University of Missouri, St. Louis
IRL @ UMSL

Theses | Graduate Works

4-19-2013

Concepts, Universals, and the Abstract

David Keith McGraw
University of Missouri-St. Louis, dkmcgraw@umsl.edu

Follow this and additional works at: http://irl.umsl.edu/thesis

Recommended Citation
McGraw, David Keith, "Concepts, Universals, and the Abstract" (2013). *Theses*. 205.
http://irl.umsl.edu/thesis/205

This Thesis is brought to you for free and open access by the Graduate Works at IRL @ UMSL. It has been accepted for inclusion in Theses by an authorized administrator of IRL @ UMSL. For more information, please contact marvinh@umsl.edu.

David McGraw
B.S., Biochemistry & Biotechnology, University of Missouri-St. Louis, 2011
B.A., Philosophy, University of Missouri-St. Louis, 2011

A thesis submitted to the Graduate School at the University of Missouri-St. Louis in partial fulfillment of the requirements for the degree

Master of Arts in Philosophy

May 2013

Advisory Committee

Gualtiero Piccinini, Ph. D.
Chairperson

Jon McGinnis, Ph. D.

John Brunero, Ph. D.

Concepts, Universals and the Abstract

David McGraw

University of Missouri-St. Louis

Abstract

Through a historical investigation into the medieval problem of universals, I come to an understanding of the philosophical notion of 'concept' that is compatible with the contemporary discussion of concepts within cognitive science. *Contra* Machery (2009), I argue that the philosophical and psychological notions of 'concept' are compatible by developing the Aristotelian solution to the problem of universals into a conceptualist position. This will then be used to subsume the various paradigms of concepts countenanced by contemporary psychologists. I emphasize a conceptual scheme of reality that is embraced by both ancient and contemporary thinkers: what I call the tetrafurcation of reality. By focusing on the cognitive activity of abstraction, it can be shown that the three paradigms of concepts discussed by Machery reflect three prominent ways in which we abstract from our experience of the world. The three paradigms exhibit a unity that is captured by the philosophical notion of 'concept'.

Concepts, Universals and the Abstract

The purpose of this discussion is multifaceted. Primarily, I wish to contribute to the contemporary and interdisciplinary discussion of concepts by critically evaluating the assumptions underlying Edouard Machery's *heterogeneity hypothesis*. *Pace* Machery, I assert the philosophical discussion of concepts really is relevant to psychological theorizing. He is inaccurate in claiming that philosophers and psychologists are discussing different things when they use the word "concept"; through noting his inaccuracies it will be shown that the abstract use of this term is not something to be eliminated but rather something to be embraced and placed in its proper role within our understanding. There are three questions I wish to address:

Question 1: What is the philosophical notion of 'concept'?
Question 2: What is common between the psychological and philosophical notions of 'concept'?
Question 3: Is the philosophical notion of 'concept' useful in psychology?

In order to answer Question 1, a brief historical detour into the *problem of universals* will be traversed. Beginning with Aristotle, moving through a few medieval philosophers, and concluding with a contemporary formulation of the issue, I will ameliorate the confusion regarding what I take to be the central issue within the problem of universals: namely, the ontological status of the entities posited to solve the problem. The medieval form of the problem is typically answered through adherence to either realism or nominalism, in one of their various forms. I will instead motivate *conceptualism*: universals are *generic concepts* that exist in the mind and are thus neither in the world (as the realists believe) nor merely an artifact of the way we express thought (as the nominalists assert). Adopting a conceptualist view of universals is tantamount to rejecting the traditional realism-nominalism debate as a false dichotomy.

This historical detour into the problem of universals will also give me the opportunity to clarify the notion of *abstraction*. The realism-nominalism dichotomy misconstrues the nature of *abstract entities* and their place in the conceptual scheme introduced by Aristotle and refined by medieval philosophers. I submit theirs is a conceptual scheme that permeates contemporary philosophy of mind: what I call the *tetrafurcation of reality*, a four-fold distinction among the

world, experience of the world, thought about experience, and expression of thought. The theory of abstraction that emerges from Aristotle's corpus can be easily imbedded within the tetrafurcation as a means of accommodating information provided by Machery (2009).

The answer to Question 2 will then become apparent. Briefly, the notion of 'generic concept' serves as the genus of the various conceptual paradigms to which Machery appeals in his hypothesis. Making explicit the tetrafurcation and the place of abstraction within it will also allow me to answer Question 3. It is answered negatively by Machery, and my discussion of the first two questions will allow me to qualify his answer appropriately. I suggest the various conceptual paradigms recognized by Machery correspond to a progression of conceptual abstraction that is identifiable within the Aristotelian tradition. By answering this third question, I will have come full circle; I will have shown that what philosophers mean by the word "concept" is the same thing that psychologists mean when using this term. By bringing psychologists and philosophers together, I will vindicate the use of the abstract notion of 'concept'.

I will end with a discussion of how my suggested unification of the accepted paradigms of concepts might be carried out. I hope to situate the philosophical view of concepts within a framework that is consistent with contemporary views in cognitive science—including computational theory of mind—and which embraces the tetrafurcation mentioned above. As an additional bonus of emphasizing the tetrafurcation, it will be suggested that the *language of thought hypothesis*, although providing useful insight into the conceptual scheme developed here, emerges out of a misunderstanding of the four-fold distinction.

The Problem of Heterogeneity

Edouard Machery presents his heterogeneity hypothesis (HH) in *Doing without Concepts*. This psychological survey goes to great lengths to motivate the need for distinct types of concept and to establish the inadequacy of the abstract term "concept" in psychological theorizing. Through presenting his consensus view of concepts and the psychological evidence of the '*exemplar*', '*prototype*' and '*theory*' conceptual paradigms, Machery concludes that the

heterogeneous and imprecise use of "concept" gives psychologists reason to disregard this excessively abstract label. To be clear, here are the five tenets of HH:

> Tenet 1: Evidence suggests that for each category an individual has several concepts.
> Tenet 2: Co-referential concepts have few properties in common; they are heterogeneous in kind.
> Tenet 3: Exemplars, prototypes, and theories are three paradigms (i.e., three kinds) of concept.
> Tenet 4: Each kind of concept seems to be used in discrete cognitive processes.
> Tenet 5: The generic term "concept" should be eliminated from psychological jargon.

I embrace Tenet 3 and Tenet 4, since they are grounded in empirical evidence that goes beyond the scope of my competency. I will not be challenging his assertion that 'concept' is a vague notion that is used to refer to a diverse group of abstract entities, and I find his overall view compelling; what I find counter-productive is the polemic nature of his view towards philosophers, and I find historical motivation to resist his characterization of the philosophy of concepts. Specifically, I think his characterization of the notion of 'concept' is confused. Thus, what I have to say will be most relevant to Tenets 1, 2 and 5.

Philosophy is Irrelevant to Cognitive Science?

Machery (2009) presents HH in a way that puts philosophers and psychologists at odds with each other, claiming that the two groups are talking about different things when they discuss concepts. Take the first two chapters of his book as evidence of the introduction of conflict between psychologists and philosophers. The first, titled "Concepts in Psychology", begins with this sentence: "The goal of this first chapter is to explain what concepts are taken to be in psychology, neuropsychology, artificial intelligence, and cognitive science" (p. 7). In the second chapter, titled "Concepts in Philosophy", Machery asserts that philosophers do not think of things in the same terms as the aforementioned cohort; while psychologists and the lot are concerned about concepts as batches of information stored in long-term memory that are recalled during higher cognitive processing, philosophers are said to be interested in our ability to have propositional attitudes about the world. Noting this, Machery (2009) moves to discard the philosophical discussion of concepts and focus on what psychologists have to say.

It is unfortunate that Machery explicitly excludes philosophers from the scientific discussion of concepts; surely, not every philosopher is equally and adequately informed about contemporary developments in cognitive science and each of its sub-disciplines, but this is no reason to exclude all philosophers from the inquiry into cognition. The human mind is as equal a part of philosophical theorizing as it is a part of psychology, and there is a coherent evolution of theoretical thought about concepts extending into psychology *and with its roots in philosophy*.

Reconciling Jargon

I find what Machery (2009) has to say about concepts to be, at times, philosophically confused: for instance, his claim that philosophers are not researching what cognitive scientists are researching when writing about concepts. Consider Machery's characterization of psychological concepts:

> A concept of *x* is a body of knowledge about *x* that is stored in long-term memory and that is used by default in the processes underlying most, if not all, higher cognitive competences when these processes result in judgments about *x*. (p. 12)

Compare this to his characterization of what philosophers discuss:

> Having a concept of *x* is being able to have propositional attitudes about *x* as *x*. (p. 32)

Suppose we take these two statements to be answers given in response to questions; they are clearly not answering the same question, as I think Machery would be happy to admit. The first could answer "What is a concept?" while the second could not. The latter statement answers "How do we know someone possesses a concept?", but the former would need to be modified in order to answer this query. A philosopher aware of contemporary research on the brain would probably agree with the psychological characterization of concepts given; in fact, a philosopher ignorant of the current state of cognitive science would probably also agree with this, granted the technical jargon was adequately explained. All this is to note how adopting the philosophical notion above assumes that an acceptable answer to the question, "What is a concept?" has already been established, or at least expects such an answer to be given at some point. I think adequately informed philosophers will find Machery's description of concepts stated above intuitively

appealing. Furthermore, I think it is also reasonable to assert that, in addition to whatever else the mind uses, producing judgments of the kind referenced in the psychological description above necessarily involves forming propositional attitudes. Just as the philosophical notion of 'concept' assumes something like the psychological notion above, the psychological notion relies on the philosophical notion for it to obtain.[1] This seems to be a prima facie reason to investigate how psychologists and philosophers might share an understanding of concepts.

