

CONCEPTS

Molly Josephson

A thesis submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Master of Arts in the Department of Philosophy in the College of Arts and Sciences.

Chapel Hill
2014

Approved by:

Dorit Bar-On

William G. Lycan

Alan Nelson

© 2014
Molly Josephson
ALL RIGHTS RESERVED

ABSTRACT

Molly Josephson: Concepts
(Under the direction of Dorit Bar-On and William Lycan)

Most contemporary theories of concepts hold that the principal function of these mental entities is that of categorizing things in the world. Such theories, however, tend to generate flawed or insufficient accounts of several cognitive phenomena. I argue for a theory which takes a different account of the nature of the mind. On minimalism, concepts are primarily relational entities, and to think conceptually is to draw relations between mental representations. The minimalist picture can easily explain numerous cognitive phenomena that other theories of concepts overlook; it can also make sense of disagreements between those theories regarding what a concept really is.

TABLE OF CONTENTS

List of Figures	v
I. Introduction.....	1
II. Systematization and Recombinability.....	3
III. Minimal Concepts.....	6
A. A More Basic Concept.....	9
B. Increasing Conceptual Complexity.....	11
C. On the Notion of Conceptual Content	17
D. Debunking the Notion of Conceptual Parts	23
IV. Abstraction and Colloquial “having the concept of”	26
V. A Critique of Sellarsian Concepts.....	31
VI. A Response to Sellars	44
VII. Concluding Remarks	46
Works Cited	48

LIST OF FIGURES

Figure

1.....	15
2.....	18
3.....	25
4.....	30
5.....	46

I. Introduction

Concept minimalism is a straightforwardly simplistic view which holds that concepts are mental representations that are systematic and recombinable. According to Elisabeth Camp's (2009) version of minimalism, they are (*a*) systematically related to what in the world they represent and to other concepts and (*b*) recombinable with one another to form other concepts (303). Systematization and recombination are supposed to be cognitive process(es), performable by whatever mental faculties a sufficiently cognitively advanced creature might possess; and concepts are thus defined as the entities to which these processes apply. Concepts may be much more than systematic and recombinable mental representations, as most philosophers maintain, but this is what they are at the very least.

Other views demand more of concepts, and they do so in various ways. For instance, Sellars (1981) claims that concepts are necessarily linguistic, Davidson (1975) claims that conceptual thought requires the concept of belief, McDowell (1996) claims that perception is conceptual, and Millikan (2005) claims that concepts are object-tracking. All of these views seem to disagree with one another, but it is not clear exactly how they disagree; it is quite difficult to see, for instance, how a Davidsonian concept compares to a Sellarsian one. These and other "intellectualist" views, as Camp calls them, seem to hold that concepts are something entirely different from what the minimalist claims. But, as I will demonstrate, the minimalist definition is answerable to the requirements of these intellectualist views and others in virtue of the simplicity of its two key criteria. And thus a developed account of minimalism will reveal that the differences between these views of concepts are merely variations in degree of

complexity. The more highly intellectualist the view, the greater the amount of systematization and recombination it requires of concepts.

So, based on the Sellarsian claim that conceptual thought requires the use of language, from which it follows that all concepts are linguistic, the minimalist regards such concepts as minimal mental entities that have been made more complex through the processes of systematic relation to and recombination with the concepts of language. And based on Davidson's claim that conceptual thought requires the concept of belief, the minimalist regards Davidsonian concepts as mental entities that have been systematically related to and recombined with the concept of belief. On the minimalist view, then, language (viz. employing linguistic structure in thought) and having the concept of belief (viz. knowing of the possibility that one is mistaken) are ways of making certain mental connections between ideas. Each time one links a thought to language or to an epistemic position, this increases the systematicity and recombining ability of the initial thought. Thus, the complexity of a concept is just the extent to which it is systematically related to other concepts or has undergone recombination with them. And the more highly complex one's concepts are, the more highly intellectual are the thoughts one is capable of entertaining.

Taking account of all such intellectualist views in this way, the minimalist can effectively draw up a continuous spectrum of conceptual complexity, on which to locate each view's "concepts" with respect to the others. As we shall see, such a spectrum would run from basic to complex, according to the amount of systematicity and recombining ability of concepts, where a "minimal concept" is the lowest possible degree of complexity and each of the views noted above—among others—locates its concepts some distance up the spectrum. The advantage of the minimalist view is that, not only can we finally make sense of all these other views on concepts,

but also, and more importantly, in positing concepts and conceptual abilities all the way down, we can use these views to derive a comprehensive and highly explanatory model of cognition.

Views like that of Sellars, as well as some other concept-intellectualists, often initially look intuitive because of a natural tendency to identify concepts with words. For instance, it is customary to talk about “the concept of justice” in philosophy or elsewhere by explicating the meaning, connotations, uses, etc. of the very word, “justice.” But if we instead allow our philosophical term for mental entities—“concept”—to track the systematic, recombinable nature of these things, then, rather than debating whether things qualify as conceptual or not, we can begin to discuss what thought itself is and how it is possible at any level. It is in this way, the minimalist view of concepts becomes a highly comprehensive account of how we think.

In this paper, I will explain the minimalist view through a cognitive story and by way of contrast from the (general) intellectualist. Ultimately, the minimalist picture will provide a new philosophical approach to the mind, one that works from the ground up and, in doing so, exposes what other views seem to be neglecting. The minimalist holds that, while the other theories on concepts are not wrong *per se*, they provide unhelpful distinctions and faulty implications about the mind as a whole and therefore cannot be used as models of cognition. Minimalism itself can provide such a model, *while* still accommodating the philosophical claims made on other views. Maintaining a minimalist view of concepts allows us to explain an immense number of cognitive phenomena and philosophical problems about the mind.

II. Systematicity and Recombinability

The minimalist equates a concept with a systematic, recombinable mental representation. Just how much is entailed in this definition may not be immediately obvious. We may first note that the term “mental representation” here can be taken to have the same denotation it does in

other literature on concepts. The terms “systematicity” and “recombinability” are the key ones in the present discussion, and they may seem relatively uncharted territory in the rest of the philosophical literature on concepts. The minimalist does not intend for them to carry any metaphysical weight; so when we claim that an entity is systematic and recombinable, we are merely claiming that it is related to other entities in a nonspecific way. What the minimalist means by these terms can best be understood through a brief example of their use.

The minimalist would claim that the concept of a tree is *systematically* related to (a) real-world tree(s) as, for instance, in appearance. When one conceptualizes a tree, one often has in one’s mind a mental image that to some extent resembles (a) real-world tree(s). The mental image may be stylized, as a cartoon tree; it may be in black-and-white, if the thinker is colorblind; it may be upside-down; it may not include roots or leaves; etc. In spite of this variability, it stands to reason that there is something—perhaps mere “tree-ness”—about the *thought* of a tree that corresponds to (an) *actual* tree(s), and that correspondence in this case is one of resemblance. This is what is meant by a concept bearing systematic relation(s) to something in the world.

Secondly, a concept is *systematically* related to other concepts. The concept of tree can be related to the concept of bush (because both trees and bushes have woody branches), to the concept of plant, to the concept of photosynthesis, etc. The systematic nature of these relations between concepts consists in (respectively) compositional and structural similarity, the fact that a tree is a type of plant, and photosynthesis being the process by which the tree lives. The structural relations between things in the real, physical world, such as between a tree and a bush, are supposed to be reflected in the relationship between the concepts of those things in a thinker’s mind (Camp₂). And while no particular one of these systematic relations to the <tree>-

concept is essential for it to be a concept of a tree (a child can have the concept ⟨tree⟩ without knowing what photosynthesis is), they do determine what that person's thoughts about trees will consist in.¹

Thirdly, a concept is *recombinable* with another concept according to the systematic relation between them. So, for instance, when two concepts are recombined, they form a third, ostensibly more complex concept; ⟨tree⟩ and ⟨bush⟩, based on the systematic relation just noted, may be combined under and thus comprise the more complex ⟨woody branched things⟩-concept. Similarly (as this may be a more intuitive example), the concept of tree and the concept of family, in virtue of the common aspects of branching, allow for ⟨tree⟩ and ⟨family⟩ to combine and form the concept ⟨family tree⟩, which is the concept of a visual apparatus that describes familial relations with a series of branching lines. And, as before, the relations from ⟨family tree⟩-concept and to any other concept to which it is related reflect the systematic relations between family trees and other real-world things.

We can think of a concept's *recombinability*—its ability to be recombined—as the extent of its systematization, since a recombination of concepts is based on the systematic relation(s) between them. That is, a concept can only be recombined with other concepts to which it bears some relation—whether that conceptual relation is a reflection of the real-world correspondents' relationship or it is a product of the thinker's drawing such a connection in her own mind. ⟨Tree⟩ can only be recombined with other concepts to which it is related for that thinker; so if one conceives the relationship between real-world trees and bushes, one can combine these concepts (as above) and develop a more complex concept for ⟨woody branched things⟩, which covers trees and bushes. If a thinker does not know what photosynthesis is or does not know that trees

¹ Note that small angled brackets (⟨⟩) to refer to the concept ⟨tree⟩, and this is very different (in my paper) from the unadorned **tree**, which refers to a real-world entity. (Quotation marks on “tree” denote verbal utterance of the word.)

perform it, however, she obviously cannot combine a <photosynthesis>-concept with her <tree>-concept because she does not possess the former. In this way, we can think of the two conditions, systematicity and recombability (or, in other words, the two processes of systematization and recombination) as essentially denoting the same thing: a relationship; systematicity focuses on the existence of the relations, and recombability focuses on their potential for use. Both are effectively describing *how that concept is deployed* in thought.

The minimalist's conditions are analogous to the claim, which hopefully sounds intuitive, that *a concept is a way of thinking*. Thinking of a tree as a relative of a bush is one way of thinking about a tree—that is, one way of deploying the <tree>-concept. Conceiving a mental image of a tree is a way of thinking about a tree. A tree can also be thought of as a plant, or as a photosynthesizer, or as a thing that shares some structural element in common with families. Each of these ways of thinking about (a) tree(s) just is the making use of two or more of the systematic relations that hold of the <tree>-concept.

III. Minimal Concepts

Consider a person who has no formal knowledge of chemistry, nor any knowledge of the periodic table of the elements. One day, as she is walking down the hall of a building on her college campus, she passes a man named Bill Lycan, who says to her, “Molybdenum,” and continues walking. A moment later, she begins to think about what just transpired. Having heard an unfamiliar word—or maybe it was two words?—or maybe nonsense?—she naturally begins to wonder about this thing and what it could be (setting aside the additional puzzlement she feels from having such a random, unexpected encounter).

