TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science 3 (2011) 222–226

ISSN: 1756-8757 print / 1756-8765 online
DOI: 10.1111/j.1756-8765.2011.01133.x

Semantics and Metaphysics in Informatics: Toward an Ontology of Tasks

Carrie Figdor

Department of Philosophy, University of Iowa

Received 28 May 2010; received in revised form 12 August 2010; accepted 22 September 2010

Abstract

This article clarifies three principles that should guide the development of any cognitive ontology. First, that an adequate cognitive ontology depends essentially on an adequate task ontology; second, that the goal of developing a cognitive ontology is independent of the goal of finding neural implementations of the processes referred to in the ontology; and third, that cognitive ontologies are neutral regarding the metaphysical relationship between cognitive and neural processes.

Keywords: Cognitive ontology; Cognitive control; Cognitive phenomics; Cognitive tasks; Informatics; Response selection; Cognitive neuroscience; Semantics

In a recent paper, Lenartowicz, Kalar, Congdon, and Poldrack (2010) seek to demonstrate the potential of data-mining in the development of a cognitive ontology—a searchable vocabulary of clearly defined cognitive terms and their relations—of use in cognitive neuroscience. Their study is prompted by a fundamental background question: Given cognitive neuroscience's goal of mapping cognitive functions to brain structures, what are the cognitive functions that should be mapped? This question would be boring if imaging experiments found neural activation in particular kinds of brain structures (at levels ranging from single neurons to networks) to be strongly correlated with and only with specific kinds of functions designated by terms in our everyday mental vocabulary (like ''memory'') or their close cousins (like ''word form recognition''). Alas, many imaging studies have found that different processes can engage the same neural structures (pluripotentiality) and different structures can subserve the same processes (degeneracy) (Figdor, 2010). Pluripotentiality is especially irksome relative to the one-to-one ideal, since it makes reverse inferences—from observed activation in brain areas to the involvement of a specific cognitive function—extremely weak as evidence for cognitive hypotheses (Poldrack, 2006, 2008).

Correspondence should be sent to Carrie Figdor, Assistant Professor, Department of Philosophy, 260 English-Philosophy Building, University of Iowa, Iowa City, IA 52242. E-mail: carrie-figdor@uiowa.edu

But pluripotentiality may reflect either a too coarse-grained taxonomy of neural structures (in which we do not have enough neural structures to match the number of cognitive processes) or a too fine-grained taxonomy of cognitive processes (in which we have too many cognitive processes for the number of neural structures). This latter possibility has a more sophisticated variant expressed as the idea that everyday mental processes ''emerge'' from combinations of more basic cognitive processes that are realized one-to-one in brain structures. This idea lies behind studies, such as Lenartowicz et al.'s, in which neuroimaging results are taken as the fixed points to which a cognitive taxonomy must conform (Price & Friston, 2005). But it is worth noting that a cognitive ontology would be tremendously useful independently of how mapping issues are resolved. That is why Lenartowicz et al.'s approach can contribute importantly toward achieving the former even though their metaphysical conclusion—that a cognitive function exists only in the minds of scientists when the use of a label for it cannot now be predicted on the basis of imaging studies—is unjustified.

Lenartowicz et al. performed a meta-analysis that used the results of text-mining techniques measuring frequency of occurrence of search terms co-occurring with the term ''cognitive control'' in the PubMed database (Sabb et al., 2008) to analyze reported patterns of brain activity associated with these terms found in studies in the BrainMap Database (Laird, Lancaster, & Fox, 2005; http://www.brainmap.org). Imaging studies involving ''response selection,'' ''response inhibition,'' ''working memory'' and ''task/set switching''—labels for putative components strongly associated with the label ''cognitive control''—were analyzed to see if any specific neural structures could be systematically linked to uses of particular component labels and so whether one could predict the use of a component label based on patterns of brain activity. This is a linguistic variant of reverse inference, in which one reasons backwards from a report of activation in a structure to the cognitive label that was probably used to describe that activity, rather than (as is typical) from the observation of activity to the kind of cognitive capacity that structure is probably performing (or helping to).