Machery (2009) urges us to consider the empirical data that psychologists have accumulated in order to understand the nature of concepts, and this is surely a reasonable thing to do. However, before investing precious resources into empirical research, philosophers would be wise to test the limits of a particular idea by first considering the consequences of adopting it. In this way, philosophers attempt to understand what concepts can *possibly* be by imagining what they *must* be, given what we know about the mind already. This is to consider the notion of 'concept' in a generic way, and there is a long line of philosophizing that demonstrates this investigation. Of course, seeing as Machery did a lot of great work in synthesizing psychological research, I will utilize his diligence to bolster my view. I am surely indebted to him.

The Problem of Universals

In his discussion, Machery (2009) uses terminology reminiscent of Aristotle in order to discuss the nature of psychological concepts. Ignoring the intricacies of his description of the notion of 'concept' for the time being, take this quote of Machery:

> Most research has focused on concepts of classes of three-dimensional, medium-sized objects, such as animals and artifacts...These classes are usually called "categories" in the psychological literature. There has also been research on concepts of events...and of substances, as well as some research on abstract concepts, such as GOOD, JUSTICE, SCIENCE, or CAUSE. (p. 12)

His appeal to notions deeply seated in philosophical discourse—e.g., categories, substance, "abstract concepts"—compels me to consider the history of these notions in order to begin my response to Machery's claims. In order to resist his attempted ostracization of philosophers, I

[1] Fodor (2008) mentions a similar position, and I agree with it (see pages 12-13 specifically).

would like to trace the roots of the philosophical notion of 'concept', thereby clarifying the relation between the entities that psychologists and philosophers reference when discussing concepts. It is my opinion that the conceptual scheme of the world which we embrace to this day was established by medieval philosophers.

Ancient and medieval philosophers adequately demarcated the terrain of the discussion of universal knowledge, but I will ultimately appeal to the work of a contemporary philosopher, Gonzalo Rodriguez-Pereyra, in order to clarify what, exactly, the problem of universals is and what a solution to the problem needs to be. With the problem properly understood, I can then move to connect the historical understanding of universals to the contemporary discussion of concepts. Through exploring the thinkers below I will show that much of what I have to say has already been said; having said that, a crucial conceptual scheme that lies behind most discussions of the mind must be made explicit, and I capture this in the tetrafurcation of reality. The place of abstract entities within this conceptual scheme will help motivate my overall view that the philosophical treatment of concepts is relevant to cognitive science.

The Ancient Roots of Realism

The problems associated with universals were discussed by both Plato and Aristotle. Both figures were concerned with understanding how we come to have general knowledge of the world. It is clear from their writings that a significant problem which both of them grappled with is that of explaining how multiple particular objects in the world can be described using a single universal term. It is unclear how to explain our ability to see two distinct objects as identical in some particular way, as *having something in common that is itself a unified whole*; this is what Rodriguez-Pereyra (2000) discusses as *the problem of the One over Many*, or P(O|M) for short.

Plato's realism. Much of the discussion found in Plato's dialogues that pertains to P(O|M) involves an emphasis on universals; how can anyone hope to acquire general and lasting knowledge when the world from which we draw our knowledge is constantly shifting from one state to another? Clearly, we are able to make universal statements about objects in our world,

and there seems to be a need for something steady and unchanging to serve as the connection between these statements and the ever-changing world. It is this need that seems to motivate realists to posit the existence of the self-subsisting entities that serve as the reference of our universal linguistic terms.

The Platonic solution to P(O|M) is this claim: distinct objects are similar because they partake in the same *Idea*.[2] I will not delve very far into the works of Plato involving his Ideas and will instead only address one key point of the stereotypical "Platonic" understanding of realism: Ideas are entities that exist outside the spatiotemporal world and independent of the objects which partake of them. This is called a *realist* view of universals because it posits the "real" existence of entities over and above the objects in which we observe the universal quality; these entities are supposed to be the proper referent of the linguistic terms used to indicate some type, property, relation or whatever else we want to attribute to multiple objects.

Aristotle's realism. The Aristotelian view refuses to admit the existence of a separate and independent realm in which these self-subsisting Ideas reside. While many suppose the existence of Ideas is a reasonable attempt to solve P(O|M), Aristotle and those that sympathize with him are also reasonable in being uncomfortable with this solution, and Aristotle finds many reasons to reject the Platonic theory of Ideas. Intuitively, it is troublesome to posit a realm of entities with which we have no direct contact in order to understand the way we think about and express our experiences. It seems strange to need these entities in order to explain something as apparently straightforward as our interaction with the world of sensible objects.

Aristotle describes a universal as "what is naturally predicated of more than one thing" (*DI* 17a38-39). Elsewhere, he says, "By 'universal' I mean what belongs to its subject in every case and in its own right, and insofar as it is itself" (*APo* 73b26-28). There is an intimate connection between our language and our knowledge of the world, and Aristotle recognizes this by emphasizing the subject-predicate form of our universal claims. Given this, Aristotle develops

[2] I will use the word "Idea" with a capital "I" when addressing Platonic universals.

an epistemology that does not appeal to innate knowledge of universals. He views the experience of particular objects and their various aspects as the basis of our knowledge, and our knowledge is of these objects *and nothing else*. The forms or essences exist in the particular objects which possess them, and any given universal term refers to the instances of a property or relation or whatever that we identify in those particular objects. Essences are not separable from their instances; the fact that we can think of them in a separate sort of way does not somehow bring them into existence independent of the particular. Aristotle clearly understands universal claims as statements applying to many particular objects, but this is due to some brute fact about the things themselves. Sir David Ross (1995) describes Aristotle's overall view thusly:

> The world which is given to us in experience is a world of concrete individual things acting and reacting on each other. In contemplating these we become aware of characters common to many individuals. These are for Aristotle as real, as objective, as the individuals. They are not in any sense the work of the mind any more than are the Forms to Plato. But he warns us to assign to them only that mode of existence which is proper to universal, viz. existence as characteristic of individuals. (p. 164)

This makes it clear that Aristotle is considered a type of realist distinct from the Platonic type, and Aristotle gives many reasons for denying the existence of Ideas.[3]

Suggestion of consensus. This is a hasty exposition of these views, but it establishes a major distinction that is drawn within the realist camp. Plato and Aristotle are often portrayed in a way that puts them at odds with each other, and the brief discussion above makes note of their differing metaphysical and epistemological views. What can also be seen is that these two philosophers agree on a significant claim: humans have an ability to form universal knowledge through repeated exposure to the world of particular objects, and we can convey this knowledge through the use of universal linguistic terms. It is clear that this point of agreement assumes three substantial demarcations: (1) that between the world and our experience of the world, (2) that between our experience and our thoughts about experience, and (3) that between our thought and our expression of thought. This is the essence of the tetrafurcation of reality, and it is the precise

[3] Many of the arguments against the Ideas are given in Plato's *Parmenides* and are repeated by Aristotle and many medieval philosophers. One in particular, the third-man argument, will be relevant later.

location of universals within this scheme that suggests the compatibility of the psychological and philosophical discussions of concepts.

It is rather obvious to me that what the ancients called "universals" are things that are represented in the mind; whether or not there are "real" entities in the world named by our abstract universal terms, there is something in our minds that we mean to communicate using each of these terms. Taking universals (as they are represented in the mind) as uncontroversial examples of concepts, I am compelled to consider the process by which we arrive at these generic concepts.

Abstraction

Universals are relatively simple when compared to particular objects. For example, blueness lacks any quality of texture, scent, sound or taste—at least for non-synesthetes—while blueberries have all of these aspects. In this way, universals are utterly different from normal objects we experience. How we arrive at these relatively impoverished notions is through some sort of cognitive act of isolation, as both Platonic and Aristotelian realists would acknowledge. Since this activity of the mind is so fundamental to the generation of universal knowledge, it is obvious why past and contemporary philosophers alike have been interested in it. It is my contention that there is a vague notion of *abstraction* tossed around rather loosely in philosophy. The relevant difference between psychologists and philosophers is that the former are trying to clear away the vagueness by cataloging the various ways in which we can abstract, while the latter clear away the vagueness by investigating the larger theoretical framework in which the notion of abstraction is imbedded. Both are valuable and necessary epistemic pursuits (though, being a philosopher, I am partial to and better trained in the latter strategy).

Abstract entities. Michael Loux (2002) explicitly attributes the above view—that universal claims are grounded in our experience of objects in the world—to realists of both Aristotelian and Platonic persuasions. What he calls *abstract singular terms* function as names of universals. Realists think our statements that include these terms make claims about entities other

than "familiar concrete particulars" (Ibid., p. 34). What kind of entity is not immediately clear, though, and the realists give at least two different answers to this prompt: the Aristotelian and Platonic solutions. I think realists pulled in either direction will agree that just as particular objects we encounter in the world are treated as the subjects of predication, abstract singular terms corresponding to predicates can play the role of a subject as well; this creates an invitation to consider each abstract singular term as referring to an object, as subjects of predication typically do. When we consider that object in itself, we are considering an *abstract entity*.

Aristotle, Cleary & aphairesis. Many philosophers have used the term 'abstraction', and there are at least two very different ways of understanding this term in modern times. One understanding has a non-physical, other-worldly connotation; when Plato's Ideas are called 'abstract', it is this understanding that is being asserted. Another understanding of the word is linked to the Greek term ἀφαίρεσις ("aphairesis") that is found in Aristotle's writings. This usage connotes a partial or incomplete nature, and it is this usage which, I submit, is the appropriate way to conceive of 'abstraction'. The tendency to assume the former connotation, or, rather, mistaking usage of the latter kind to be of the former, is the source of the confusion surrounding the ontological status of universals and disguises the consensus that lurks behind the scenes.