Let us take a closer look at these thoughts of hers. She is thinking about what was said and what happened; and so we can say she is mentally representing something. Moreover, this

mental representation is systematically related to (if only in that it is directly caused by) the utterance event, the physical sound, and the physical state of affairs at the time and place of utterance. For instance, her thought (call this thought \langle molybdenum \rangle) reflects the sound she thinks she heard from Bill's mouth: each time she deploys \langle molybdenum \rangle , she represents it as the sound /mə'libdɪnəm/ in her mind. Furthermore, as she entertains the possibility that what she heard may have been a nonsensical utterance and that it may have been a real, meaningful word, it becomes evident that this thought is systematically related to two other concepts already in her cognitive repertoire: her concept of nonsense (\langle nonsense \rangle) and her concept of a meaningful word (\langle word \rangle). It is also related to her \langle Bill Lycan \rangle -concept: she perceived Bill saying it and thinks on this. So, the mental representation bears the requisite systematic relations to something(s) in the world and to other concepts.

Additionally, the recombining ability of her \langle molybdenum \rangle -concept (that is, its potential for being combined with other concepts) is a product of the extent of its systematization. At this point, its recombining ability for the subject is relatively limited because of the scant nature of the original context in which it was deployed. Since "molybdenum" was uttered with little or no contextual situation, the thinker only recombines it with concepts having to do with this relative lack of contextual information. She deploys it in combination with her \langle nonsense \rangle -concept when she entertains the thought that "molybdenum" was a nonsensical utterance. Likewise, to entertain the thought that it could have been an actual, meaningful utterance is to deploy \langle molybdenum \rangle and \langle word \rangle / \langle meaning \rangle in combination; to think of Bill Lycan having said the utterance is to combine \langle molybdenum \rangle and \langle Bill Lycan \rangle ; and so on. So, according to the minimalist's criteria, she has a concept. (This concept is not necessarily a "concept *of* molybdenum" in the everyday

sense, since she has never taken chemistry, hence its designation in brackets.²)

Since our subject is not familiar with the periodic table, nor has she ever otherwise encountered actual molybdenum (whether she knew it or not), when she hears and subsequently thinks about “molybdenum,” this thought does not involve what one might ordinarily think upon hearing the word. If we know what molybdenum is, then our own concept of ⟨molybdenum⟩ looks to have certain attributes (relations) that the subject’s does not (yet) have. For instance, if we know molybdenum has the atomic number 42, then we can represent our concept ⟨molybdenum⟩ by its relation to the ⟨atomic number 42⟩-concept. Similarly, what we know about its being a chemical element, metallic, etc. would be reflected in our ability to deploy the ⟨molybdenum⟩-concept as it relates to our concepts ⟨chemical element⟩, ⟨metal⟩, etc. But our subject, who does not know enough to systematically relate these concepts to one another, does not deploy (her) ⟨molybdenum⟩ that way. When she deploys ⟨molybdenum⟩, she is thinking about the sound, the possibility of its being a nonsense word, puzzlement over why someone would randomly utter “molybdenum,” etc. If she is a particularly proficient English-speaker, though, she may by sheer luck make some of the same conceptual connections we make to our ⟨molybdenum⟩-concept: for instance, she may be capable of discerning the suffix “-um” from what she heard and in turn of relating “molybdenum” to other terms possessing this suffix; some such terms denote chemical elements, and some of these may be slightly familiar to her through everyday experience. If this series of relations is evident to our thinker, then she, like a chemistry-educated thinker, may be able to think about “molybdenum” as related to “aluminum” and to “helium,” for instance. And insofar as she can relate her concept ⟨molybdenum⟩ to her

² We could have designated the concept with any meaningless symbol, but it will eventually become evident why the bracketed word is an appropriate one. For our present purposes, it is just a coincidence that the symbolic representation of the concept looks a little like the word—linguistic representation—for the real-world element molybdenum.

concepts of other chemical elements, her concept shares something in common with that of chemistry-educated thinkers (people we would ordinarily consider to have a sufficient concept of molybdenum).






So, the minimalist's point here is that, even in the early stages of thought about something, when one is merely conjecturing possible meanings of an utterance for example, one's cognitive processes are indeed conceptual. These cognitive processes are effectively efforts to develop systematic relations between her new concept, ⟨molybdenum⟩, and concepts more familiar to her, such as ⟨word⟩ and ⟨nonsense⟩, etc.

A. A More Basic Concept

This is about as close as we can get to talking about a truly minimal concept for ⟨molybdenum⟩: the systematic relations it possesses are few in number and very loose. For the sake of exposition, I have included more systematic relations above than are necessary to minimally satisfy the minimalist's requirements. The definition of minimalism on Page 1 looks to require merely two relations of a concept: one relation to something in the world, and one to another concept. So perhaps we can yet imagine a truly minimal ⟨molybdenum⟩ concept, which is systematically related to only one other concept and one real-world thing. We might posit, then, a minimal-⟨molybdenum⟩-concept which is systematized as follows: ⟨molybdenum⟩ to (the sound) /mə'liɪbdɪnəm/ and ⟨molybdenum⟩ to ⟨nonsense⟩.

Admittedly, given the story just told, it may seem unlikely that a person would ever have such a simplistic concept (of “molybdenum” or of any thing) because, as has been demonstrated, even mere utterance of a word can initiate a long stream of thought about that utterance, allowing for the construction of numerous systematic relations almost immediately upon the subject's hearing it. Nevertheless, to call such an entity ‘conceptual’ is still permissible on minimalism.

And it is perfectly conceivable that a person may *only* deploy a concept this way. (This would have been the case if our subject had deemed the utterance to be nonsense and left it at that; i.e. immediately started thinking about something else.) But, it seems, we rarely have such minimal thoughts even about apparently nonsensical utterances, and anyone who hears “molybdenum” for the first time under such circumstances as our subject did will likely immediately try to relate her newly acquired concept to as many of her more highly developed concepts (i.e. <word>, <aluminum>, etc.) as she can, resulting in that concept’s rapid increase in complexity beyond the minimal.

For the sake of precision, then, what may be a better example of a minimal concept is , the mental image of a tree. This <tree>-concept is a way of thinking about real-world trees by representing them visually. More precisely, this concept is minimal if (and only if) systematically related as follows:  to real-world tree(s), and  to the idea of vague visual resemblance between it and (a) tree(s). (To make this look like an extremely lightweight relation, for present purposes, we can consider the idea of vague visual resemblance to be a very basic concept itself, i.e. the concept <vague visual resemblance>.) Lastly, it satisfies the recombability criterion in virtue of its systematic relation to the <vague visual resemblance>-concept:  can further be combined with the concept of a bush, since bushes may vaguely resemble trees in appearance. So in this sense, mental images, mental sounds, and the like are our best approximation to the least possible entity that satisfies the requirements of the minimalist: it is mental, it is a representation, and it is systematically related and recombable. So, as long as the thinker can deploy a concept akin to <vague visual resemblance>, which, we can imagine, even a toddler can do, we can see that  suffices as a minimal concept of a tree.

B. Increasing Conceptual Complexity

Let us now return to our story about the minimal-⟨molybdenum⟩ thinker. Since the systematic relations that the concept first has are not what we (as people familiar with the term and the element) usually think of when we think/deploy ⟨molybdenum⟩, and because they are fairly loose conceptual relations (to ⟨meaning⟩, ⟨word⟩, and ⟨nonsense⟩, mainly concepts about language), and for reasons yet to be explained, the early ⟨molybdenum⟩-concept is of very little substance. That is, after only hearing “molybdenum” once, one’s deployment of ⟨molybdenum⟩ is limited in scope by the relatively few concepts that one is capable of relating it to. After all, our subject’s thoughts are mainly just musings about whether it even is a word or nonsense or anything, for that matter. With so few conceptual relations, and such loose ones at that, her ⟨molybdenum⟩-concept looks as if it’s waiting to be systematized.


And our subject can ratify her concept ⟨molybdenum⟩ by, for instance, asking Bill Lycan why he said what he did or what it meant, or by typing what she thought she heard into the search bar of Wikipedia. When she does so, she makes use of its assumed referential quality (in other words, its conjectured systematic relations to the concepts of language—⟨word⟩, ⟨meaning⟩, etc.) as a basis for furnishing new systematic relations. As she reads, she will use written language, the communicative medium of the Wikipedia page, to furnish systematic relations between her concepts. Linguistic communication (language) is ideal for systematizing one’s concepts, since, as Camp suggests, words are highly effective vehicles for discrete concepts, and structured linguistic communication requires that words be deployed in combination with one another. So a word, such as “molybdenum,” thought of as a vehicle for an individual concept, ⟨molybdenum⟩, combines with other words, which convey other concepts, to

form a sentence.³

We can derive an alternate version of the minimalist’s principle of recombination from Camp’s considerations about language. In order to distinguish concepts from “undifferentiated thoughts,” she claims, “concepts must be *compositional*: re-deployable not just in isolation but in combination” (2013). An “undifferentiated thought” would be, for example, ⟨molybdenum⟩ without any of its aforementioned systematic relations. So, on a fully minimalist reading of Camp’s claim, in order for this entity to be properly conceptual, everywhere it is deployed it must be deployed alongside other concepts.

Words, such as those that make up the content of a sentence on a Wikipedia page, provide exemplary vehicles for representation in which concepts can be combined with one another, and a reader or hearer can furnish systematic relations between her own concepts deriving from those combinations. The effective conveyance of concepts (mental entities) looks to require a common set of vehicles; and exactly why language provides some of the best is also explained by Camp: “If we assume that stable representational abilities⁴ require underlying representational vehicles, then this entails that concepts’ vehicles must be at least somewhat arbitrary. Language offers a paradigmatic case of semantic arbitrariness . . . [and] syntactic neutrality,” (2013). It is the balance between structure and arbitrariness that makes language a particularly good representational system for deploying concepts together and furnishing systematic relations between them. This is why reading the Wikipedia page on molybdenum and asking Bill about “molybdenum” are ideal actions to take in order to develop one’s ⟨molybdenum⟩-concept.

³ Noteworthy, too, is that the sentence itself as a whole can be thought of as a single complex concept, ostensibly comprised of two or more more basic ones.

⁴ We can read “stable representational abilities” as the ability to have a concept (⟨molybdenum⟩) after hearing the utterance “molybdenum” or as the ability to mentally envisage a tree () after seeing a tree.

So, Wikipedia or Bill combines certain concepts, via their linguistic representations, and the subject apprehends the concepts from those words in combination (and her ability to use language), using each new combination to inform her own concept \langle molybdenum \rangle . So, when a sentence such as “Molybdenum’s atomic number is 42” appears, the reader furnishes a relation between her \langle molybdenum \rangle -concept and her \langle atomic number 42 \rangle -concept that reflects this particular combination. She also gains from her perusal of Wikipedia (or conversation with Bill) new systematic relations between \langle molybdenum \rangle and concepts from chemistry such as \langle metal \rangle , \langle chemistry \rangle , \langle atom \rangle , etc., in virtue of their deployment in combination with one another. In this way, a greater number of co-occurrences of terms is reflected in the mind of the reader (or hearer) as tighter systematic relations between the corresponding concepts. This is how she comes to develop a more complex concept \langle molybdenum \rangle , and eventually a concept of molybdenum (the real-world substance).