Their results showed that ''response selection'' stood out as being sufficiently neurally distinguishable to earn its keep in a cognitive ontology. They conclude that the process denoted by ''response selection'' likely has ontological reality while other labels that fared less well may denote functions that are ''manifest only in the minds of cognitive scientists.'' That is, the term ''ontology'' is ambiguous between its meaning in informatics and its meaning in metaphysics, in which ontology is the study of basic categories of existence (such as ''object,'' ''relation,'' ''event,'' or ''property''). In these terms, their conclusion is that if a term is useful in an informatics ontology, then it probably refers to an existing process, and if not, then it probably does not.

But let's take things a bit slower. A useful informatics ontology will contain syntactic items that are reliably semantically stable and distinguishable; if it is also intended to reflect metaphysical-ontological reality, its terms should denote real (rather than ''emergent'' or ''epiphenomenal'') processes (Cooper, 2010). Focus on the first, semantic, problem and set aside, momentarily, the metaphysical one. Objective semantic identity for syntactically identical cognitive terms used in cognitive neuroscience is determined by the sameness and stability of tasks used within and across cognitive neuroscientific studies over time. In short, a useful cognitive ontology for cognitive neuroscience depends essentially on a useful task ontology.

The results of Lenartowicz et al. show that combining neuroimaging data and text-mining techniques can suggest hypotheses for developing a task ontology. They provide empirical justification for focusing on tasks grouped as indicators of ''response selection'' (as opposed to, say, those associated with ''task/set switching'') as a starting point. This is an informatics variant of Newell's (1973) second proposal for unifying psychology: Focus on a single complex task to develop ''a sufficient theory of a genuine slab of human behavior'' and work outward from there. Newell's unification problem is still with us, in the form of a ''menagerie'' of putative processes of cognitive control (Cooper, 2010) and of psychological processes in general (Uttal, 2001). In the informatics context, however, the problem is a menagerie of meanings of cognitive-process terms. A cognitive ontology is critically unlike the Gene Ontology in that the objective meaning of its terms cannot be fixed by direct observation. This is why the semantics and metaphysics of an informatics ontology require separate careful consideration in psychology but not in biology, why text-mining is a worthwhile technique in psychology only if syntactically identical cognitive terms have the same operational meaning, and why a cognitive ontology requires a task ontology. In short, the basic problem is that we now have a menagerie of tasks, and Lenartowicz et al.'s approach is an interesting way to motivate hypotheses about where to start systematizing tasks.

Some of the issues that now arise (which Lenartowicz et al. do not address, and which I will only describe) can be illustrated by considering the term ''response selection'' and the task labels—''Choice RT,'' ''RM tasks,'' and ''Simon Task''—associated with it (Sabb et al., 2008). The first issue is distinguishing technical from ordinary meanings. Task and cognitive-process labels, such as ''Simon Task'' and ''response selection,'' are technical terms, even though they may be identical syntactically to a description of various things people may do, singly or in groups, in nonlaboratory contexts. (This semantic gap is obvious for terms used in social cognitive neuroscience and comparative psychology, such as ''deception'' (Sip, Roepstorff, McGregor, & Frith, 2007). ''Response selection''—rather than a wholly invented term, such as ''gyring''—is introduced as a label of a putative subpersonal process precisely because the description's ordinary meaning is being used to indicate the general kind of process it is hypothesized to be. But for scientific purposes it must now be given a precise technical meaning in terms of associated indicator tasks (whose labels may be coined with the same purpose in mind). Lenartowicz et al.'s results show that the indicators of ''response selection,'' unlike those for the other component labels, appear to form a homogeneous set from the neural point of view. The first issue amounts to determining that this result is not an artifact of a small sample, and then, assuming the result is robust, determining why the (perhaps expanded) set is homogeneous.

The second issue is the semantic identity of syntactically identical task labels. ''Response selection'' will be semantically identical across studies only if its indicator tasks are homogeneous across laboratories and time. Every task labeled ''Simon Task'' must be identical in the relevant ways. Which are the relevant ways? Unfortunately, task identification and individuation are no more principled now than when Newell raised the unification problem. Whether tasks are individuated by cues, or stimulus-response associations, or both, or something else entirely, the important semantic issue is to ensure the same principle of individuation and the same relevant features are used in every task that falls under the same syntactic label. That is, all tasks sharing the same label should be classified at the same level of generalization using the same relevant features. To illustrate with a familiar example: All tasks labeled ''Stroop'' might be fixed to be ''name the ink color and name the color word'' tasks, and variations—both more specific and more general—should be given different (though related) labels. For the same reason that an informatics ontology in biology that used ''cat'' for cats, mammals, and Siamese cats would be useless, we do not want semantically ambiguous task labels. There may be no one right task ontology, but there needs to be a *systematic* one in which syntactic identity captures semantic identity and so results are in fact comparable (and searchable).

