conclusion

After having introduced a novel taxonomy for scientific representations, I defended the view that understanding actual scientific practice can shed light on what should be considered a virtue or vice. One of the methods that can be used for that purpose is to rely on language technology to extract linguistic patterns describing features of cognitive artifacts involved in scientific inquiry. By systematically relying on the functional relationship between the use of a given artifact and the general goal of inquiry, one can substantiate normative claims about virtues and vices of various artifacts.

In this paper, I relied on collocation extraction for this purpose, but insights into scientific practice need not be restricted to explicit linguistic patterns. An exciting future perspective is to complement empirical hands-on research on scientific practice (e.g., from Science and Technology Studies) with big data processing to understand how certain practices might be reflected by science communication. Arguably, however, the most important tasks for future studies are (1) to distinguish distinct kinds of scientific artifacts depending on their features included in the proposed taxonomy, and (2) to provide a systematic study of their virtues and vices. Moreover, the very process of scientific inquiry and its aims require a much more extensive analysis in terms of problem-solving than could be provided here for reasons of space.

To sum up, the aim of this paper was to defend the view that scientific representations are cognitive artifacts, whose functioning depends on computations we may perform on/using them. Their correct functioning depends on their having certain crucial features, which instantiate epistemic virtues (or, alternatively, vices, if they do not function properly).

Acknowledgements
For helpful comments, I thank Witold Hensel, Mateusz Hohol, Wiktor Lachowski, Szymon Miłkoś, Przemysław Nowakowski, Daniel Piecka, Wiktor Rorot, and Witold Wachowski. The work on this paper was funded from National Science Center Grant under decision DEC-2020/39/O/HS1/02239 (PI: Marcin Miłkowski).

References

Afeltowicz, Ł., & Wachowski, W. (2015). How Far we Can Go Without Looking Under the Skin: The Bounds of Cognitive Science. *Studies in Logic, Grammar and Rhetoric*, *40*(1), 91–109. doi: 10.1515/slgr-2015-0005

Becvar, A., Hollan, J., & Hutchins, E. (2008). Representational Gestures as Cognitive Artifacts for Developing Theories in a Scientific Laboratory. In M. S. Ackerman, C. A. Halverson, T. Erickson, & W. A. Kellogg (Eds.), *Resources, Co-Evolution and Artifacts: Theory in CSCW* (pp. 117–143). London: Springer. doi: 10.1007/978-1-84628-901-9_5

Brey, P. A. E. (2005). The Epistemology and Ontology of Human-Computer Interaction. *Minds and Machines*, *15*(3), 383–398. doi: 10.1007/s11023-005-9003-1

Brown, G., Wyatt, J., Harris, R., & Yao, X. (2005). Diversity creation methods: A survey and categorisation. *Information Fusion*, *6*(1), 5–20. doi: 10.1016/j.inffus.2004.04.004

Campbell, D. T., & Fiske, D. W. (1959). Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychological Bulletin*, *56*(2), 81–105. doi: 10.1037/h0046016

Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. New York: Oxford University Press.

Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, *58*(1), 7–19.

Davidson, D. (1991). Epistemology externalized. *Dialectica*, *45*(2–3), 191–202. doi: 10.1111/j.1746-8361.1991.tb00986.x

Facchin, M. (2021). Structural representations do not meet the job description challenge. *Synthese*. doi: 10.1007/s11229-021-03032-8

Fasoli, M. (2018). Substitutive, Complementary and Constitutive Cognitive Artifacts: Developing an Interaction-Centered Approach. *Review of Philosophy and Psychology*, *9*(3), 671–687. doi: 10.1007/s13164-017-0363-2

Fodor, J. A. (1991). The Dogma that Didn't Bark (A Fragment of a Naturalized Epistemology). *Mind*, *100*(2), 201–220. doi: 10.1093/mind/LI.202.200

French, S. (2020). *There Are No Such Things As Theories*. Oxford: Oxford University Press. doi: 10.1093/oso/9780198848158.001.0001

Frixione, M. (2001). Tractable competence. *Minds and Machines*, 379–397.

Gablasova, D., Brezina, V., & McEnery, T. (2017). Collocations in Corpus-Based Language

Learning Research: Identifying, Comparing, and Interpreting the Evidence. *Language Learning*, *67*(S1), 155–179. doi: 10.1111/lang.12225

Gitelman, L. (Ed.). (2013). *“Raw data” is an oxymoron*. Cambridge, Massachusetts: The MIT Press.