\langle Molybdenum \rangle becomes more highly systematized in two ways: she furnishes tighter relations to it and a greater number of them. The tightness of a single systematic relation between two concepts reflects the number of times certain words or ideas appear together, for instance, in the same sentence or paragraph or section. And the number of systematic relations to \langle molybdenum \rangle derives from the number of words or ideas that are associated, e.g., anywhere on the Wikipedia page with “molybdenum.” So each occurrence of the word or utterance “molybdenum” is reflected in a modification of the reader or hearer’s concept \langle molybdenum \rangle , either by the tightening of a relation it already has or by the furnishing of a new one.

Importantly, this picture does not entail that the \langle molybdenum \rangle -concept loses any of its older systematic relations. The furnishing of new systematic relations upon a mental entity does not eliminate its older systematic relations, and this has important implications for the structure

of the mind. The subject's conceptual relations from ⟨molybdenum⟩ to ⟨nonsense⟩ and to ⟨word⟩, etc. still exist in virtue of the original context in which the concept was deployed. The destruction of these relations would mean, for instance, that she had forgotten her earlier experience: if ⟨molybdenum⟩ and ⟨Bill Lycan⟩ are no longer connected, then she can no longer think about Bill Lycan having uttered “molybdenum.” But since she retains this ability (nothing about the new context causes her to forget her original experience), the respective concepts retain their original systematic connection.

Although nothing on Wikipedia causes her concept to completely lose its older systematic relations, they do look to become less relevant to the thinker as she learns more about molybdenum. Relative to its connections to ⟨chemical⟩, ⟨metal⟩, and to all the other concepts deployed by Wikipedia, ⟨molybdenum⟩'s connections to ⟨Bill Lycan⟩, ⟨nonsense⟩, ⟨word⟩, etc. look to become looser, less relevant by comparison. Since the subject generally trusts Wikipedia as a good source of information and does so more strongly than she trusts her own intuitions, and since the Wikipedia page makes no mention of Bill Lycan or any of the other things she had earlier associated with the concept (though Wikipedia does give pronunciations), she takes her concept ⟨molybdenum⟩ to be more closely related to the concepts from chemistry than it is to the concepts she apprehended on her first encounter with the idea. The status of her concept ⟨molybdenum⟩ before and after reading Wikipedia are as in Figure 1.⁵

⁵ Note that this diagrammatic scheme is only representative of the difference in tightness of systematic relations between concepts; the concepts themselves will be pictured differently upon further exposition of the minimalist view. Also note that only one of the original (nearly-minimal) systematic relations is pictured, and ⟨meaningless sound⟩ is meant to stand for all of the loosely related concepts that arise (for the subject) from the first context of utterance. So, not shown are the conceptual relations between, e.g. ⟨Bill Lycan⟩, ⟨word⟩, and ⟨molybdenum⟩; and this is just to save space. Thirdly, note that the systematic relations between the four non-⟨molybdenum⟩ concepts in 1b are not pictured, although they do (or at least probably do) exist.

Figure 1. Relative tightness of systematic relations for a concept of increasing complexity.

1a. One (of multiple) conceptual relation at $t_{initial}$.



1b. Conceptual relations at t_{final} .



Lines connecting two concepts indicate systematic relations between those two concepts. The shorter the line, the tighter the relation. Bracketed words in rounded boxes denote concepts. The time $t_{initial}$ is before, e.g., reading Wikipedia, and the time t_{final} is after.

Initially, when the concept $\langle \text{molybdenum} \rangle$ is basic, it bears systematic relation(s) to the concepts deployed in the original context that $\langle \text{molybdenum} \rangle$ was deployed (by Bill Lycan and by way of the verbal utterance “molybdenum”). Through subsequent and repeated deployment of the concept $\langle \text{molybdenum} \rangle$ along with the concepts from chemistry, such as $\langle \text{metal} \rangle$, $\langle \text{chemical} \rangle$, and $\langle \text{atom} \rangle$, the Wikipedia reader takes these new relations to be closer and tighter than the others.

Once her $\langle \text{molybdenum} \rangle$ -concept has been systematized according to the recombinations she ascertains on Wikipedia, the new relations she furnishes become potential ways for her, too, of recombining $\langle \text{molybdenum} \rangle$. For instance, now that she knows real-world molybdenum to be a metal, she can deploy $\langle \text{molybdenum} \rangle$ herself with concepts from chemistry, as by saying, “Molybdenum’s atomic number is 42, and it is a metal.” In fact, the greater tightness of these new systematic relations makes it more likely that she will deploy $\langle \text{molybdenum} \rangle$ in this way than that she will deploy $\langle \text{molybdenum} \rangle$ with $\langle \text{meaningless sound} \rangle$, as by saying, “Molybdenum, molybdenum, molybdenum...what a strange sound...,” or, “Colorless green molybdenum sleeps furiously,” or a Jabberwocky-type utterance. She has effectively systematized her concept

⟨molybdenum⟩ according to its relations to her other concepts of real-world things to which molybdenum is (really) related. This makes the older systematic relations less relevant to her, though still not altogether absent.

Learning what molybdenum is then, is effectively the process of apprehending certain concepts deployed with others (as through linguistic representation), and using these recombinations to inform one's own concept, increasing its number and tightness of systematic relations to other concepts. This, in turn, enables one to instantiate one's own recombinations of that concept. If our subject claims to know that molybdenum's atomic number is 42, as by saying, unprompted and a sufficient amount of time after reading Wikipedia, "The atomic number of molybdenum is 42," this confirms that she has learned this fact.

The process of systematizing a concept by furnishing it with systematic relations increases its complexity. So the complex ⟨molybdenum⟩ in 1b is just a more highly systematized, and consequently more recombinable, 1a. So, 1a and 1b are still the same entity in the mind of the thinker, just with different degrees of complexity.⁶ Importantly, they cannot be two different concepts because the entity to which new systematic relations are fixed is the same entity to which the old ones are. That means the subject still retains the ability to variously recombine ⟨molybdenum⟩ with ⟨meaningless sound⟩ or to recombine it with ⟨metal⟩, as by uttering one of the aforementioned sentences. The continuity of a single concept, then, persists through its increases in systematicity, and through its various recombinations. Adding new relations, creating new avenues for recombination, is just a way of increasing a concept's complexity and

⁶ Across two different people at the same time, we have a slightly different story: the greater the number of systematic relations to ⟨molybdenum⟩ that one person shares in common with another (and his concept of ⟨molybdenum⟩), the more similar their concepts are. So, across two different conceptualizers at the same time, 1a and 1b are the same concept insofar as they agree on some (set of) systematic relations, and they are different to the extent that they disagree. The upshot of this aspect of the minimalist's view is very compelling: the main aim of a conceptualizer who wishes to communicate successfully with others is to have her ⟨molybdenum⟩-concept be as similar as possible to (i.e. the same as) other people's ⟨molybdenum⟩-concept.

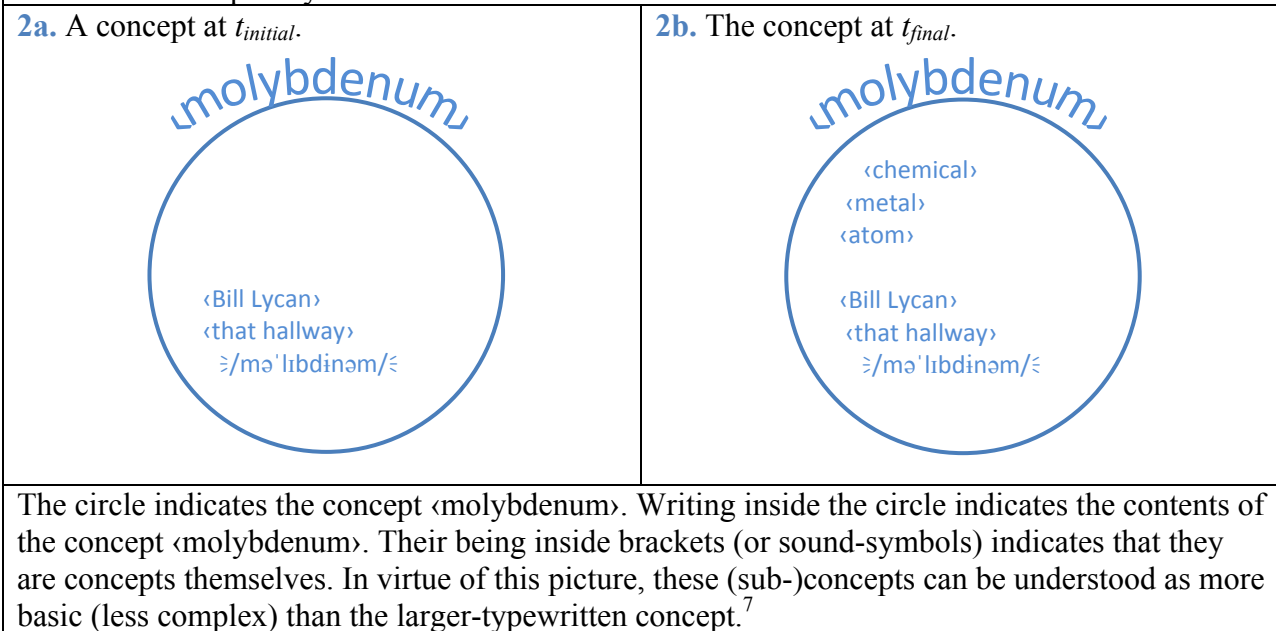
conceptualizing its real-world correspondent entity (whether that entity is molybdenum or Bill Lycan's utterance) in new ways.

C. On the Notion of Conceptual Content

Many views of concepts call for (conceptual) content as opposed to (conceptual) relations. The difference between these two ways of viewing concepts may initially appear a mere difference in perspective, but, as I will show, it is more substantial. An understanding of concepts as things-with-content is quite unhelpful as a model on which to build an account of our mental lives. The way we think, the way our minds work, is much more like the minimalist's relational picture than a picture based on conceptual content, and minimalism is much better suited toward modeling thought. Even so, the two types of accounts are not all that different, and in order to interpret most intellectualist views, the minimalist picture can actually accommodate the wishes of an intellectualist who maintains that concepts have content.

First, though, it will be most important to explain their differences. To someone who thinks of concepts as having content (necessarily), the move from basic to complex that was described in Figure 1 would look more like Figure 2.