To alleviate the tension within this notion, I appeal to John Cleary's "On the Terminology of 'Abstraction' in Aristotle". Cleary (1985) asserts that translating the Greek "aphairesis" as "abstraction" is an unfortunate occurrence, for the Greek term had a technical usage at the time of Aristotle's life that becomes hidden when the word "abstract" is used with the non-physical connotation. He argues that 'aphairesis' is the counterpart of another technical notion which should be more familiar to the English-speaking world: πρόσθεσις ("prosthesis"). This word is primarily used in the context of *prosthetic limbs*: artificial body parts used to make a body complete again. The limbs are being *added to* a body, thus preserving the ancient usage of "prosthesis" to indicate a logical process of addition that forms a new whole from multiple unified parts. The opposite of this would be the division of a whole into parts that are

subsequently treated as wholes themselves; this leads Cleary to assert that aphairesis, being the counterpart of prosthesis, is a logical process of *subtraction*. Cleary is asserting that Aristotle uses "aphairesis" to indicate the process by which we come to understand the reference of so called "attributes", what I call *aspects* of objects:

> …subtraction (or abstraction) is a logical method which allows one to intellectually isolate the primary subject of a given attribute in the following manner. One focuses upon a particular attribute (e.g. having internal angles equal to two right angles) and asks: to which aspect of a concrete sensible triangle does this attribute belong primarily? (p. 22)

In other words, abstract singular terms ultimately refer to the particular objects and their various aspects that we experience through our senses, and abstraction is, for example, the consideration of a particular aspect of an object as a thing in itself. While Platonic realists wish to say that this consideration necessitates the independent existence of a referent for abstract singular terms, this is not something inherent to the notion of abstraction.[4] Instead, realists can make the weaker claim that abstract singular terms refer to some sort of partial understanding of objects and their aspects: an understanding existing in the mind that is about the world as it is experienced. Partial understanding is the product of the logical process of abstraction (in the vague sense of the world), but this is not to say that only one kind of concept is produced through this process. The world appears heterogeneous, so it is to be expected that our mental life accurately reflect that.

Medieval Musings

Aristotle captures the heterogeneity of the world within a handful of ways of being. In the *Categories*, he introduces ten ways of being—called "categories"—that seem to be exactly the kinds of things abstraction produces in the mind. The categories serve as a list of the various ways we experience things in the world. Aristotle makes an explicit distinction between the categories of things in the world and the linguistic terms we use to refer to those categories.[5] Within Aristotle's thought in general, the primary sense of existence is 'being the subject of predication',

[4] The Platonic solution involves a disregard for the tetrafurcation of reality; it supposes our expression necessarily refer to objects in the world, when it only truly refers to thought about experience of the world.
[5] For Aristotle, "category" refers to the basic ways of being; Machery (2009) uses "category" and "class" interchangeably and is thus deviating from the traditional philosophical use of the word.

i.e. being a "substance"; of the ten kinds of existence that are the categories, only 'substance' is neither said of nor is in another thing; in this regard, only a substance is an object—a "this"—in the truest sense.[6] Because, according to Aristotle, the accidents of a substance depend on the substance's existence for their own, what we really know when we acknowledge properties, relations, and other instances of the categories are *aspects of particular objects*. Our ability to recognize these various aspects and consider them as things in themselves just is the process of abstraction. To put it roughly: the objects in the world are the basis of the knowledge we express with our words because, usually, our words refer to objects and their aspects. The distinction between our thoughts about the world and the world itself is what I take to be the distinction between the abstract and the concrete; the world is concrete, and our thoughts are "abstract", i.e., incomplete and partial conceptions of the concrete world. Though we can make this distinction, there is an implicit and typical correspondence between the abstract and the concrete that is captured by the tetrafurcation.

Medieval philosophers in the western world latched onto Aristotle's philosophical works, particularly his *Categories*. Some of the more significant figures of the Middle Ages gave their attention to the problem of universals as it arose from the *Categories*; their work foreshadows contemporary discussions, so it will be relevant to discuss these thinkers. From these medieval musings we can see an expression of my suggested tetrafurcation and the centrality of the notion of abstraction in explaining generic concepts.

The medieval problem. Porphyry wrote an introduction to the *Categories* as a means of clarifying the logic of Aristotle. In the first few statements of this work, Porphyry explicitly sets aside three questions: "(1) whether genera or species exist in themselves or reside in mere concepts alone; (2) whether, if they exist, they are corporeal or incorporeal; and (3) whether they exist apart or in [sensible] objects and in dependence on them" (Warren, 1975, pp. 27-28). It is these three questions which Boethius addresses in his second commentary on Porphyry's

[6] "Substance" is used in the *Categories* to indicate individual objects: 'this particular man' or 'Socrates'.

Introduction, and it is here that the medieval problem of universals begins to take shape. The categories for which Aristotle's work is named represent many of the things which we call universals; what Porphyry and Boethius say about 'genus' and 'species' can be extended to universals in general and, I submit, to most (if not all) concepts.

Realism per Boethius. In his second commentary on the *Introduction*, Boethius begins discussing these questions by stating the initial problem to be addressed: "Genera and species either exist and subsist or are formed by the understanding and by thought alone" (Spade, 1994, p. 21). Mirroring several criticisms of Ideas that can be found in the ancients' writings, Boethius proceeds to argue that universals do not "exist and subsist". If we assume universals exist independently and are common to more than one object, we end up asserting that something that is singular (the universal) is also multiple (since it is wholly present in each particular), or we conclude that what is multiple and identical across particulars necessitates a further universal to unite the instances with this additional entity (a version of the third-man argument—see note 5). Given this, we can assume that universals are "formed by the understanding and by thought alone".[7]

Boethius then notes that our understanding can deviate from the way things actually are, but this does not necessarily imply our understanding is false and empty.[8] He says:

> False opinion rather than intelligence occurs only in those cases that arise from composition… But if this understanding arises *from division and from abstraction*, then the thing is not disposed the way it is understood, and yet that understanding is not false at all. (Ibid., p. 23, emphasis added)

He offers us an example to make this clear. The notion of a line is something which we can consider on its own, independent of any particular object in the world; considering these things on their own is to consider them differently from how they are actually disposed, though these things nevertheless still exist as they do (i.e., as aspects of particular objects in the world). That is to say, we can conceive of things like lines, shapes and colors apart from any instance of these things,

[7] I submit this is tantamount to denying the existence of Plato's Ideas, although Boethius claims to be neutral towards this conclusion. See note 10 (Spade, 1994, p. 25).
[8] The philosophical import of the notion of opinions arising from "composition' will be noted later.

yet this in no way make lines or shapes or colors subsist on their own. Boethius summarizes his conclusion thusly: "…for genera and species—that is, for singularity and universality—there is one subject. But it is universal in one way, when it is thought, and singular in another, when it is sensed in the things in which it has being" (Ibid., p. 25).

Boethius refrains from judging either the Platonic or Aristotelian view to be superior. He clearly sympathizes with the criticisms of Ideas and chooses to discuss universals as those things that populate our understanding rather than as self-subsisting entities, but he also claims to be agnostic by refusing to pass judgment. What I find to be most insightful is captured in his explanation of how a universal arises as a species and then comes to be considered as a genus. We think about the "likeness" or similarity among particular objects, and we see this resemblance as a legitimate aspect by which to classify objects in the world. Then, we consider the likeness in itself, knowing full well that the likeness is *amongst particular objects*, and we come to see it as a genus: a class of particulars, a universal. This is how I see Aristotle and Plato coming together, through the words of Boethius: "[Likenesses] subsist therefore in the realm of sensibles, but are understood apart from bodies" (Ibid., p.25).We experience the various aspects of particular objects, acknowledge the similarity amongst the instances, and are able to shift our consideration from the external objects to the experience of their aspects as objects themselves.

Conceptualism. This shifting consideration is what I call *abstraction*. Boethius also labels this process "abstraction", and many subsequent philosophers continue to use that word in order to indicate the cognitive process that is occurring during such shifts in consideration. As mentioned above, realists believe universals are signified by *abstract* singular terms, and, to be sure, Ideas are the canonical abstract entities—albeit an idiosyncratic usage of "abstract", as noted earlier. But the word "abstract" simply means partial and incomplete, which may by happenstance mean non-physical. There is nothing essential to the process of abstraction that necessitates a non-physical existence; although the term "abstract" is the alternative to "concrete", this is to emphasize the fact that abstract entities are partial and incomplete. Abstract

entities are not in the concrete world in the same way that particular objects are because some feature of the concrete instances from which we produce the abstractions has been subtracted.

Because of the way Boethius concludes his commentary, I consider his discussion to be one that attempts to understand precisely what Aristotle meant in his philosophizing.[9] This tempts me to consider Aristotle as less of a realist and more of a *conceptualist*; while it seems clear that Aristotle is received as a realist, Boethius's treatment of the problem of universals seems to indicate that Aristotle's views can be accommodated by a conceptualist understanding of universals. By "conceptualism" I mean the view that universals are what arise when we consider our experience of the world and attempt to understand our varied experiences in a general way. They arise within our thought as unified entities, but they are the unification of multiple distinct experiences; since the existence of universals is typically parasitic on the existence of its instances, the extent to which universals exist is the extent to which we unify our experience through some sort of cognitive activity.[10]

Clearly, there is some "real thing" in the world to which our universal statements refer, but the unity that we conceive is just that: a *conception* of our world that emphasizes a particular aspect of our experience, to the exclusion of all other aspects. We typically label such conceptions with linguistic terms, but the linguistic term is not the complete extent of the existence of universals, nor is it necessary for the concept to be said to exist; the labels are labels *of something in the mind* that is about the world as it is experienced, and the subtle differences among our experiences are set aside while we focus on the similarity. So, universals are more than just words, and they are less than particular objects. Universals are in the world only insofar as our experiences are of things in the world.