Gładziejewski, P. (2016). Predictive coding and representationalism. *Synthese*, *193*(2), 559–582. doi: 10.1007/s11229-015-0762-9

Haugeland, J. (1998). *Having Thought. Essays in the metaphysics of mind*. Cambridge, Mass. and London: Harvard University Press.

Heersmink, R. (2013). A Taxonomy of Cognitive Artifacts: Function, Information, and Categories. *Review of Philosophy and Psychology*, (4), 465–481. doi: 10.1007/s13164-013-0148-1

Heersmink, R. (2021). Varieties of Artifacts: Embodied, Perceptual, Cognitive, and Affective. *Topics in Cognitive Science*, *13*(4), 573–596. doi: 10.1111/tops.12549

Hellman, M. (1980). A cryptanalytic time-memory trade-off. *IEEE Transactions on Information Theory*, *26*(4), 401–406. doi: 10.1109/TIT.1980.1056220

Hochstein, E. (2016). One mechanism, many models: A distributed theory of mechanistic explanation. *Synthese*, *193*(5), 1387–1407. doi: 10.1007/s11229-015-0844-8

Hohol, M. (2020). *Foundations of Geometric Cognition*. New York: Routledge.

Hohol, M., & Miłkowski, M. (2019). Cognitive Artifacts for Geometric Reasoning. *Foundations of Science*, *24*(4), 657–680. doi: 10.1007/s10699-019-09603-w

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, Mass.: MIT Press.

Ioannidis, J. P. a. (2005). Why most published research findings are false. *PLoS Medicine*, *2*(8), e124–e124. doi: 10.1371/journal.pmed.0020124

Jurgens, D., Kumar, S., Hoover, R., McFarland, D., & Jurafsky, D. (2018). Measuring the Evolution of a Scientific Field through Citation Frames. *Transactions of the Association for Computational Linguistics*, *6*, 391–406. doi: 10.1162/tacl_a_00028

Keas, M. N. (2018). Systematizing the theoretical virtues. *Synthese*, *195*(6), 2761–2793. doi: 10.1007/s11229-017-1355-6

Kilgarriff, A., Baisa, V., Bušta, J., Jakubíček, M., Kovář, V., Michelfeit, J., … Suchomel, V. (2014). The Sketch Engine: Ten years on. *Lexicography*, 7–36. doi: 10.1007/s40607-014-0009-9

Kuhn, T. S. (1977). *The essential tension: Selected studies in scientific tradition and change*. Chicago: The Univ. of Chicago Press.

Larkin, J., & Simon, H. A. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*, *11*(1), 65–100. doi: 10.1016/S0364-0213(87)80026-5

Laudan, L. (1977). *Progress and Its Problem: Towards a Theory of Scientific Growth*. Berkeley, Calif: University of California Press.

Lean, O. M., Rivelli, L., & Pence, C. H. (2021). Digital Literature Analysis for Empirical Philosophy of Science. *The British Journal for the Philosophy of Science*. doi: 10.1086/715049

Litwin, P., & Miłkowski, M. (2020). Unification by Fiat: Arrested Development of Predictive Processing. *Cognitive Science*, *44*(7), e12867. doi: 10.1111/cogs.12867

Marghetis, T., & Núñez, R. (2013). The Motion Behind the Symbols: A Vital Role for Dynamism in the Conceptualization of Limits and Continuity in Expert Mathematics. *Topics in Cognitive Science*, *5*(2), 299–316. doi: 10.1111/tops.12013

Medawar, P. (1963). Is Scientific Paper a Fraud? *The Listener*, *70*, 377–378.

Miłkowski, M. (2010). Making Naturalised Epistemology (Slightly) Normative. In M. Miłkowski & K. Talmont-Kamiński (Eds.), *Beyond Description: Naturalism and Normativity* (pp. 72–84). London: College Publications.