Figure 2. The conceptual-content-believer’s version of Fig. 1, or, the process by which a concept increases in complexity of content.



Setting aside for the moment what has already been laid out of the minimalist framework, the pictographic representation of a concept in Figure 2 is meant to be largely intuitive to the reader: at first, one’s concept of anything is very basic, practically void of content; then, when one learns about that thing and gains a better understanding of it, one ‘fills in’ its conceptual content. So what were represented as systematic relations in Figure 1 are now represented as relatively unstructured conceptual content(s) because the conceptual-content-believer may not appeal to systematic relations. (We could, for what it’s worth, come up with a hybrid figure, containing filled-in circles with lines between them, respectively indicating conceptual contents *and* conceptual relations; but such a view runs the risk of redundancy and still falls victim to the objections yet to be raised to the views that give rise to Figure 2.)

⁷ Some content-intellectualists may reflexively disagree to varying degrees with my denoting the content itself as concepts, but, for the sake of argument, I wish to take them literally: if they believe in the recombining of concepts (which they do, as I have explained) and they refer to the content of a concept as “conceptual content,” then the content itself should indeed be a (re)combination of concepts.

Most views that posit conceptual content essentially take for granted a concept's relation to a specified real-world thing and are at once intellectualist. That is, they talk only of 'the concept of molybdenum,' and they claim that to have this concept, one must understand certain things about real-world molybdenum, i.e. that its atomic number is 42, that it is a metal, etc. Since 2a (and 1a) does not include these factors in any way, it is not a concept of molybdenum. This much is acceptable, even to the minimalist: the minimalist would agree that before we know what real-world molybdenum is, our concept is not a concept *of* that real-world thing.

A problem arises for the intellectualist, however, when accounting for the molybdenum story we have just told—specifically, when a thinker acquires knowledge about real-world molybdenum. The conceptual-content-intellectualist who deems 2a either (i) not a concept of molybdenum or (ii) not conceptual at all (or not even pre-conceptual) must hold that there is a deep cognitive divide sometime between $t_{initial}$ and t_{final} . Essentially, he claims that, at some point in the subject's reading of Wikipedia, she acquires a new concept.⁸

If the claim is that 2a is not a concept of molybdenum but is rather the concept of something else, such as the concept of the sound /mə'libdɪnəm/ or of a nonsensical utterance, then, after the subject were to finally discover what molybdenum is and becomes capable of conceptualizing molybdenum, it seems she would no longer be furnishing systematic relations to (or filling in the content of) the same entity as before. This means that she should be incapable of seeing "molybdenum" for its auditory weirdness and as a chemical element; it means that "molybdenum" should cease to remind her of Bill Lycan wherever it indicates a chemical element; and so on. Basically, it should be impossible for a thinker to generate a single thought which deploys the (⟨molybdenum⟩-)concept cross-contextually. In order to entertain such a

⁸ This can also be put in terms of an epistemic qualification: until the subject has (some amount of) knowledge about molybdenum, she does not have a concept.

thought, the intellectualist would claim, the central concept must be swapped: the thinker quickly shifts between deploying a ⟨molybdenum⟩-concept and a ⟨something else⟩-concept; and since these two entities have so little in common, it really is a huge cognitive (and conceptual) leap to do so. We must conclude, furthermore, that such a thinker takes “molybdenum” to represent two distinct concepts whose only link is through this very word, a rather complex form of representation (and not through something more basic).

But this is not necessarily so. It often just so happens that common words of which we already have sufficiently developed concepts may on occasion sound like non-words⁹, or that a word reminds us of the first instance on which we heard it¹⁰. One can think of “molybdenum” in any number of ways, in virtue of the systematic relations that hold of the central concept ⟨molybdenum⟩. We needn’t take account of some profound conceptual shift because our initial mental entity, ⟨molybdenum⟩, remains the same throughout. So, just as we can doubly deploy ⟨molybdenum⟩ as ⟨chemical element⟩ and ⟨atomic number 42⟩, we can doubly deploy it as ⟨chemical element⟩ and ⟨weird-sounding⟩. The mere fact that such seemingly idiosyncratic thought is possible from reading the same word on a page is strong evidence for a relational picture of concepts over a content-based one. And it is evidence that the systematic relations whose employments result in these widely various thoughts are fastened upon the same concept, rather than to different ones.

If the content-intellectualist’s claim is rather that 2a is not a concept at all (nor conceptual nor pre-conceptual), then the problem is even worse. To say this is to deny that we are systematically relating *anything* before we acquire a concept of molybdenum. That is, until we know what real-world molybdenum is, our thoughts about what Bill Lycan said, our thoughts

⁹ Try repeating “molybdenum” five times. Did you forget that molybdenum is a metal?

¹⁰ (“Molybdenum” will forever remind the author of Bill Lycan.)

about the word, our thoughts about the sound, etc., cannot be related to one another; on assumption, then, they are random: they bear no connection, whether statistical or incidental, to one another, and any time they occur simultaneously (as by being triggered by the same word) is pure coincidence.

But, of course, it is possible to have systematic thoughts about a word (or anything) before we know what it means (or is). We can think of “molybdenum” as a weird-sounding assortment of voiced plosives and nasals; this thought is highly intellectual, exhibiting substantial deployment of linguistic concepts, and is very plainly a systematic way of parsing “molybdenum.” And, as in cases like that of our story, its relation to Bill Lycan’s utterance needn’t be random or coincidental at all. But these systematic relations obtain before the entity is considered a concept of molybdenum, and thus before it is considered a concept, for this kind of intellectualist.

The fact that both ways of thinking can be systematic *and* coexistent (without requiring much cognitive effort) is a ruling again in favor of the minimalist’s relational picture of concepts. It is crucial to understanding the cognitive story that all of the systematic relations furnished are furnished to the same entity and hold of that entity indefinitely. \langle Molybdenum \rangle is related both to the concepts from chemistry and to the concepts from the original context of utterance (until and unless there is some sort of cognitive degeneration). Moreover, the existence of some systematic and recombinable representational entity—the minimal \langle molybdenum \rangle -concept—is a necessary precursor to the complex \langle molybdenum \rangle -concept: the subject would have no reason to go on Wikipedia and learn about molybdenum if she had not formerly entertained the possibility that what Bill Lycan uttered was a word. Therefore, the complex concept can only exist as an extension of its more basic, but still conceptual, precursor.

Thus, a more appropriate model of concepts to account for this cognitive story is one in which the new systematic relations are furnished to the same entity as the old ones are. And the content-intellectualist cannot accommodate this. The systematic and recombinatory nature of thought is not restricted to a particularly high level of thought, or what intellectualists may say is conceptual. A thought can be systematic and recombining even prior to its being a concept *of* the real-world thing. In fact, the systematic nature of the precursor (to a concept of a thing) is necessary for the fully-filled-out concept to exist. And there is no evidence of a cognitive leap in the process of furnishing a concept with systematic relations. As one absorbs the words on Wikipedia, what changes is an increase in the number and tightness of systematic relations to the concept, not the concept itself. So we see that, once a thinker has acquired any mental representation from its initial context, the causal story ensures that the systematic relations she subsequently makes are to that same mental representation. From the pre-Wikipedia cognition to post-Wikipedia cognition, the difference is in the concept's complexity, not which concept it is. The same mental representation persists throughout, consistently triggered by any occurrence of the word "molybdenum," or its utterance "molybdenum." And what that word brings to mind, for some thinker, is a product of which systematic relation(s) the thinker deploys on that occasion.

On the minimalist picture, then, there is no cognitive leap in the reading of Wikipedia. (There is indeed a rapid systematization of ⟨molybdenum⟩, but nowhere is this process discontinuous.) The cognitive processes of systematicity and recombining are best captured on the minimalist's relational picture. According to this picture, we do not need to know exactly what molybdenum is in order to think systematically upon hearing or seeing the word. And since, as the minimalist holds, concepts are none other than systematic and recombining mental

representations, our pre-Wikipedia ⟨molybdenum⟩ is no less of a concept than our post-Wikipedia ⟨molybdenum⟩: both are sufficiently systematic (for the minimalist), just to different degrees.

The minimal picture, whereupon thought is relational, is much better acct of cognition than the notion that concepts have definitive contents. Only on the minimalist account can we maintain that someone's thoughts about that thing Bill Lycan said are not completely unrelated to his thoughts about a chemical element. Each respective thought is a different way of deploying the same concept, ⟨molybdenum⟩. In one, it is deployed in combination with the concept ⟨Bill Lycan⟩, and in the other, it is deployed in combination with the ⟨chemical⟩-concept. These two concepts are (as in Fig. 1) systematized to ⟨molybdenum⟩ and thus have only one degree of separation between them. They are related *by* or *through* ⟨molybdenum⟩, and a thought process that goes from ⟨chemical⟩ to ⟨molybdenum⟩ to ⟨Bill Lycan⟩ such as, in the thought that Bill Lycan could have been uttering the name of a chemical, seems not so far-fetched.

Systematization is evidently a constructive process¹¹; and conceptual complexity is a difference of degree. A more complex (or more intellectual) concept just has a greater number of tighter systematic relations upon which it is combinable. A more detailed demonstration of the continuity between basic and complex concepts is given as a critique of Sellars.

D. Debunking the Notion of Conceptual Parts

Before moving on to a theoretical application of the minimalist view, there is one last illusion to dispel about the view that concepts have content, and that has to do with seeing concepts as having parts. For this, we should take a new example concept. Consider what is the concept of a tree. On a first pass, it seems reasonable to consider ⟨branch⟩, ⟨leaf⟩, ⟨plant⟩,

¹¹ There does exist a reverse process as evidenced by memory loss, which appears to be the subtraction or removal of systematic relations. But, given the confines of this paper, we must be content that this and other kinds of reverse-systematization are extremely slow relative to systematization and move on with the exposition of minimalism.