---

[9] "I did not regard it as appropriate to decide between [Plato's and Aristotle's] views. For that belongs to a higher philosophy. But we have carefully followed out Aristotle's view here, not because we would recommend it the most, but because this book, [the *Introduction*], is written about the *Categories*, of which Aristotle is the author" (Spade, 1994, p. 25).

[10] I say universals are *typically* parasitic on their instances for their existence because I do not deny the possibility of somehow acquiring universal knowledge of non-instantiated traits.

Whether or not Aristotle himself was a conceptualist is irrelevant; what I deem significant is that the epistemological view presented by Aristotle is conducive to a conceptualist interpretation. There is a distinction between the words used to indicate universals and the things in the world that are the basis of our knowledge of universals, and the tendency to believe our words refer to objects leads to realist opinions. What I assert is that our words express thoughts about our experiences, and our experiences are fragmented into chunks that are not always metaphysically significant; while we experience objects, we experience them *with various aspects* that are *thought of as* objects themselves.

Nominalism and conceptualism. One philosopher in particular introduces a fourth question to Porphyry's original three in order to address the relation that holds between the world and the linguistic terms corresponding to universals. Peter Abelard's "Glosses on Porphyry" introduces this question: "Do genera and species, as long as they are genera and species, necessarily have some *thing* subject to them by nomination?" (p. 27, emphasis added). I take this to be asking whether or not the concepts we use to designate genera and species—as well as the other categories—are designating more than the particular objects to which we apply the corresponding predicates; is there an object for every word used in such a way? Abelard makes note of the fact that philosophers chronologically preceding him attribute universality and particularity to both objects *and* words, and he subsequently confronts two different views: one which corresponds to the Platonic view of Ideas, and another being what I take as a resemblance class theory that, he suggests, is consistent with the views of Boethius. Abelard claims these views attempt to make objects out of universals, and they fail to do so; universals-as-words is the remaining option, which he then proceeds to investigate.[11]

Reminiscent of Aristotle, Abelard begins the presentation of his own view with a discussion of predication. Predication, he says, is a sort of truthful conjoining of a particular

[11] While I do not completely agree with Abelard's refutation of the resemblance class theory and feel his own view appeals to some sort of basic resemblance among particulars, I will not address this here.

object and a predicable feature, expressed in the form of an existential verb of some kind and conjoined to the object through some sort of convention or institution.[12] A word signifies its object through some convention, but the convention gets applied to many different particulars; the distinct particulars all have what Abelard calls a *common cause* that makes the predication appropriate. This common cause among the particular objects amounts to a likeness conceived in the mind: some understanding that the mind generates within itself and for itself, drawn from the various particular objects the mind encounters through sensation.

Abelard claims anyone who invented names "meant to impose them in accordance with certain natures or characteristics of things" (Ibid., p. 46). It is this set of characteristics—the common cause—that the mind conceives and identifies as the subject of universals. This sounds similar to how concepts are characterized, and I hope from this discussion it can be seen that calling Abelard a nominalist is not quite accurate. While he has been labeled as both a nominalist and a conceptualist (cf. Honderich, 2005), Abelard's views do not seem to support the claim that universals are *merely* words, since he explicitly claims universal terms signify something else, namely, concepts in the mind, that derive from particular objects in the world. Thus, I consider Abelard to be a *conceptualist* and not a nominalist. This is significant because it emphasizes the importance of clarifying the role language plays in our understanding of the world: words refer to the mental stuff that results from experience of the external world, and similarities within this mental stuff is the source of the universal quality of our abstract singular terms.

What we see in Abelard is the tetrafurcation that I am emphasizing. He acknowledges the distinction between words and the things they refer to, but he also qualifies this in a way that illustrates the distinction between, on the one hand, what we think and, on the other, what causes our thoughts. His "common cause" leads to a "likeness" in the mind; we then form a word used to name that likeness and treat it as a subject, although it is not truly an object (as the particular

[12] There is what Abelard calls *syntactical conjoining*, what Boethius calls "composition", but this is not the conjoining that is evident in universal claims.

objects of our experience are). Abelard claims the world of objects impinges on our senses in a way that produces entities in the mind, which we then isolate from the total experience and label with linguistic terms. Similarity among multiple mental entities—which arises from the similar way in which particular objects impinge on us—leads to the linguistic term being applied to multiple objects univocally. Just as the quote of Ross (1995) given above characterizes Aristotle as asserting that "we become aware of characters common to many individuals", Abelard holds the similarity we recognize among particular objects to be a brute fact relative to our thought about the world.

Conclusions Thus Far

From our philosophical predecessors of the ancient and medieval periods, we receive an explanation of universals that emphasizes the necessary relation between the particular objects in the world and the universal knowledge we acquire form them. Aristotle has an obvious grasp of the abstract/concrete distinction, and he notes that our linguistic systems are used to express the way in which the concrete world impinges on us. As was mentioned, Boethius's discussion includes the recognition of the process of abstraction and its place in Aristotle's system, and he clearly means this to be a process that happens within our minds; the result of this cognitive activity must be mental structures that correspond to the concrete world of objects and their aspects. Abelard reiterates this, claiming the resemblance we note among particular objects grounds the use of labels to communicate these features of our experience. As we now will see, Avicenna imbeds the theory of abstraction within this conceptual scheme of the world which is, I submit, still embraced to this day.

Avicenna on Abstraction

As should be clear now, the medieval philosophers emphasized the fact that we abstract our world into intelligible chunks that are then assigned generic labels. One Arabic philosopher,

Avicenna, made note of this process and situates it within a broader theory of abstraction.[13] He discusses *essences* (instead of universals), which basically correspond to what Abelard calls the "nature or characteristic of things" and what Aristotle calls the forms of things; essences also happen to be what Plato calls Ideas, and it is here that the notion of abstraction becomes essential for properly understanding universals. Avicenna says essences can be considered in three ways: in the particular objects, in our minds as those objects are conceptualized, and finally in a way that captures what the previous two considerations have in common. This gives us a hierarchy of abstraction, where the essence-in-particular is most individual—i.e. most infused with particularizing features of physical existence—and the essence-in-itself is the most abstract.

The essence-as-conceptualized is an intermediate level of abstraction, containing logical relations, degree of commonality of predication, and other accidental features. This, then, gives us a progression of our conceptual knowledge of the world. First, we encounter objects in *sensation*, with essences particularized to the material things being sensed. Next, we can *imaginatively* encounter objects, where the essences are abstracted from the material accidents and are particularized only by its logical relations, among other such things that allow these essences to be used in thought processes. These essences-as-conceptualized are what Jon McGinnis (2007) calls "likenesses of essences as those essences exist in concrete particulars and sensible objects" (p. 172). Finally, we can abstract away all the accidental features of the instances of an essence and consider the likeness as a thing in itself.

McGinnis (2007) says, "…abstraction for Avicenna simply seems to be a process of selectively attending to certain essential features of the sensible object…to the exclusion of other accidental features" (p. 173). This gives us a little clarity on how the transition to higher levels of abstraction occurs. Through some sort of cognitive activity, we sequentially abstract away aspects of any particular object we encounter in the world until we reach a point where we can no longer remove anything from the notion and still capture the aspects of our experience we mean to

[13] McGinnis (2007) provides an accessible discussion of Avicenna's philosophy.

capture. This irremovable aspect or set of aspects becomes the universal notion, the essence, or the similarity we find among particular objects, considered in itself. Thus, the sequential abstraction from experience of the world leads to concepts of entities in the world, and we can consider these in their own right and apart from the material source of the concept.

Question 1: Answered

What is the philosophical notion of 'concept'? Concepts are mental stuff that results from the processing of our experience of the world. There are unified mental entities generated by the activity of the mind that allow us to make meaningful expressions that correspond to the world in some reasonably robust way. This is an abstract way of thinking about things, and this is precisely the point. Whatever the nature of this activity is, there must be something inside the mind that allows us to connect our expressions to the world. The world is diverse, so the mental stuff that mediates between the world and our expressions must be of a nature that represents this diversity. With this in mind, it seems quite appropriate to suppose 'concept' refers to a genus with species. What unifies the genus is the fact that it must exist in a way that captures the variety in the world and allows for the diversity of expression as well as instantiation and functionality in the mind.

The Proposed Solution: Tetrafurcation of Reality

The world is populated with particular objects with various aspects and relations. Our experience is populated with coordinated sensations of those aspects, and our thoughts are populated by concepts that represent those coordinated sensations. Finally, expression is populated by objects with intentionality; whether or not they are written language, spoken sounds, or gestures, these things serve as vehicles of communication. What these various modes of expression actually communicate is difficult to elucidate, but at the very least we can say they express thoughts about one's experience of the world. This is the abstract understanding of the world that we receive from antiquity. I submit that this is also the conceptual scheme that contemporary researchers adopt. Though it is vague, it captures a significant progression of information transmission that must be recognized and explained. Also, it emphasizes the

necessary basis of our understanding: experience of particular objects, in all the ways they appear to our mind. There is much more to be said about each of the four partitions—much more than can be said here—but I hope it is clear that these are not arbitrary distinctions.

The Problem of the Many over One

Aristotle and those that embrace his general view shift the focus of the problem of universals away from P(O|M) and towards another issue: *how is it that a single, unified object can be described in so many ways, or said to be so many things?* This is what Rodriguez-Pereyra (2000) calls *the problem of the Many over One*, or P(M|O). By emphasizing the basis of the universal application of predicates to be the individual instances of predication, P(M|O) is getting closer to the heart of the issue. You cannot predicate something of many objects without first predicating something of one, so we must come to understand how it is that we can predicate something—and, more specifically, how we predicate *many things*—of a single object. Rodriguez-Pereyra establishes this claim by noting that the various forms of universal predication, such as '*a* and *b* are of the same type *F*' and '*a* and *b* share the property *F*', can be broken down into statements of the form '*x* is of type *F*' and '*x* has the property *F*'; explaining these simpler statements will in turn explain the more complex ones, so a proper explanation of the problem of universals must account for how we understand single objects to have multiple aspects to their existence.