Miłkowski, M., Clowes, R. W., Rucińska, Z., Przegalińska, A., Zawidzki, T., Gies, A., … Hohol, M. (2018). From Wide Cognition to Mechanisms: A Silent Revolution. *Frontiers in Psychology*, *9*, 2393. doi: 10.3389/fpsyg.2018.02393

Miłkowski, M., Hensel, W. M., & Hohol, M. (2018). Replicability or reproducibility? On the replication crisis in computational neuroscience and sharing only relevant detail. *Journal of Computational Neuroscience*, *45*(3), 163–172. doi: 10.1007/s10827-018-0702-z

Mizrahi, M. (2021). Theoretical Virtues in Scientific Practice: An Empirical Study. *The British Journal for the Philosophy of Science*. doi: 10.1086/714790

Morgan, A. (2013). Representations gone mental. *Synthese*, *191*(2), 213–244. doi: 10.1007/s11229-013-0328-7

Nersessian, N. J. (2016). The Cognitive-Cultural Systems of the Research Laboratory. *Organization Studies*. doi: 10.1177/0170840606061842

Netz, R. (2011). *The shaping of deduction in Greek mathematics: A study in cognitive history*. Cambridge: Cambridge Univ. Press.

Nirshberg, G., & Shapiro, L. (2020). Structural and indicator representations: A difference in degree, not kind. *Synthese*. doi: 10.1007/s11229-020-02537-y

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

Norman, D. A. (1993). *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading, Mass.: Addison-Wesley Pub. Co.

Pence, C. H., & Ramsey, G. (2018). How to Do Digital Philosophy of Science. *Philosophy of Science*, *85*(5), 930–941. doi: 10.1086/699697

Pessoa, L., Medina, L., & Desfilis, E. (2021). *Refocusing Neuroscience: Moving Away from Mental Categories and Toward Complex Behaviors* [Preprint]. Open Science Framework. doi: 10.31219/osf.io/8cmhg

Piccinini, G., & Anderson, N. G. (2018). Ontic Pancomputationalism. In M. E. Cuffaro & S. C. Fletcher (Eds.), *Physical Perspectives on Computation, Computational Perspectives on Physics* (pp. 23–38). Cambridge: Cambridge University Press. doi: 10.1017/9781316759745.002

Pitcher, D., & Ungerleider, L. G. (2021). Evidence for a Third Visual Pathway Specialized for Social Perception. *Trends in Cognitive Sciences*, *25*(2), 100–110. doi: 10.1016/j.tics.2020.11.006

Poldrack, R. A., Kittur, A., Kalar, D., Miller, E., Seppa, C., Gil, Y., … Bilder, R. M. (2011). The

Cognitive Atlas: Toward a Knowledge Foundation for Cognitive Neuroscience. *Frontiers in Neuroinformatics*, *5*. doi: 10.3389/fninf.2011.00017

Regt, H. W. de. (2017). *Understanding scientific understanding*. New York: Oxford University Press.

Rupert, R. D. (2013). Distributed cognition and extended-mind theory. In *Encyclopedia of philosophy and the social sciences.* Los Angeles: SAGE Publications.

Rychlý, P. (2008). A Lexicographer-friendly Association Score. *Proceedings of Second Workshop on Recent Advances in Slavonic Natural Languages Processing*, 6–9. Brno: Masaryk University.

Short, T. L. (2007). *Peirce's theory of signs*. Cambridge; New York: Cambridge University Press.

Soo Ko, B. (2019). ImageNet Classification Leaderboard. Retrieved March 13, 2022, from Computer-Vision-Leaderboard website: https://kobiso.github.io/Computer-Vision-Leaderboard/imagenet.html

Stenning, K., & Lambalgen, M. V. (2008). *Human Reasoning and Cognitive Science*. Cambridge, Mass.: MIT Press.

Tan, M., & Le, Q. V. (2020). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *ArXiv:1905.11946 [Cs, Stat]*. Retrieved from http://arxiv.org/abs/1905.11946

Vaccari, A. P. (2017). Against cognitive artifacts: Extended cognition and the problem of defining 'artifact.' *Phenomenology and the Cognitive Sciences*, *16*(5), 879–892. doi: 10.1007/s11097-016-9484-9

Van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, *32*(6), 939–984. doi: 10.1080/03640210801897856

Wachowski, W. M. (2018). Commentary: Distributed Cognition and Distributed Morality: Agency, Artifacts and Systems. *Frontiers in Psychology*, *9*. doi: 10/gdcbs5

Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Cambridge, Mass.: Harvard University Press.

Wray, A. (2012). What Do We (Think We) Know About Formulaic Language? An Evaluation of the Current State of Play. *Annual Review of Applied Linguistics*, *32*, 231–254. doi: 10.1017/S026719051200013X

Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, *21*(2), 179–217. doi: 10.1016/S0364-0213(99)80022-6

Zhang, J., & Norman, D. A. (1994). Representations in Distributed Cognitive Tasks. *Cognitive Science*, *18*(1), 87–122. doi: 10.1207/s15516709cog1801_3