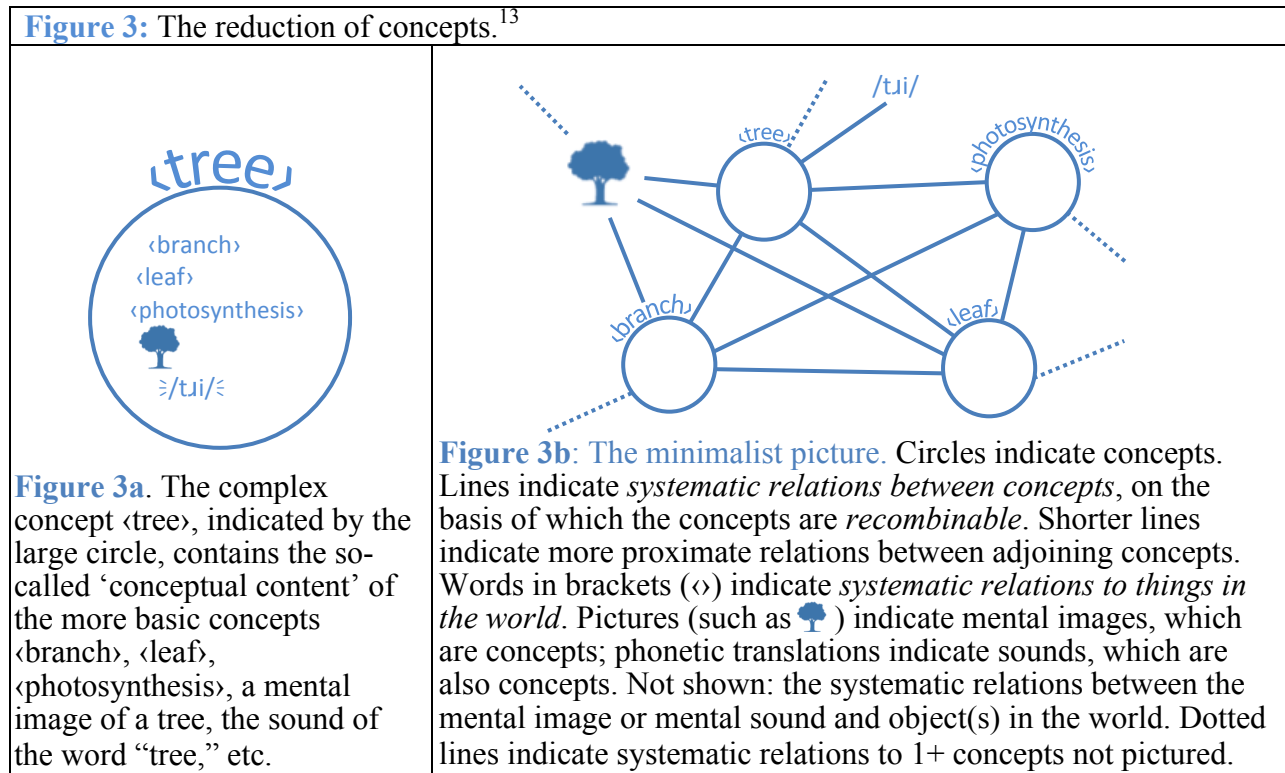
⟨photosynthesis⟩, etc., part of the content of a sufficiently complex ⟨tree⟩-concept, since, as we know, trees are plants with branches and leaves, perform photosynthesis, and so on. And it seems that we would want to say that a person who grasps these very basic building blocks and has seen a tree before should “get the concept of a tree,” so to speak. While this is not blatantly false (it is just manner of speaking; and, of course, the minimalist’s conditions have been satisfied), it ostensibly implies that, in virtue of their relations to one another, some concepts are intrinsically more basic than others. The key notion here, revised in light of the relational picture, is that ⟨tree⟩ is a product of the combination of the concepts of ⟨branch⟩, ⟨leaf⟩, etc., and therefore our ⟨tree⟩-concept is inherently more complex than the concepts of ⟨branch⟩, ⟨leaf⟩, etc. Such a view looks like an attempt to track a hierarchy of concepts: it is natural to think of a ⟨branch⟩ as more basic than a ⟨tree⟩, because, not only are real-world trees composed of branches, but also having the concept ⟨tree⟩ looks to require having the concept ⟨branch⟩, among others.

But, repeats the minimalist, there are no specific relations to a particular concept that hold of necessity. In fact, I am sure I learned what a tree was before I learned what a branch was; meaning, I systematized the concept ⟨tree⟩ before I systematized the concept ⟨branch⟩, and then at some point the conceptions must have intersected or one must have diverged from the other. It seems there are two (or potentially more) ways of looking at the relationship between these concepts; as a result, neither looks to be intrinsically prior to the other. If no concepts are inherently more basic than others, then the relations to a minimal concept are nonspecific; we can view each individual concept as an (empty¹²) mental node in a web of related concepts. See Figure 3. Until we apply more than the minimal, basic relations, all concepts are on a par with each other. Thus, it is only in virtue of thinking (of) ⟨tree⟩ *as* a thing with ⟨branch⟩es, ⟨leaf(v)⟩es,

¹² Note that what had been depicted as a circle with things inside it (in 2a, 2b, and 3a) effectively reduces to an *empty* circle (in 3b).

that ⟨photosynthesizes⟩, etc. and thinking (of) ⟨branch⟩es *as* (solely) parts of ⟨tree⟩s, that one concept is apparently more complex than the other. The concept ⟨tree⟩, in this case, must have undergone a greater amount of systematization (or in other words more recombinations) than ⟨branch⟩ has.

Figure 3: The reduction of concepts.¹³



Having eliminated several intellectualist misconceptions about how we think, we can now see that a generalized intellectualist thesis—one that is directly opposite the minimalist—would claim that there is some amount or tightness of systematic relations required to have a concept. As demonstrated earlier, this would entail that there is some particular point in the subject’s learning about molybdenum at which her thoughts become conceptual. While I do not wish to prove them wrong in this regard (since they are free to conceptualize concepts however works for them), as I will demonstrate at the expense of Wilfrid Sellars, the distinction between

¹³ Note that Figure 3b should not be confused with Camp’s account of “thinking with maps.” More on this later.

what is not conceptual and what is conceptual on any such picture is inherently arbitrary.

IV. Abstraction and Colloquial “having the concept of”

As we saw from the molybdenum example, a basic (almost absolutely minimal) concept begins with a small number of systematic relations to other concepts; these relations increase in number each time that concept is deployed in a new way, and they increase in tightness each time it is deployed in a similar way as before. The process of furnishing systematic relations between concepts increases their complexity. A more complex concept has a greater number of systematic relations, which represent a thinker’s ability to think that concept in a bunch of different contexts. Having highly complex concepts of things allows us to think more abstractly about whatever that concept represents. For the moment, each systematic relation to a central concept can be thought of as an additional degree of *abstraction from* that concept (though, like systematization, we do not wish to think of abstraction as coming in discrete degrees).

For example, abstraction with respect to ⟨molybdenum⟩ might look like the following. [Disclaimer: Though this is wildly—almost egregiously—abstract, and because it is so abstract it may be difficult for the reader to follow, I do think it is necessary to motivate an account of abstraction on the minimalist view. To understand it, I suggest the reader read it through first simply as a random sequence of ideas, (where each idea is a concept, so *only* pay attention to what’s inside ⟨⟩, and) where it doesn’t matter why or how two consecutive ideas are related to each other. It will just be as if we are tracing the lines in someone else’s systematic web of concepts (cf. Figure 1b)—someone who happens to be quite imaginative and perhaps a little deranged.] If we think about ⟨molybdenum⟩ as (having) ⟨atomic number 42⟩ (or, reducing this, as ⟨atomic number⟩ ⟨42⟩), and we can think of ⟨42⟩ as ⟨the meaning of life, the universe, and

everything› (from *Hitchhiker's Guide*), then we can think of ⟨molybdenum› as ⟨the meaning of life, the universe, and everything›. Read:

⟨molybdenum› → ⟨atomic number 42› (→ ⟨atomic number› + ⟨42›) →
⟨the-meaning-of-life,-the-universe,-and-everything›.

Of course, it seems random or insane for someone to consider ⟨molybdenum› as ⟨the meaning of life, the universe, and everything›, especially if she is familiar with real-world molybdenum and the meaning of life, the universe, and everything (within reason for the latter of course). But it is not random in the sense of those two concepts being unrelated; it is merely *highly abstract* (compared to thinking (of) ⟨molybdenum› as ⟨chemical element›). (It is less abstract than thinking (of) ⟨molybdenum› as ⟨*x*›, where the move from ⟨molybdenum› to ⟨*x*› requires a greater number of systematic relations or more-distant ones.) This is the notion of abstraction to be explained on the minimalist picture.

Abstract thought is essentially the deployment of a concept in a novel context, such as in thinking ⟨molybdenum› is ⟨the meaning of life›, which is based on what the concepts share in common (respectively) about the number 42. The ability to think so abstractly requires certain highly complex concepts, concepts with a large number of both proximate and distant systematic relations. So, such a thought-process would be impossible for someone who only conceives ⟨42› as a ⟨number›, ⟨6×7›, ⟨age of my cousin›, etc., and who has never read Douglas Adams. Likewise, it would be impossible for someone who only conceives ⟨molybdenum› as ⟨nonsense›. The abstractness of this particular thought and the subsequent entertaining of that thought requires a relatively high level of complexity of the concepts ⟨molybdenum›, ⟨42›, and ⟨the meaning of life›. It is likely that someone who has such complex concepts as these (of molybdenum, 42, and the meaning of life) has other highly complex concepts of other things.

That is to say, we can assume that a person who is capable of making the connection between <molybdenum> and <the meaning of life> likely has other concepts that are highly complex. In order to have such a complex concept <42>, for instance, she must have, at the very least, read a substantive amount of the *Hitchhiker's Guide*, which is at least two hundred pages of co-deployed concepts. So systematicity is to some extent generalizable.

Moreover, the “train of thought” (that is, the systematic, connected steps) that started at <molybdenum> and ended at <the meaning of life> now underwrites a new systematic connection directly between <molybdenum> and <the meaning of life>. Through the co-deployment of those two concepts in the same train of thought, the thinker's original concept <molybdenum> is once again made more complex by the furnishing of yet another systematic relation directly to it. Admittedly, this relation is a very loose one, and under normal circumstances, one may choose not to consider <molybdenum> <the meaning of life> or vice versa, but it is nonetheless a new systematic relation between those two concepts for this conceptualizer (and now for you, the reader).

Here is a less extreme example. As children or other nonintellectuals we may only conceive¹⁴ of trees (have a <tree>-concept) by thinking about those large wooden, leafy things in our backyards—that is, according to the systematic relations and recombinations of the central <tree>-concept in Figure 3. Over time, we observe the concept deployed in far wider contexts than this. For instance, the concept <tree> is deployed in (applicable to) the idea of family trees, (which are not plants and have only figurative—not physical—branches), stumps (which are not living and have neither branches nor leaves), dichotomous keys (which include stepwise instructions functioning as branches), syntax trees, pixilated trees, etc. Each of these real-world

¹⁴ It may help to note the careful consistency with which we use various inflections of the word concept: conceive (of), conceptualize, conception, conceptual, etc. These particular words are not to be taken lightly, as some philosophers are apt to do.


entities may comprise its own distinct concept, and hence new conceptual connections may be furnished upon the singular concept <tree>. (See Figure 4.) Moreover, thinking about these more abstract relations, such as the one between <tree> and <family tree> and their representations as “tree” and “family tree,” where one is a part of the other, we can deduce that there must be some conceptual relation between trees and families. Indeed, we find one: both trees and families involve a branching structure.¹⁵ At the point of this realization, a new systematic relation is formed between our concepts <tree> and <family>. Once again, this new conceptual connection informs our singular concept <tree>, enabling us to broaden the scope of this concept’s recombining ability and generally make for more abstract thought(s) about trees. The <tree>-concept is now shown to be quite complex.¹⁶

This is far more high-level than the move from basic to complex in Figure 1a to Figure 1b or in 2a to 2b, which only demonstrated abstraction from the original context of utterance. Here we are invoking far more comprehensive systematization than depicted in those diagrams, but the principle is exactly the same. Abstraction comes in degrees, in terms of number and closeness of systematic relations between concepts (and of potential recombinations); each time a new systematic relation is discovered or created between two or more concepts, the complexity of one’s conceptual system increases. The complex, widely-applicable concept of <tree> (Figure 4) is an abstraction of the more-strictly-defined (3b) <tree>, which was itself abstracted from the original-context-of-utterance (1a-level) <tree>.