One contemporary philosopher, David Armstrong (1989), discusses the problem of universals in terms of the *token-type distinction*; he defines the problem as one of explaining the sameness of type among multiple tokens. As a means of clarifying this distinction, he also addresses the notion of a *class*, which is more or less a group of tokens that is formed based on certain criteria—variously described as resemblance, similarity, or identity among the tokens. He then says, "…we can formulate the Problem of Universals in the following way: *What distinguishes the classes of tokens that mark off a type from those classes that do not*?" (p. 13). This way of framing the problem makes it clear that Armstrong is addressing P(O|M) and not the

more basic P(M|O), but it emphasizes the centrality of the notion of resemblance among particulars to the problem of universals. Thus, it can be seen that we should follow Rodriguez-Pereyra (2000) in answering P(M|O) rather than P(O|M). For it is only after explaining P(M|O) that we can begin to answer P(O|M). Only after recognizing particularized tokens can we understand a type based on the resemblance of those tokens.

The development of *truthmaker theory* is something Rodriguez-Pereyra and Armstrong both embrace.[14] This theory claims that what makes a proposition true is that there is something in the world—a *state of affairs*, in the preferred terminology of Armstrong (2010)—that makes that proposition true. Because universal propositions can be reduced to propositions about particular tokens, the truthmakers of a universal proposition are just the truthmakers of the various propositions about the tokens referenced in the universal claim. Our ability to form propositions about token states of affairs relies on our ability to abstract from our experience. This is all compatible with the tetrafurcation presented above.

Embracing the Tetrafurcation

I have said before that the conceptual scheme developed by medieval philosophers is still embraced today, and I can now formulate a consensus view that takes all of the above into account. There is a world outside us, and it consists of objects with various aspects and relations to each other—the truthmakers of our propositions—that we experience through our senses. We then fragment our experience and put labels on the different parts, storing this information internally. We are able to think about these parts and their relations; we can dissect them further, and we can combine them. Usually, we express these thoughts, and we then must establish some sort of conventional system used to allow interpersonal understanding. Regardless of whether our thoughts are expressed, it is true that our experience is processed into intelligible chunks to make further manipulation manageable. Thus, we acquire abstract concepts: representations that apply to multiple objects in a way true to the resemblance of the experiences of those objects. This

[14] See Rodriguez-Pereyra (2000) and Armstrong (2010).

picture contains the tetrafurcation of reality because it makes explicit appeals to the world and the particular features of it, our experience of those features, our thought about our experience, and the (primarily linguistic) expression of our thought.

Abstraction is supposed to be the process by which we come to understand concepts, and we can here begin to see the ambiguity in this term; for starters, our concepts are abstract, but so are the experiences to which those concepts refer. Abstraction *itself* is an abstract notion, meaning it is a process devoid of any features that would allow us to make distinctions within the various entities the process produces; it also ignores the possibility that what we call *the* process of abstraction might be *several* processes. Regardless of this ambiguity, we can note that concepts must be abstractions from our experience of the world. Even though we possess various types of concept, they all must be represented in the mind and contain information gleaned from our environment through our senses.

With the proposed tetrafurcation made explicit, it becomes apparent where many people have gone wrong when adressing the problem of universals. By taking the problem of universals to be P(O|M), philosophers have clouded the distinction between the vehicle of expression and the world itself. Our language serves to express our thoughts about our experience of the world. We typically expect that our statements correspond to structures in the actual world, since we use language to express our thoughts, our thoughts are about our experience of the world, and our experience is reliable. There is, of course, nothing to guarantee the reliability of experience, nor is there a guarantee that our statements actually align with our thoughts; however, the reliable correspondence among the world, our experience, our thoughts, and our expressions is precisely what someone with a desire for systematic understanding seeks.

Clearly, we are able to carve our experience in a way that fails to capture the structure of reality. There is reason to believe a correct and an incorrect way to abstract from our experience both exist. Moreover, As Boethius notes and Cleary (1985) reiterates, there is at least one other way to logically produce concepts besides abstraction, and that is through *prosthesis*: composition

of our experiences in ways that deviate from the way actual objects are constituted. This is where things like unicorns, centaurs, and the like arise; it is easy to see how these thoughts are neither empty nor do they refer to actual objects in the world. They rely on our sensory information but are the result of assembling abstract concepts into artificial wholes. Contrary to the Platonic realists, this does not thereby bring new creatures into existence. Clear and careful conception is not sufficient for metaphysical actuality, but this is precisely what realists in the Platonic tradition believe. Thus, I diagnose the error of the Platonic realists as assuming that anything we conceptualize exists in the external world.

The way in which we experience the world may be subject to evaluation based on its correspondence with the other realms of reality, but this is more opaque and harder to elucidate. Therefore, for most intents and purposes, one of the more basic assumptions that we must make in order to begin investigating the world is that our experience is a reliable source, indeed the only source, of true information. Also, we must be able to access significant portions of this information consciously in many situations. As Machery (2009) notes, "although introspection occasionally misleads, there is little reason to doubt that when concepts can be introspected, we have partial, but accurate introspective grasp of the [information] they store" (p. 87). There is usually no reason to suspect our senses are being deceived, so we can usually assume a direct correspondence between our senses and the actual world.[15]

Abstraction in Cognitive Science: An answer to Questions 2 & 3

All this is to say that in discussing universals, medieval philosophers in the Aristotelian tradition were actually discussing a type of concept; for these philosophers, concepts are simply the mental entities that result from our experience of the world. Abstraction at least partially consists of whatever manipulations we happen to perform on our experience to arrive at the unified understanding captured by abstract singular terms. It seems clear to me that the abstract notion of 'concept' that surfaces from ancient theorizing is consistent with the psychological

[15] I realize this is a controversial suggestion, but I do not wish to discuss this any further, as it is tangential.

notion of 'concept' offered by Machery (2009). From Avicenna, we receive a view of abstraction that implies a sort of hierarchy of conceptualization; if we embrace this scheme, we can begin to see that the various types of concepts Machery and other contemporary cognitive scientists endorse might correspond to significant demarcations within the sequential process of abstraction. Even if this does not hold—and I think it does—it is clear that philosophers, just like psychologists, are concerned with the mental stuff that results from organizing experience, and this is what is labeled "concepts".

Concepts[16]

Contemporary theories of mind help us bring this abstract theorizing closer to the realm of the concrete. Having understood universals to be concepts in the mind, I would like to present one way to understand this abstract notion consistent with contemporary research in cognitive science. I propose my understanding of concepts has these features: it is compatible with the computational theory of mind; it can function in a manner compatible with conceptual atomism; and it can be used to integrate the prototype, exemplar and theory paradigms of concepts. My contention is that the different paradigms of concepts reflect the demands that psychological studies put on the participants, but they also represent significant aspects of the way we learn about our world. The mind uses only the information that is pertinent to a given task in order to complete that task. In some cases we only need exemplary information, others only prototypical information, still others only theoretical information; this does not mean that there is a lack of unity within this information.

What I propose is this: CONCEPT is a notion that refers to entities that must perform a variety of functions in our thought, and there are pragmatic reasons for keeping this notion vague. I am compelled to believe each of the conceptual paradigms is correct in certain situations and inappropriate in others. Below, I attempt to show how the prominent views of concepts can be

[16] In this section, I adopt a modified nomenclature in order to make more readable: SOME CONCEPT means what I have previously called "the notion of 'some concept'".

combined. I wish to emphasize that I doubt the stuff in our minds can be cleanly and precisely carved into the various conceptual paradigms that will be presented; Machery (2009) discusses the *ecological validity* of psychological experiments on categorization and concept learning, and these sorts of concerns will help me tailor my unifying view.

Theorizing in the Background

In order to present my understanding of concepts, it will help to see the theoretical foundation upon which I am constructing it. Luckily, much of what the ancient and medieval philosophers had to say about abstraction and the tetrafurcation of reality is compatible with contemporary theorizing about the mind. The *representational theory of mind* (RTM) and the *computational theory of mind* (CTM) are two of the main views that I intend to accommodate with my proposed conceptual organization. I also mean to discount a prominent view towards the mind: the *language of thought hypothesis* (LTH); I believe the tetrafurcation of reality will allow me the resources to do so.

Computational theory of mind. This theory first came to be as a result of the combined efforts a neurophysiologist—Warren McCulloch—and a biophysicist—Walter Pitts. The duo was greatly influenced by the work of Nicolas Rashevsky, a prominent biophysicist and founder of the Committee on Mathematical Biology. McCulloch and Pitts joined together to produce an interdisciplinary theory that employed both of their techniques to explain how the mind can function in a way that produces mental phenomena. Their theory employs simplified neurons linked together in "nets" that were argued to be capable of performing digital computations. McCulloch and Pitts's theory supposedly demonstrated how the mind could be generated by a (properly organized) brain, thus solving the mind-body problem. This was the first explicit theory of mind that made a claim to explain mental phenomena via computation.[17]

So, what exactly *is* CTM? Roughly, it is the idea that the mind performs computations in order to produce outputs in response to inputs; whatever is involved in cognition is a type of

[17] All of this is reviewed in Piccinini (2004).

computation. While computation in the mathematical sense is a process that uses an algorithm to produce correct outputs from inputs, Piccinini & Scarantino (2011) distinguish between two ways in which the definition of "computation" can vary in the relevant literature: how restrictive the term is, and whether or not "being the vehicle of computation requires possessing semantic properties" (Ibid., p. 6). *Generic* computation is a manipulation of vehicles according to rules defined over these vehicles. These vehicles needn't be a certain type of medium, so long as they are diverse enough to facilitate an appropriate level of manipulation. *Digital* computation is a subcategory of generic computation; it is more restrictive (i.e., less abstract). A digital computing system must receive a string of digits as its vehicle, and it must manipulate the string of digits according to rules that *exclusively determine* the relation between the input string of digits and the output string of digits. Turing machines are a famous type of digital computing system.