There are certainly more issues to be explored (see, e.g., Bilder et al., 2009). However, ensuring systematicity of task labels and homogeneity of indicator-task groups would go a long way toward providing a stable semantics for an informatics ontology of use in cognitive neuroscience.

Now let's turn to metaphysics. What is the relation between a semantically valid cognitive ontology and the metaphysical issue of when cognitive labels refer to real processes?

The vast majority of experimentalists (at least judging by their peer-reviewed publications) are agnostics regarding the metaphysical relation between mind and brain, and rightly so. Brain-behavior correlations are neutral on mind-brain metaphysics: Substance dualism, emergentism, and reductionism, among other metaphysical positions, are equally consistent with these correlations. Lenartowicz et al.'s study provides no reason to abandon this agnosticism. If inferences to the real existence of cognitive processes based on correlations between imaging studies and uses of cognitive labels were valid, it would follow that syntactic regularity is sufficient to determine reference (as well as semantics, if reference determines semantics). Instead, it is plausible to think that when we can not predict the use of a cognitive label on the basis of neuroimaging data, the label is probably exceptionally poorly operationalized. Their results show that despite the lack of any systematic task ontology, we got lucky with ''response selection.'' Eventually, of course, cognitive labels that cannot be stably operationalized on the basis of a systematic task ontology will not earn a place in our cognitive ontology. Such labels may be eliminated from our informatics ontology or relegated to a nonessential role, and the putative processes they refer to will be eliminated from our metaphysical ontology or at best deemed to exist ''only in the minds of cognitive scientists.''

More generally, the search for a cognitive ontology for cognitive neuroscience is not the same project as the search for neural implementations of cognitive functions. One-to-one mappings in particular would provide the most predictive power, but a cognitive ontology would be extremely useful without them—just as the Gene Ontology is useful without mapping phenotypes one-to-one with genotypes. In fact, the more complex the neural implementation of cognitive processes turns out to be, the more critical a cognitive ontology will be for unifying cognitive neuroscience.

References

Bilder, R. M., Sabb, F. W., Stott Parker, D., Kalar, D., Chu, W. W., Fox, J., Freimer, N. B., & Poldrack, R. A. (2009). Cognitive ontologies for neuropsychiatric phenomics research. *Cognitive Neuropsychiatry*, *14*(45), 419–450.

Cooper, R. P. (2010). Cognitive control: Componential or emergent? *Topics in Cognitive Science*, *2*(4), 598–613.

Figdor, C. (2010). Neuroscience and the multiple realization of cognitive functions. *Philosophy of Science*, *77*(3), 419–456.

Laird, A. R., Lancaster, J. L., & Fox, P. T. (2005). BrainMap: The social evolution of a human brain mapping database. *Neuroinformatics*, *3*, 65–77.

Lenartowicz, A., Kalar, D., Congdon, E., & Poldrack, R. A. (2010). Towards an ontology of cognitive control. *Topics in Cognitive Science*, *2*(4), 678–692.

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283–308). New York: Academic Press.

Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, *10*(2), 59–63.

Poldrack, R. A. (2008). The role of fMRI in cognitive neuroscience: Where do we stand? *Current Opinion in Neurobiology*, *18*, 223–227.

Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology*, *22*(3), 262–275.

Sabb, F. W., Bearden, C. E., Glahn, D. C., Parker, D. S., Freimer, N., & Bilder, R. M. (2008). A collaborative knowledge base for cognitive phenomics. *Molecular Psychiatry*, *13*, 350–360.

Sip, K. E., Roepstorff, A., McGregor, W., & Frith, C. D. (2007). Detecting deception: The scope and limits. *Trends in Cognitive Sciences*, *12*(2), 48–53.

Uttal, W. (2001). *The new phrenology: The limits of localizing cognitive processes in the brain*. Cambridge, MA: The MIT Press.