¹⁵ (Note: we are intentionally bracketing the connection between <branch> and <family tree> for the moment.)

¹⁶ Note that there is no real difference between what might be called “figurative” and “literal” deployments of a concept on the minimalist picture; the “figurative” versus “literal” distinction corresponds to the difference in tightness of the respective relations between various concepts and the <tree>-concept. Camp claims that there is a substantive difference between these two ways of deploying <tree>, and thus in the two ways of thinking about trees (2013). She introduces the notion of a “characterization” (as the flesh on the concept-skeleton) to explain how concepts are deployed in fiction and poetry. I do not wish to argue the point. On the purely minimalist picture, what Camp deems “characterizations” are effectively highly complex, highly abstract concepts. And as will be shown below, abstraction tracks complexity.

Figure 4: Highly complex/abstract <tree>. A highly complex concept (an intellectualist’s or colloquial concept, or Campian characterization) is systematically related to a greater number of other concepts. For ease of understanding, I have removed relations between concepts other than <tree>. Zig-zags are shorthand for longer lines.



Interestingly, the ability to use a concept in a relatively abstract way is exactly the meaning of “concept” invoked when, colloquially, we accuse someone of having “a warped concept of justice,” or when we say “the whole concept behind *James Bond* is so-and-so,” or when we say “I agree with your conceptual analysis of the painting.” The concepts being referenced here—call them <justice>, <*James Bond*>, <the painting>—are highly complex, each with numerous systematic relations. In virtue of such extensive systematization and potential for recombination, one who thinks about the concept of justice in an abstract way has a potentially infinite number of ways to deploy this concept. Analogously, a highly complex concept <tree> applies not only to <branched leafy plant>s but also to <family tree>s, <stump>s, etc.; and being able to think of these things as they relate to the concept <tree> indicates a highly intellectual and abstract level of thought. (This is because, after all, <tree> is also systematically related to real-world trees; and if real-world trees are related to real-world familial structure, then it must be in a

highly abstract way.)

This, I think, is the general intellectualist's motivation for insisting on a particular, high degree of complexity. And, certainly, the ability to engage in abstract thought is a good goal for any conceptualizer; but what the intellectualists miss is that this ability comes directly from the systematic, relational picture of concepts described on the minimalist view. That is, our ability to deploy concepts abstractly is due to the relations that hold between them, and these relations are no different in nature from any others. Just like any other relation to our <tree>-concept, they increase in number, tighten, loosen, branch off from one another, etc., according to how they are recombined. As a result, there is nothing about a concept that disallows abstract deployment; we could always furnish new systematic relations, like that between <molybdenum> and <the meaning of life> (even if the relation is rather weak), or like that between <family> and <tree>. As we have seen, these new connections can be forged by way of relating a series of (other) concepts. This process yields the highly complex concepts <family tree> and <molybdenum as the meaning of life>. And being able to engage in this level of conceptualization, that is, entertaining further thoughts about these (level of) concepts, drawing on or furnishing more systematic relations to them, is precisely what we mean when, colloquially, we say someone "gets the whole concept of a tree." This degree of thought is often employed in, e.g., metaphor, humor, sarcasm, idiom, hinting, etc.

V. A Critique of Sellarsian Concepts

Similar to the minimalist's processes of systematizing and recombining, Wilfrid Sellars claims that mental representations are sorted and concatenated when we think. Sellars is an intellectualist, however, which derives from his belief that language is necessary for this sorting and concatenating to occur. He considers linguistic (or logical) thought to be different in kind

from nonlinguistic thought. (And he claims that only humans are capable of the former.) In this section, I will demonstrate that linguistic concepts (that is, concepts that require the thinker to have or use language) are not a different *type* of concept, as Sellars claims, but are rather concepts of a particularly high degree of complexity. And as such, they fit (perfectly well) on the minimalist's spectrum. It helps that Sellars and the Camp-style minimalist have strikingly similar accounts of cognition; however, as I will show, since Sellars sees linguistic thought as different in kind from nonlinguistic thought, he looks to have an artificial break in the complexity spectrum.

A person learns to ski through verbal instruction. A beginning skier is told to lean forward, put her weight on the downhill ski, keep her shoulders square, control speed by turning, plant her pole, etc. Each of these verbal instructions functions as an explicit mental representation to determine her behavior. As she becomes a competent skier, she internalizes these instructions and becomes less conscious of them while skiing. Eventually, the mental representations of the rules she needs to follow and the bodily movements dictated by these rules function only in the 'background.' If a skilled skier were to consciously think about every individual movement, every instruction she had received, then her thoughts would get in the way, and she'd barely be able to move. The representations are functioning unconsciously: when she sees a bump in the snow, she skis around it without much thought.¹⁷

As thinking humans, we possess the ability to represent things in the world and deploy, relate, and act on these representations; this is what separates us (at least) from non-thinking things like plants. Even subconsciously, when we may not be fully aware of their presence, mental representations are functioning in determining our behavior. Certain other types of human behavior are not resultant from our ability to mentally represent, such as removing one's hand

¹⁷ The use of a skier in this example is inspired by Searle (1983).

from a hot stove and blinking. These kinds of behaviors are (in the relevant cases) strictly physical¹⁸ responses to physical stimuli; in other words, one does not have to think in order to remove one's hand from a hot stove. One does have to do some (however small) amount of cognitive work in order to avoid a ski mogul, even if one is a highly skilled skier. The difference between these two types of behavior is the response most immediately following sensory input.

When a skier sees a mogul in front of her, her response—carving a turn just so—is, in Sellarsian terms, mental and then physical; when the cook feels a painful burning on his hand, his response—removing his hand—is immediately physical. For the cook, the mental component may come in later: the physical reaction of removing his hand is quite immediate, it occurs before he has time to think. By contrast, something intercedes between the skier's seeing the mogul and her turn around it. The skier has (to some extent) drawn on or formed the relevant mental representations, as of ⟨mogul⟩ being a ⟨thing to ski around⟩, which allow her to perform her task. It may seem that she reacts just as quickly to seeing the mogul as the cook does to feeling extreme heat, but the reactions are nonetheless starkly different in nature.

So the skier conceptualized where the cook did not. (Recall that we are only considering the time between the stimulus and the reaction, not the cognitive processes that may have taken place after the reaction.) According to Camp, minimal conceptual abilities “encompass a type of engagement with the world” (Camp₁, 303). So the skier is (somewhat passively) engaging parts of her environment by representing them and reacting according to those representations. The reaction of the cook, on the other hand, does not look to require any sort of critical engagement with his surroundings. Sellars has a bit of a different account of conceptuality: he writes of concepts as somewhat of a gateway between two domains. He writes, “our *concepts* pertaining to

¹⁸ The term “physical” is intended as a contrast to “mental.” It is meant to include “biological,” since the response of removing one's hand from extreme heat is also a feature of human biology.

the middle-sized objects of the perceptual world are our conceptual point of entry into the domain of the physical” (§14). So, since the concepts of heat or pain were not deployed before the cook removed his hand, he did not enter the mental domain (until, possibly, he considered them afterward); and the skier did enter the mental domain by deploying the relevant concepts originally introduced to her by her ski instructor.

Sellars and Camp have similar functional accounts of concepts. As “a medium in which we think,” (Sellars, §11) our concepts allow us to recast items in the physical domain as systematically corresponding entities of the mental domain. Just how that correspondence is systematized is explained most succinctly by Camp: “the...features of the [mental representations] themselves often reflect salient physical features of the objects or properties being represented” (Camp₂, 158–9). Furthermore, the structural relations between the concepts themselves reflect the structural relations among the things in the world we perceive. The skier uses both of these systematic qualities when assessing the mogul and making her turn: salient physical features of the mogul, such as its dimensions, its being made of snow and being located on a ski mountain, etc., are cued to particular aspects of her (complex concept of) ⟨mogul⟩, which helps her recognize the object, and, in also recognizing that it is a certain distance away from her and that she is moving at a certain speed, she is able to time her movements accordingly.

Camp’s account of the function of concepts is slightly more symmetrical across “domains” than Sellars’s: both consider concepts to be intermediaries between the mental and the physical (in that they are *mental* and representational of *physical* things), but for Sellars alone, they are tied more closely to the mental. A significant component of Camp’s overall view is humans’ ability to “think with maps.” When we think in this way, Camp claims, our concepts are very much like discretely located and physically extended icons on a map, and the relations

between concepts directly reflect the physical, spatial relations of objects in the world (Camp₂, 160–1).¹⁹ Sellars’s view does not seem to allow for any correlates of physical spatial relations (at least not in the way Camp’s does). He acknowledges mapping as a form of mental representation, but, as we will become evident later, it does not qualify as conceptual representation.

Therefore, the main functional application of concepts is different for each philosopher: rather than systematizing (and recombining) mental representations, as Camp holds, Sellarsian conceptual abilities are said to “sort” (and “concatenate”) representational *states*. One operates on representational states, while the other operates on representations themselves. So a representational state, or mental event, on Sellars’s account, is a state of awareness, and the sorting that occurs is a type of association or inference (or primitive inference, as is the case with animals) between different mental events. “Primary concepts . . . can be shown not to concern unique modes of relationships between *mental events* and *reality*, but rather to provide a technique for classifying *mental events* by reference to paradigms in our *background language*” (§35, emphasis added). So, this kind of systematization concerns not the relationship of mental to physical, but rather the relationship of the mental to another mental entity, the background language.²⁰

For Camp, the task of the most basic concepts is to “enable a creature to represent . . . the world,” and the most basic cognitive processes “produce systematically related representations of various . . . [physical] states of affairs” (Camp₁, 282). This is systematization of the mental by

¹⁹ This is not to be confused with the slightly map-looking thing in Figure 3b. Thinking with maps utilizes a particular representational system, as does my depiction of the minimalist picture, but the similarities end there. Thinking with maps is much more complex than thinking minimally, and so the systematic relations of 3b must be more highly systematized and recombined before they become icons or symbols on a map.