Computation can be even more restrictive than this; *classical* computation is digital computation whose rules of manipulation are a step-by-step algorithm that is represented within the computing system. All of this is rather technical, and it is often difficult to distinguish what type of computation needs to be used when. Digital computation was what McCulloch and Pitts had in mind when formulating CTM, but it is now clear that if the mind is performing computations, it is a form of computation that is not clearly or exclusively digital. We can simply call it *neural computation* (Piccinini and Bahar, 2012). As we learn more about how the different parts of the brain communicate and the specific details of the signals being used, we will be able to characterize neural computation; that being said, there is enough evidence to discount the claim that neural computation is digital.

Generic computation is a notion that seems to necessitate the tetrafurcation of reality. A computing system receives input from an external source, manipulates the information carried by that input, and generates some output signal. Thus, we have a world in which the computing system exists. Stimulation from this world then gives rise to internal reactions; the internal reactions are processed, and external reactions are emitted back into the environment. This

abstract analogy suits things well, and it seems reasonable to take the imbedded functional organization seriously. Concepts are things in the mind and so, by analogy, correspond to the computing system's internal composition. Of course, CTM is meant to be more than an analogy, but even as an analogy it is a powerful thought.

Representational theory of mind. RTM fits with CTM quite nicely. CTM holds that computation is performed according to rules sensitive to a particular type of vehicle. RTM is the idea that the mind *represents* the external world internally; thus it is an easy connection to establish that the vehicles of computation are mental representations of some kind. When RTM is combined with CTM, you get the idea that the mind manipulates mental representations during the execution of computations in order to respond to external stimuli. These representations are held to have semantic content; for example, DOG refers to *dogs*. In other words, cognition is the manipulation of meaningful representations that map onto the world. Being an integral part of the tetrafurcation as well as an assumption of most researchers, I find RTM uncontroversial and will simply assume it to be the case.

What follows from RTM, CTM and the tetrafurcation is that our mind represents the world via our experience and produces its thought through performing computations on those representations. Because we have a limited ability to experience the world—our five senses—it is easy to appreciate the idea that the world is substantially more complicated than we can detect, and our scientific research surely supports this claim. Likewise, the way in which our experiences of the chaotic events of the world are ordered in thought gives use reason to suppose our experience is more complicated than what we can think about it; I think investigations into the brain support this idea with the various subconscious activities of the mind that have been discovered. Furthermore, it seems reasonable to me to suppose our expressible thoughts are probably relatively impoverished relative to the total number of thoughts we have.

Fodor's Linguistic Turn

The *language of thought hypothesis*, or LTH, is a further extension of CTM. Jerry Fodor first proposed his form of this hypothesis in his book *The Language of Thought*, and later discussed in his book *LOT 2: the Language of Thought Revisited*.[18] This view assumes CTM and RTM; Fodor proposes that the representations found in the mind have syntactic structure in the same way language has it. The representations are combinations of simpler entities that give the representation its syntactic and semantic content, and the computations performed on these representations are sensitive to (only) the syntax. Because thought exhibits traits that language also exhibits, such as systematicity and productivity, it is supposed to help to understand thought by postulating that the mind thinks in its own language. We then have a way in which thought is systematic, productive *and* meaningful. This is meant to be more than an analogy. LTH is not only claiming thought can be understood as if it were a linguistic system; it is the claim that thought itself *is* linguistic.

Compositionality. Fodor (2008) claims systematicity and productivity come from the *compositionality of thought*. What would allow thought to be systematic and productive is to propose it is a combinatorial system in which representations in thought are composed of simpler representations according to rules, just like linguistic representations. If philosophers are correct in asserting RTM, it is necessary for the theory of mental representations to account for all the phenomena that are observed in cognition, which then makes it necessary that the theory explain compositionality. The representations of language exhibit compositionality, making it easy to jump from RTM to LTH.

Fodor thinks that the demands of compositionality make it hard to *not* adopt LTH. While this claim is tempting, I resist it. I believe proponents of LTH reverse the order of things. Thought is not compositional because it is linguistic; rather, language has the features it does because of two things: the way in which we experience the world and our need to communicate our thoughts

[18] There are versions of LTH that precede Fodor's discussion. The earliest I encountered is traceable to William of Ockham; this is mentioned by Kenny (2010).

about our experience. The complexity of the world leads to the complexity of our experience, which then feeds the complexity of our thought. Since all of this is prior to communication, I see language as the attempt to externalize our internal understanding after that understanding has been acquired. We need to communicate our thoughts reliably and in a way that another individual can experience, which means our communication must be of a structure conducive to sensory perception and translation into thought. Going back to Abelard, it can be seen that conventions that establish a correspondence between labels and the world are necessary. Conventions allow for communication; by appealing to the world with which all interlocutors are acquainted, we can express our thoughts in a reliable way, given we carry out any dialogue necessary to make conventions salient.

In a note in *LOT 2*, Fodor mentions the possibility of what I suppose is actuality. He says:

> One can imagine the view according to which *only* thought is compositional in the first instance and the apparent productivity, systematicity, etc. of language is parasitic on that of the thoughts they are used to express. In fact, I am inclined to think that is the right view. (p. 55)

This is a consequence of embracing the tetrafurcation of reality; first, the world exists, then we sequentially experience it, think about it, and express our thoughts about it. Each transition involves a reduction, and *abstraction*, of information; the world is surely richer than our experience leads on, and whatever relation holds between thought and experience, it painfully obvious that our expression of thought is relatively bankrupt compared to thought itself.[19] This is a prima facie reason to doubt that thought is *merely* linguistic. Our linguistic expressions seem unable to express the complexity of our thoughts, just as our thought most likely fails to capture the richness of our experience and our experience fails to capture the totality of reality. Anyone who has been at a loss for words or has struggled to formulate a unique thought should have an intuition that our thought must be structured linguistically *after the thoughts have occurred*.

[19] I hope the deficiencies of my expressive ability in communicating my thoughts are not peculiar to me but are a common to all individuals who attempt to express their thought.

I must qualify the preceding comments. First, communication has been a part of humanity for a *long time*—probably longer that humanity has been distinguishable from our ancestors. It is obvious that our brains have large regions devoted to language comprehension and production, and the way in which our young are submersed in a community of language users leads to a complicated picture of the role of language in thought. It is also clear that we develop concepts of things that would never be thought if not for our use of language. Language is so intertwined with thought that it is surely impossible in practice to separate the two. Suffice it to say that thought can in principle exist without natural language or conventions of communication; even if they lack a reason or the means to communicate, it is reasonable to suppose people (and, therefore, other animals and maybe computers, too) can think. Thus, whatever features language has, it has those features because our expression of thought requires it. This is not to say that thought is linguistic, but, rather, that language is an attempt to express our thought. Thought *at least* has the features of language, and surely additional features, too.

Referentialism. Fodor (2008) claims that the compositionality of thought demands one to maintain *referentialism* as well as the view that thoughts have constituent structure. This is the claim that the content of a representation is completely determined by the referent of the representation; the content of representations compose because the constituents of a representation compose, which determines the reference of the complex representation. Since we are assuming CTM and RTM, it is asserted that computations are sensitive to the semantic content of mental representations. Fodor holds that the alternative way to account for semantic sensitivity (two-factor semantics, such as inferential-role semantics) is unable to account for the compositionality of representations, so referentialism is to be preferred.

Fodor analyzes (and eliminates) the possibility that something else besides reference is needed to determine the content of a representation; he claims that the *senses* which are posited by an inferential-role semantics (to explain cases where co-referring terms are not substitutable) are unnecessary for determining the reference of these terms. The syntax of complex

representations is enough to determine the difference between supposedly co-referring terms like those found in Frege cases, so referentialism is sufficient for concept individuation. I maintain that Fodor's arguments for referentialism are independent of his adherence to LTH and stem instead from his commitment to RTM, and so I adopt referentialism while remaining uncommitted to LTH.

Conceptual atomism. Fodor (2008) thinks compositionality requires *conceptual atomism* to be adopted. Atomism is the view that (most) concepts are *atomic*, meaning that their possession does not require one to possess any other concept. If CONCEPT-X is said to have another concept as a constituent, the identity of CONCEPT-X would be (partially) constituted by its relation to this other concept. Conceptual atomists believe complex representations are individuated by their constituent concepts, which are themselves atomic—or, if not, eventually we bottom out the atomic constituents. In light of the tetrafurcation, it seems as though this "bottoming out" will be a fundamental appeal to experience; our thoughts ultimately reduce to information resulting from our experience of objects in the world, and this information is not conceptualized in any way. It appears as though conceptual atomism simply turns into a re-emphasis on referentialism, since our thoughts are, at the lowest level, reflections on our experiences, and our experiences are of the world.

To take stock: CTM provides a generic account of what the mind does in order to produce thought: processing of discrete entities in accordance with rules that govern the manipulation of those entities. RTM provides the entities which undergoes computations: mental representations that correspond to things in the world and have semantic content. Fodor's discussion of LTH provides us with two other ideas: conceptual atomism (atomic concepts are individuated independent of all other concepts) and referentialism (atomic conceptual content is determined by reference). This is all consistent with the tetrafurcation of reality and the imbedded theory of abstraction, and I say these points are independent of Fodor's commitment to LTH.

Since LTH is an empirical claim, I will let the data speak for itself and abstain from any further judgments about it.

Structures of Concepts

Adopting CTM and RTM puts constraints on what concepts can be and what kind of information they contain. Concepts function in a variety of ways within our thought, and any proposed conceptual structure must be able to explain all the functions demanded of them. Machery (2009) seems to think abstract conceptual organizations that philosophers discuss cannot account for the heterogeneity of conceptual function, but I hope to show that the conceptual scheme presented above does allow for heterogeneity.