²⁰ This is the basis for McDowell’s objection that the Sellarsian picture looks like “frictionless spinning in a void,” (1996) though I think Sellars avoids this objection sufficiently in his account of language-entry (1969).

relation to the physical, as required by the minimalist in Section II. Camp's concepts systematize mental representations according to individuals in the world; Sellarsian concepts systematically sort mental events according to the background language. The idea that concepts systematize mental events as opposed to systematizing individual representations is intimately tied to the importance, on Sellars's view, of language (and, interestingly enough, *not* to a different account of cognition).

Sellars relies heavily on the systematic nature of language. “[I]n the domain of the mental, language is primary in the order of knowing” (§4). Put more concretely: we conceptualize with language. Language permits the sorting of thoughts (aka mental “episodes,” which themselves are non-linguistic) because, he claims, thoughts are “analogous to . . . linguistic structures, and are functionally connected with linguistic behavior” (§7). This distinctive parallel presumably makes it only natural for us to systematize our thoughts by using language as a sortal. So although we can have thoughts that are not linguistic according to Sellars, language itself is naturally, functionally superimposed on thought through repeated instances of linguistic behavior so that, by and large, we think primarily *with* language.

Thus, on Sellars's account of conceptualization, mental representational events (thoughts) are filtered through the background language. Linguistic concepts are the so-called ‘sortals’ of mental events, which is to say, the dual function of a concept is to (1) “refer to” an object and (2) “characteriz[e] it as thus-and-so” (§60).²¹ In yet as many words, Sellars also writes, objects are assigned “referential symbol[s],” and the entire representational event is “given . . . a character” (§72e). This latter claim is a significantly stronger one on his view because it invokes symbolic representation as necessary for certain cognitive processes. The referential symbols Sellars mentions here are either logical or linguistic symbols, and the character[ization] is effectively a

²¹ Basic representational states perform both these functions as well, which I will address below.

linguistic description (in the background language).

Moreover, the relations between representations are said to be of either a logical or properly inferential type. And this is distinct from mere causal, associative relations, which can hold between mental states. Representational systems that cannot use logic or language (what Sellars calls “pre-logical” but he will later refine to “propositional” representational systems) can only produce “a chain of representational states” from goal states to action states (§69). In these systems, thoughts are represented as, for instance,

‘Smoke here’ / ‘Fire nearby’,

but not as

‘If smoke anywhere, then fire nearby there’ / ‘Smoke here’ / ‘Fire nearby’,

(§85). The difference between these two toy thought processes (which are supposed to be in contrasting representational systems, but we are limited here by the medium of this essay) is the presence of a logical relation between the two mental states. In the first type of representational system, the state of thinking ‘Fire nearby’ is only causally linked to the state of thinking ‘Smoke here,’ we can assume as a result of the creature’s experience with real-world fire and smoke. In the second, the two states are logically related (but also distinguishable from that logical structure), which allows for a proper logical inference to take place. So in order for a thinking creature to do anything other than associate causally between mental states, language and/or logic is necessary. It should be quite evident now that conceptual abilities are firmly grounded in linguistic abilities on Sellars’s account. Ultimately, he defines our entire conceptual representational system as a “linguistic” one.

The concept-minimalist could object to Sellars’s ontology of concepts and conceptual abilities at this point, citing that the requirement of language is too strong: we are perfectly

capable of systematizing and recombining (or sorting and concatenating) our mental entities without the use of words or logic. For example, the skier was able to deploy visual and tactile (mental) representations in a complex and functional way without having to make explicit (or linguistic) reference to the instructions of how to ski. That is, she did not need to (to use Sellars's own terms) "think out loud" and tell herself, "knees forward, plant left pole, shift weight to right ski," etc. If she did not think linguistically (think-out-loud), Sellars would say she did not conceptualize. But it seems that, in order to perform such a highly calculated maneuver—timing her turn, balancing her weight, turning around the mogul just so—the relevant concepts must have been deployed to some extent. Without making use of a background language or logic, it looks like we still have the ability to sort mental representations.

An answer for Sellars would come by way of his account of the more primitive representational systems just described. Though he does allow for language-less ways of sorting mental events, he does not consider these cognitive processes to be conceptual. So, it would seem, it comes down to logical and associative inference versus causal association. This is the fundamental difference, for Sellars, between human and animal representational states, respectively, and most likely, he would say the state of the skier (though not properly conceptual) is somewhere in between.

A closer look at Sellars's picture reveals that he acknowledges different types of complexity of representational systems and therefore cognitive abilities. Though animals do not possess concepts (Sellars claims), they look to effectively possess a more primitive sortal equivalent. Animal representational systems have a "propositional" structure (§73): like the human representational systems of language and logic, they can represent an object in the world as being of a character, but that's about all they can do. (Sellars calls them "propositional"

because the object and its character are analogous to the syntactic subject and predicate, respectively, of a *proposition*.) They are less complex than linguistic or logical representational systems because they do not allow an animal to carve up its “perceptual awareness” very far. The basic units of mental representation are not discrete linguistic concepts but rather whole representational states (§56).

Bee thought, for instance, works this way on his account (and similarly on Camp’s). A bee represents itself and the locations of some objects around it (the sun, nectar, the hive) relative to its own location, constructing a cognitive map. But this map is a “system of dispositions and propensities” to move in particular directions (§56)—not a system of discrete concepts. And although a bee can “[be] aware of something as something,” (“pre-linguistic awareness;” §57) the bee only possesses a singular, aggregate mental representation of all the (relevant) items in its environment together. And since a bee’s sorting function operates on representational states as wholes, the output is simply a causal association of two or more representational states. Thus, we have that the bee is unable to separate out the individual objects from each other or from the physical state as a whole and rather individuates between whole physical states. Yet it can still refer and characterize. So the bee must be using concepts’ more primitive cousin.

Sellars attributes similar cognitive capabilities to rats as well, possibly with a slightly higher degree of complexity: their output is likewise just a set of causally associated states, but rats are also able to represent spatial structure (§63). He also claims that this primitive form of representation is innate to human beings (§57), which is evidence that he may recognize a spectrum of complexity of representation similar to the minimalist’s for concepts.

Once again, the minimalist is not satisfied. That cannot be what the skier is doing; the skier is capable of responding to minor changes in her environment, such as different-sized

moguls or different trail grades, and adjusting her responses accordingly. Her mental map is certainly more malleable and more fine-grained than the bee's or the rat's. She can accommodate changes in her environment that bees and rats cannot. So why should the skier's conceptual abilities be so severely limited by her lack of applying linguistic concepts? Because of Sellars's reliance on language, and the fact that he attributes to rats—mammals like us!—more or less the same cognitive abilities ascribed to bees, he does not seem to provide an appropriate middle ground.

Similar to Sellarsian concepts' sortal function, but different in one important aspect, the systematicity of concepts on Camp's picture applies to different items within a mental (or physical) state, rather than to mental states as wholes. And even though both would claim that the bee's cognitive map is not a conceptual one, they reason from theoretically opposed angles. If a map is said to be composed of concepts, then physical objects and properties in the organism's vicinity must be individually represented (or representable) on the map by the organism; the mental state itself (the map) and its components must be "segmentable into re-deployable parts" (Camp₃). This means the (representations of) entities on the map are distinct from (the representations of) their particular locations on the map and can be represented apart from the map or in different locations on it. So, if the bee needs to construct a whole new mental map each time one item in its environment changes location, then the maps are not segmentable, and thus not comprised of distinct, recombinable concepts. The minimalist rules the bee's mental states non-conceptual because they are not segmentable; Sellars rules them so because they are not logically or linguistically represented—which would have effectively made them 'segmented,' though he does not have an equivalent term for this. This is a surprisingly subtle difference between their accounts.

While Camp hesitates to call the bee's mental representation a "map" for the reasons just mentioned, Sellars does explicitly refer to the bee as a cognitive map-maker. This seems to indicate that he considers the bee (and other animals) capable of a bit more than strictly non-systematic thought. Propositional thought, by comparison to linguistic or logical thought, is less structured and less systematic. But because he claims that humans can innately represent as bees do, the bee looks like a proto-conceptualizer on Sellars's account. Presumably, then, we represent in somewhat the same way as bees do whenever we think non-conceptually. This suggests a particular, distinct 'starting-point' for human cognition, and it clearly indicates that (Sellars thinks) we think differently before we acquire language and therefore concepts. Therefore, any human who perceives or experiences the physical world without language to sort (and systematize) the individual components of each mental state, is thinking only "propositionally"—including the skier.

Linguistic concepts allow us to divide up our representational states more finely; propositional representational states are less fine-grained than logical or linguistic ones. So it seems Sellars would say of the skier that if she is not sorting with linguistic concepts, then there are some physically individuated objects (objects in the physical domain) is that she not capable of mentally representing individually: for instance, she is not thinking about the movement of her knees as distinct from the movement of her shoulders, or she is not thinking of each individual instruction she was given as a beginner, or she is not thinking of the mogul as a thing-to-ski-around.

A different possible answer for Sellars would be to consider the skier's cognitive state to be finely systematized (antecedently) but still not conceptual. Since she had formerly (when she was learning how to ski) used linguistic concepts to think about skiing around a mogul, her

environment now when she approaches a new mogul has already been systematized (and recombined) for her, such that she no longer has to conceptualize the ski slope and what she is doing. So she is still not actively working with concepts, and she still cannot adjust her mental map without linguistic or logical concepts.

So Sellars and the minimalist have come to an impasse about the conceptuality of the skier. The minimalist would claim that the skier is, strictly speaking, representing conceptually; Sellars would say she is not. And yet their analyses are strikingly similar: they both place the experienced skier somewhere between the pre-conceptual bee and the highly conceptual beginner skier; and they do so by assessing the relative complexities of their representational states, the amount of things they are capable of individuating.

At this point, the root of the minimalist's objection mentioned in Section III above becomes clear. The requirement that concepts be linguistic (or logical) makes Sellars's distinction between conceptual and nonconceptual seem superficial. It is unclear where to draw the line between a conceptual skier and a nonconceptual one. At what specific point in the learning process did the skier stop skiing conceptually? Or if the experienced skier were to start calculating her turns just a little more carefully, why should this be conceptual where before it wasn't? It seems wrong to so highly privilege linguistic representation, particularly when any other kind of representation we do quickly reduces to mere insect-thought.

The problem is his requirement of a particular system of representation. On the minimalist view, and on Camp's view, language is just one of many representational systems that employs concepts. For Sellars, it is the only conceptual one. According to the latter, when we think without language, we are representing the world as animals do, and therefore all thought is propositional: it systematizes mental states, not individuals within a state.