Associationism. The claim of conceptual atomism is motivated by a distinction between the different possible types of conceptual structure. Suppose concepts have structure. This structure can either be one of *containment* or one of *inference*; either a concept can be said to literally contain other concepts as its parts, such as BIRD containing FEATHER, or a concept can be related to other concepts via inferential dispositions, such as RED being inferentially connected to COLOR. The first instance implies that to represent BIRD necessitates representation of FEATHER; the second type of structure holds that while RED implies COLOR, one can represent something as red without needing to represent it as colored (Laurence and Margolis, 1999, Chapter 1). Both of these types of structure are essentially ways of identifying concepts by appealing to their relations to other concepts. One could argue that inferential dispositions and constituents are both aspects of an *associationist* account of conceptual content. It is intuitively uncontroversial to say that some representations are associated with others—DOG associated with CAT, FISH with GILLS, and so on. In some cases it appears this associative content is required to perform computation. Associationists hold that the strength of the association determines whether a certain associate of a concept is used to individuate a concept during computation. A constituent of a concept would be a strong association, implying that the complex concept cannot be represented without also representing the constituent. Other

associates, such as inferentially associated concepts, are not necessarily represented along with the original concept.

The problem for associationist accounts of conceptual content arises when trying to determine the strength of these associations; it is hard to see the difference among associative bonds, a concern that results in what Weiskopf (2009) discusses as the *principled basis argument* (PBA). There is no principled way to determine the boundaries between relevant and irrelevant associations, so it seems *all* associative content must be used. Holism is achieved, which is not a good situation since it makes it very difficult, if not impossible, to share concepts with other individuals. This argument from Weiskopf is meant to motivate one to adopt conceptual atomism.

Weiskopf (2009) proposes adopting *localism* in order to reply to PBA. He claims that many models that posit complex structure within concepts are well-supported by psychological evidence. As Fodor points out in *LOT 2*, computations (the way a LTH theorist understands them) are by definition local processes, being sensitive to (and only to) the syntax of representations. Weiskopf uses this as a way to support the cognitive models of concepts found in psychology; he claims that this locality is what gives the models positing complex structure their effectiveness. The models to which Weiskopf is referring define the complex structure in different ways, but each has a *local* description of semantic content. Before jumping into these models, I will explore the previously favored theory of concepts—concepts as definitions.

Concept-as-definition. What is sometimes referred to as the *classical theory* of concepts—concepts being equal to definitions, or being "constituted by their defining inferences" (Fodor, 2008, p. 25)—was a widely accepted theory of concepts at one point in time. The classical theory is something like this. Assume that some concepts can be defined; there must be some primitive set of concepts which cannot be decomposed that serve as a pool of terms that can be used to define a complex (non-primitive) concept. This view implies that complex concepts have constituent concepts. Complex concepts, being individuated by the relations held among their constituents, are said to allow for analytic inferences to be made (Ibid., p. 29). This is a

strong point of the classical theory of concepts; the definitional account of concepts is beneficial because it allows one to explain analyticity.

There are other benefits one gains by supporting the classical theory of concepts, but there are also many challenges this theory faces: enough challenges to compel some to look for a more accurate way to describe concepts. As stated before, CTM entails that computations (involving complex concepts) consist in manipulating the constituent (syntactic) structure of complex concepts. Computations are thus local events; holding this assumption, there are three popular theories of concepts in psychology. These are the *prototype*, *exemplar*, and *theory* paradigms. These three paradigms are often held to be running in opposition to one another; some think only one of these can be the "correct" way we use concepts. On the contrary, psychological experiments seem to support the existence of categorization processes that use each of the three versions of concepts (Machery, 2009). In order to see if these three views can be integrated, we must first expound to what each view amounts.

Concept-as-prototype. The prototype theory of concepts claims that concepts are bundles of stereotypical properties that are usually possessed by members in a given class of objects; these properties are sometimes referred to as "statistical bodies of knowledge" (Machery, 2009, p. 83) in order to emphasize that these traits are not *all* necessary properties, but rather the correct number of co-instantiated properties provides sufficient conditions for class membership. A prototype for PLANT would include such things as GREEN, GROWS IN SOIL, PHOTOSYNTHESIZES, IMMOBILE, and PRODUCES SEEDS, all weighted appropriately.

Measures to determine *typicality effects* are often cited as support for the prototype paradigm. Typicality effects are the phenomena that result when subjects are asked to determine how well a certain item fits into a given class; prototype theorists believe that the degree to which members of a certain class fit the prototype can vary from member to member, so categorization that involves prototypes are similarity assessments. The various degrees of similarity implies that the rate at and accuracy with which we categorize typical members will be different from the rate

of atypical members; the experimental results are consistent with this claim (we are faster and more accurate at classifying typical items compared to atypical items); also, when given examples of a class of objects, then prompted with new items to be categorized, the more typical new items are classified more quickly than less typical ones (Machery, 2009, p. 163-168), which is taken to be evidence that prototypes are extracted from the exemplars of the class, thus making items with traits similar to the statistically prominent traits of exemplars easier to match to the class. Typicality effects observed in certain situations lead me to believe this is one valid way in which we conceptualize our world. Comparing a prototype's properties and a new object's properties to categorize the new object is an important aspect of concept usage explained by this paradigm.

Concept-as-exemplar. Exemplar theory is similar to prototype theory in that it too proposes that categorization is based on similarity, though it is between a putative member of a class and the representations of members of the class rather than a prototype of the class. For example, I may have stored in my memory a representation of the particular 1967 Ford Mustang that my father used to own. The way proponents of the exemplar paradigm tell the story of categorization is that when I come across a car I have never seen before, I compare the representation of this new object to the representation of the exemplars I have in my memory and determine how similar they are; I look at such things as the grill of the car, the unique grooves in the hood and fenders, the headlight position, etc, until I am able to determine if the new object is similar (enough) to the exemplar 1967 Ford Mustang I have in my memory (or some other exemplar I have stored—I potentially compare the new object to *every* exemplar that has some similarity to the new object).

Evidence for exemplar theory consists of two phenomena; the *old-items advantage* and the discrepancy between prototype-predicted results and actual results of categorization. Old-items advantage is that items that are "older" (or more familiar) are processed more quickly than new (less familiar) items, regardless of typicality. The second phenomenon taken as evidence for the exemplar paradigm is that sometimes less typical members are categorized more quickly than

more typical members, even though the prototype paradigm predicts the opposite; familiarity also plays a role in these cases (having a pet dog with three legs makes a three-legged dog quicker to be categorized than a more typical four-legged dog) (Machery, 2009, pp. 173-178). While there are, of course, challenges to this evidence, I believe there is enough support for the idea that we sometimes store representations of particular objects classified as an instance of a class in our mind for categorization purposes. Familiarity is an important idea for which exemplar theory accounts.

Concept-as-theory. Following the development of the prototype and exemplar paradigms, a third paradigm arose. This asserts that concepts are *theories*, in the sense that they are bodies of knowledge that are able to explain what (and, importantly, *why*) properties are possessed by members in a class of objects. This is a vague idea, and it is hard to tell what exactly psychologists who support this paradigm mean by "theory". There is one fundamental difference between the theory paradigm and prototype/exemplar paradigms: categorization based on prototypes and exemplars rely on measuring *similarity* between a new object and old information, while categorization based on theories relies on determining whether a new object satisfies causal/functional roles. Machery (2009) says, "[p]yschologists assume that laws, causal propositions, functional propositions (for instance, the propositions that birds have wings to fly), and generic propositions (for instance, the proposition that dogs bark) explain why things happen" (p. 101), and these types of things can be accommodated by the theory paradigm.

The support for the theory paradigm is mostly in the form of showing how similarity (the main force behind prototype and exemplar theories) is sometimes *not* used to categorize objects, but rather some kind of rule or law relating properties of a thing to category membership is used. Lance Rips designed the pizza experiment with the intention of demonstrating that categorization sometimes does not rely on the similarity judgments (an integral part of both prototype and exemplar paradigms) but rather rely on appeals to rules about category membership. Subjects were told to imagine an object larger than a quarter *and* smaller than the smallest pizza they had

ever seen, then they were asked to (a) tell whether the object was more similar to either a quarter or a pizza, and (b) whether the object was more likely to be a quarter or pizza. While the intermediate size object was randomly determined to be more similar to either a quarter or a pizza, most subjects determined that the object was more likely to be a *pizza*. Rips concluded that this effect was the result of subjects appealing to the rules which define members of the quarter category, namely necessarily being a certain size. A further experiment was performed by Smith and Sloman; half of the subjects were told that the intermediate size object was also silver-colored, while the other half were only told the size. None of the subjects were instructed to talk out loud. The results were that most of the half told that the object was silver stated that the object was both more similar to and more likely to be a quarter, while the other half gave random answers to both questions. These studies provide evidence that at least sometimes people use rules about category membership to make categorization judgments, which is a phenomenon that should be explained.[20]

Doing without Entities beyond Necessity

Having evidence that all three proposed paradigms of concept are employed to make categorization judgments depending on the circumstances, Machery (2009) concludes that the term "concept" does not designate a homogeneous class of entities. He thinks "concept" should no longer be used by psychologists (and should be replaced with "prototypes", "exemplars", and "theories"). I disagree with Machery; while I *do* think conceptual information is used in different ways depending on the demands placed on an individual, I think what "concept" refers to is similar to what Machery assumes CONCEPT to be, which, when properly modified, can be used to unite these three paradigms of concepts.

Recall the description of concepts that Machery (2009) gives us (page 6). To begin, I agree that concepts are bodies of knowledge, if by "knowledge" it is meant the information drawn

[20] These studies are reviewed by Machery (2009): see pages 183-187.

from experiences of the external world.[21] Also, I would agree that this information must be stored in our memory, though I will leave the distinction between long-term memory and other forms of memory to be made by those who care to distinguish these. That these bodies of information are used in higher competences is something which I find uncontroversial; I think it is obvious that our claims as to what an object is and what properties it possesses are important information that is used when thinking about said object. The last phrase, "when these processes result in judgments about *x*", seems that it was added on to make this description apply more specifically to the way in which psychologists investigate concepts, and it seems unnecessary in order to grasp what is meant by "concept". I will amend his description by first removing this restriction.

"Used by default" is a problematic qualifier. I believe this is part of what makes Machery (2009) able to claim that concepts are not homogeneous entities. It is obvious from the research he presents to support this claim that the information used to make categorization judgments depends on the particular context in which the judgment must be made; since different circumstances cause people to use different information to categorize, it would seem the understanding of concepts Machery presents would cause one to get different answers when asked what body of information is used to individuate a concept. In one situation I may use a prototype of SEDAN to categorize a new model of automobile, while in another situation I may compare the new auto to a particular instance of SEDAN I have stored in my memory, and still in another situation I may use rules that determine what it means to be a sedan in order to classify the new automobile. Does this mean I have three different concepts of sedans? Machery says, "Yes!", but I think this is a hasty judgment; all three bodies of information could be said to individuate SEDAN, and the fact that an exemplar, a prototype, and a theory of SEDAN are all called concepts of sedans gives me reason to think there is some sort of unity within these entities. What if these paradigms are three bodies of information extracted from a larger body of information that is the *generic concept* SEDAN, yet due to the particular circumstances I may

[21] Machery makes note of this (p. 8).

only need to retrieve a limited set of information to perform the necessary processes? I suggest these three seemingly different concepts, which correspond to prototype, exemplar, and theory, are really three ways of processing the same thing: three partially overlapping bodies of information that all refer to the same class of entities.

Concept Functionalism

Assuming RTM, the mind represents the external world internally. It seems fitting to think the relations we express in spoken words should be represented in the mind, and it also seems safe to conclude that the structure that our thoughts represent is the structure of the world in our experience. Any theory of concepts that cannot account for all the different ways in which our mind forms and uses concepts is an unsatisfactory theory. What is immediately obvious to me is that there is a certain functional role that concepts must play in thought, and that role is that of serving as a representation of the external world that is accurate in the ways relevant to the problem being solved. Implicit in this entire discussion is that we are communicating our thoughts. Since, ex hypothesi, the mind solves problems through computation, the mental representations of the world must be computable and easily communicable. Concepts, being those mental representations, must be vehicles that can represent the world as it is, have computations performed over them, and be easily communicated.

The tetrafurcation of reality contributes a way to understand how concepts acquire their structure: language is a means of communicating thought, our thought is about our experience of the world, and the world is a state of affairs that is there to be regularly observed. Any structure we observe and recognize in the world is reflected in the structure of our thought (of course, after being shoved through the bottleneck that is our capacity for sensory experience), and our language is structured so as to facilitate the communication of the structure of our thought. Since we can discuss any particular object in at least the three distinct ways that Machery presents, it seems appropriate to say that there is a generic concept we can think about in different ways.

Motivating Concept Functionalism

I would like to show how such a structure can arise. This sketch embraces the notion of abstraction which Avicenna develops, i.e., the stepwise abstraction from experience of objects. Assume you are a person who has never encountered birds, and don't know what it means to be a bird. You see a cardinal fly by (though you don't know it is a cardinal). You store a representation of this object, which would presumably consist of such things as how it moved, where it went, what it looked like, any noises or smells perceived, and also the time and place you saw it. You now possess an exemplar of an unknown category. The only way to refer to this object would be to say something like, "that creature I saw on such-and-such day at such-and-such place," or "that thing which flies and is red and went in a tree". Suppose the next day you see another cardinal, and notice it had mostly the same characteristics. Imagine that the following day, you see a robin and a finch sitting in a tree and then watch them fly away; you might note that while these have some shared characteristics to the other creatures you saw, they are not identical. You are now in a position to extract a prototype; after storing several exemplars in your memory, you are able to pull out common characteristics among these distinct yet similar representations. Shared characteristics between similar exemplars are weighted based on their frequency to produce a prototype.

Later, let's say a friend describes an object he just saw; it was orange, had wings, and flew into a nest in a tree. You are able to take this description and attempt to fit it to your prototype, and upon doing so you judge that this object is the same type of thing you saw several times before. You are now capable of formulating a theory. There exist creatures that vary in color and that fly in order to get from place to place, specifically trees. Presumably, the more acquainted with exemplars of a concept you become, the more refined your prototype becomes, thus the more obsolete using exemplars to categorize becomes. Eventually the prototype is refined enough to become a theory; this process involves becoming more acquainted with a certain concept, most likely through discussion with others, so a person's level of acquaintance

(partly) determines which method of categorization one uses. No one would say your theory of BIRD is not about birds, and it seems painfully straightforward to say any exemplary information you use to develop the prototype or theory is just as much a part of your understanding of BIRD as the prototypical or theoretical information.

Ecological Validity

One consequence of adopting this view is that any given concept can be considered in various ways, so the way in which we learn concepts becomes a crucial aspect to be considered. Machery (2009) notes that there certainly are discrepancies between experimentally confirmed concept acquisition processes and real-world concept acquisition processes. He says, "…the ecological validity of most concept-learning experiments remains unclear" (p. 162). What I think must be added to my suggestion of conceptual hierarchy is that we rarely acquire concepts in the proposed exemplar-prototype-theory sequence. We often learn things in a muddled way: maybe we see a token of a type, are instructed that it is labeled in a particular way, and are then given various reasons why it is a token of the indicated type. The idea is that we eventually accumulate enough experience to confirm and revise the partial prototype-theory we are initially given, and we can always revert back to less theoretical (i.e., more concrete, less abstract) exemplars.

What I find wrong with the claims Machery (2009) makes is that they seem to take the experimental evidence of the various ways of categorizing as evidence that there are distinct bodies of information in the mind. I suggest this is an effect of the experimental setup; the capacity to perform these various tasks is does not establish that the information to which we appeal while solving the distinct problems is neatly divided into exemplary, prototypical, and theoretical information. The psychological studies are designed to detect specific capacities, and the design is such that it elicits only the desired responses. Since abstract universal terms exist, and since the same term is used to identify the exemplar, prototype, and theory of a class of objects, there seems to be a lack of any sort of strong division within the information. Our ability

to adjust the information we are using during categorization tasks is not evidence that the information is disconnected.

Conclusion

The philosophical notion of ‘concept’ is an abstract notion that serves as the genus of all the particular concepts we form. Because of this generic usage, the need to differentiate among the various types of concepts is obvious; this in no way undermines the usefulness of the generic usage, but it does put emphasis on understanding the distinct ways in which we use concepts in thought. The psychological investigation of concepts is surely valuable, and there are definitely differences between the aims of philosophers and the aims of psychologists when investigating concepts. Clarifying what “concept” is supposed to mean helps to show that philosophers and psychologists may be talking past each other when they disagree about what concepts are. If I have fulfilled my aims, I have shown that there is common ground between these two research programs, and we can hopefully move past polemics and move towards synergism.

References

Armstrong, D. (1989). *Universals: An opinionated introduction*. Boulder: Westview Press.

Armstrong, D. (2010). *Sketch of a Systematic Metaphysics*. New York: Oxford University Press.

Barnes, J., Schofield, M. & Sorabji, R. (1975). *Articles on Aristotle*. London: Duckworth.

Cleary, J. (1985). On the Terminology of 'Abstraction' in Aristotle. *Phronesis* **30**, 13-45.

Fodor, J. (1998). *Concepts*. New York: Oxford Univesity Press.

Fodor, J. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Kenny, A. (2010). *A New History of Western Philosophy*. New York: Oxford University Press.

Honderich, T. (Ed.). (2005). *The Oxford Guide to Philosophy* (2nd ed.). New York: Oxford University Press.

Loux, M. J. (2002). *Metaphysics: A contemporary introduction* (2nd ed.). New York: Routledge.

Machery, E. (2009). *Doing without Concepts*. Oxford : Oxford University Press.

Margolis, E., & Laurence S. (1999). *Concepts: Core Readings*. Cambridge: MIT Press.

McGinnis, G. (2007). Making Abstraction Less Abstract. *Proceedings of the ACPA* **80**, 171-183.

Piccinini, G. (2004). The First Computational Theory of Mind and Brain: A close look at McCulloch and Pitt's "Logical Calculus of Ideas Immanent in Nervous Activity". *Synthese* **141**, 175-215.

Piccinini, G. & Bahar, S. (2012). Neural Computation and the Computational Theory of Cognition. *Cognitive Science*, 1-36.

Piccinini, G. & Scarantino A. (2010). Computation vs. Information Processing. *Studies in the History of Philosophy of Science* **41**, 237-246.

Piccinini, G. & Scarantino, A. (2011). Information Processing, Computation, and Cognition. *Journal of Biological Physics* **37**, 1-38.

Piccinini, G. & Scott, S. (2006). Splitting Concepts. *Philosophy of Science* **73**, 390-409.

Rodriguez-Pereyra, G. (2000). What is the Problem of Universals?. *Mind* **109**, 255-273.

Ross, D. (1995). *Aristotle* (6th ed.). New York: Routledge.

Russell, B. (1945). *A History of Western Philosophy*. New York: Touchstone.

Scruton, R. (2002). *A Short History of Modern Philosophy* (2nd ed.). New York: Routledge.

Spade, P. V. (1994). *Five Texts on the Medieval Problem of Universals: Porphyry, Boethius, Abelard, Duns Scotus, Ockham*. Indianapolis: Hackett Publishing Company, Inc.

Weiskopf, D. (2009). Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research* **79**, 130-162.