According to the Camp-type minimalist, there are many—possibly infinitely many—different representational systems, most of which humans are capable of using to represent the world, and most of which we share with animals. Camp considers language to be a distinguished form of representation only in the sense that it may allow uniquely high levels of abstraction that we might not be able to achieve without the tools afforded by linguistic representation. But the notion of abstraction is not addressed by Sellars; if he were to accommodate the minimalist, he would have to posit a specific degree of abstraction that only language can achieve, and he would have to claim that attaining this degree of abstraction is somehow necessary for conceptualization. Since Sellars does not address the notion of representational abstraction, however, language is on par with potentially any other representational system that can systematize and recombine to the same degree.

Since Sellars argues that language and logic are *not* the only means by which thoughts can properly be systematized and recombined²², it is indeed a bit flat-footed of him to claim that only this way of representing things is conceptual. On his view, we begin as propositional thinkers, the same as some animals, equally capable as they are of systematizing the physical world into mental representations. When we acquire language, though, we suddenly put ourselves above the rest of the thinking world. With linguistic concepts applied to mental representations, we can finally think *conceptually*. If we can set aside this ill-advised distinction, though, we can see that his account of the range of human cognitive abilities (since we start on the non-conceptual side of things) easily agrees with a continuum of complexity. What he considers ‘basic cognitive abilities’ effectively count as minimally conceptual abilities for the minimalist.

²² The reader may note that I have not discussed the Sellarsian version of the recombability criterion: I have mainly analyzed the systematicity component, “sorting.” Nonetheless, his criterion of “concatenation” is open to a similar line of attack as that of the present critique.

Our ability to use linguistic concepts is less of a deviation from the cognitive abilities of our fellow animals than Sellars makes it out to be. To be able to think in language is no marked departure from the rest of our cognitive abilities, because it is just another set of concepts. There are potentially many other ways to systematize and recombine mental representations which could be equally complex.

VI. A Response to Sellars

By the minimalist's lights, language or any representational system is just another set of concepts by which to systematize a concept: ⟨meaning⟩, ⟨part of speech⟩, ⟨definition⟩, ⟨connotation⟩, ⟨orthography⟩, ⟨pronunciation⟩, ⟨rhyming⟩, etc., are each concepts themselves.²³ And using language to systematize our otherwise non-linguistic concepts, 'thinking out loud,' as Sellars calls this process, is nothing other than furnishing new systematic relations between the central concept and the relevant concepts of language in a highly abstract form, greatly increasing its complexity. For example, when we systematize the otherwise more basic ⟨tree⟩-concept as an English noun by furnishing two new relations between ⟨tree⟩ and the complex ⟨noun⟩ and ⟨English⟩ concepts, the realization of these two new relations entails a specification of how the ⟨tree⟩-concept is to be subsequently deployed.

One can observe these specifications at work when, for instance, it would sound strange to the reader if someone were to say, "The sequoia tree-ed for over two thousand years." The weirdness of this sentence stems from the concept ⟨tree⟩ deployed as a ⟨verb⟩ instead of as a ⟨noun⟩, as it usually is. And it is not strictly incorrect to deploy ⟨tree⟩ in this way, unless it is specified that the ⟨tree⟩-concept is to be deployed as a noun in a sentence. The effect of furnishing a relation between ⟨tree⟩ and ⟨noun⟩ and making use of this relation—call this deploying the complex ⟨tree⟩-⟨noun⟩—is that real-world trees and their linguistic representations

²³ The "ways of thinking" approach is probably the most helpful angle here.

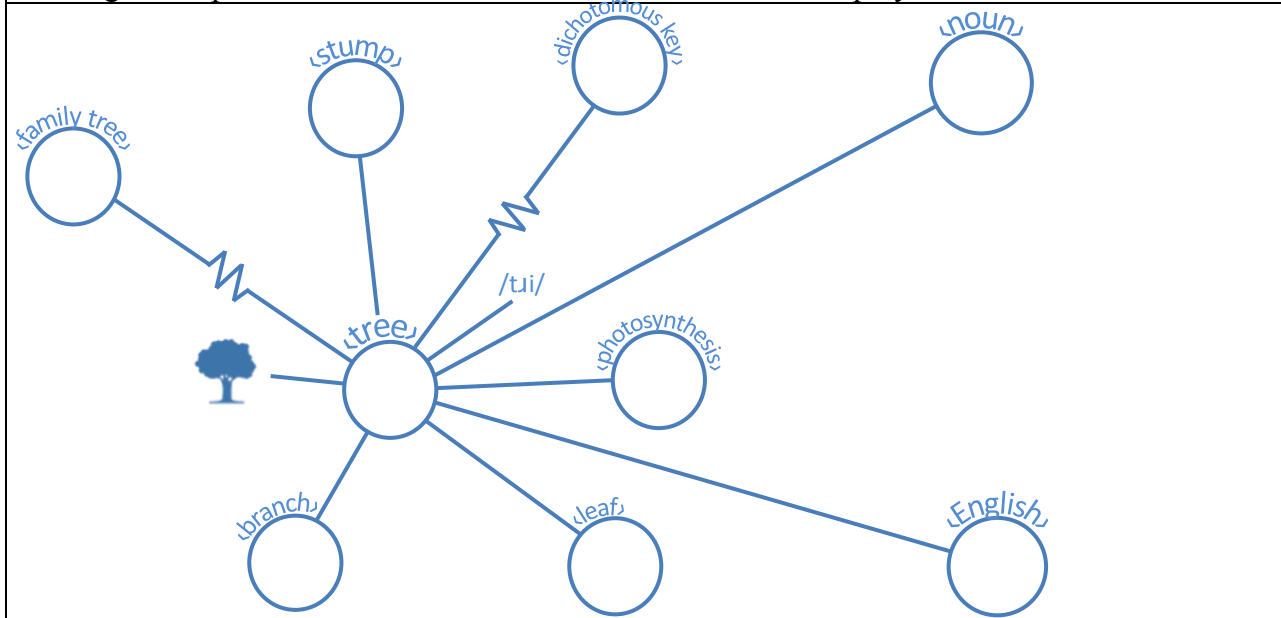
may only occupy certain positions in a sentence and take on certain suffixes. So when the above sentence sounds strange to the reader, it is because his/her <tree>-concept is tightly bound to his/her complex <noun>-concept and all the syntactic and morphological rules this relation entails.

So when we use language as a representational system for our otherwise more basic concepts, like <tree>, we are effectively taking account of certain systematic relations between them and the highly complex concepts having to do with language. Elevating our concept <tree> to the word, “tree,” entails the establishment of systematic relations between <tree> to concepts such as <noun> and <English>. Indeed, the process of using any representational vehicle for a concept is nothing more than systematically relating a more basic concept to a particular set of concepts encompassing that representational system. Each representational system is itself just a set of rules for how to systematize and recombine concepts; and each set of rules can be triggered by a singular concept, as was the case with the complex <noun>-concept. So when we think (of) <tree> as rather a <tree-noun>, we are specifying that this concept can only undergo certain recombinations: e.g., before a verb, or after a preposition.

It only looks to be different in kind from the nonlinguistic concept because of the apparent complexity of the rules involved. But when we realize these rules are themselves just more concepts (and concepts about recombination, no less), the uniformity of complexity rematerializes. And each time we make our <tree>-concept more complex by systematizing it to <noun> and then to <English> (or vice versa), we are indeed making it more complex through exactly the same process as before. See Figure 5. (Note that this is, specifically, the linguistic version of <tree> from Fig. 4, but that for a concept to qualify as linguistic, it needn't necessarily be furnished with all the relations from Fig. 4; for a concept to be linguistic, it must be

systematically related to the concepts from language, i.e. ⟨noun⟩, ⟨English⟩, etc.) Thus, there is no difference in kind between a concept for a thinker who uses language to systematize her concepts and a concept for one who does not. We can therefore preserve our continuum of conceptual complexity.²⁴

Figure 5: A linguistic ⟨tree⟩. The concept is now additionally systematically related to certain concepts from language: ⟨noun⟩ and ⟨English⟩. As such, we can think of ⟨noun⟩ and ⟨English⟩ as relatively complex concepts themselves; and when ⟨tree⟩ becomes systematically related to them, this engenders particular sets of rules about how ⟨tree⟩ is to be deployed.



VII. Concluding Remarks

It is widely acknowledged that language is the most complex and abstract way to represent thoughts. To be sure, it is very unlikely that any other representational system can convey the level of complexity required for thoughts such as, “I know that you know that I know that you know . . .” (I can certainly not depict that in a diagram!) And language gives us puns, metaphor, figurative speech, and other mechanisms for conceptual recombination. But language is far from perfect, and our conceptual capacities far outrun what can be represented with words,

²⁴ A minimalist critique in a similar vein could be applied to numerous other accounts of concepts, including Davidson (based on his claim that, in order to have a concept, one must have the concept of belief), McDowell (who claims that concepts come to us from the Given), and Millikan (based on the idea of object-tracking).

sentences, and essays. The numerous additional rules that language involves, which are exactly what make linguistic concepts so complex, also—rather paradoxically—restrict the way in which the concepts can be systematized or deployed any further. As ⟨tree-noun⟩, the concept can only be deployed, e.g., before a verb or after a preposition. But as ⟨tree⟩ (without a part of speech or meaning or connotations) it is infinitely deployable.

While it may be true that language is the most complex and abstract way to represent thought, this does not mean it allows us to express everything we can think. As we have seen, there is strong reason to believe that thought is most fundamentally a relational process, where mental entities are somehow systematically connected to one another such that they may be readily combined to form new thoughts—with a surprisingly rapid increase in complexity. The minimalist's relational picture is the best approximation to the true nature of cognition.

WORKS CITED

- Camp, Elisabeth. "Putting Thoughts to Work: Concepts, Systematicity, and Stimulus-Independence." *Philosophy and Phenomenological Research*. Vol. 78, No. 2. Philosophy and Phenomenological Research, LLC. 2009. Web.
- Camp, Elisabeth. "Thinking with Maps." *Philosophical Perspectives*. Vol. 21. 2007. Web.
- Camp, Elisabeth. "Playing with Concepts: Logic, Association, and Imagination." UNC Chapel Hill. Lecture. 15 Nov., 2013.
- Davidson, Donald. *On the Very Idea of a Conceptual Scheme. Proceedings and Addresses of the American Philosophical Association*, Vol. 47, 1974. pp 5-20.
- McDowell, John. *Mind and World*. Harvard UP. Cambridge, MA. 1996.
- Searle, John R. "The Background." *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge UP, 1983. 150–1. Web.
- Sellars, Wilfrid. "Language as Thought and as Communication." *Philosophy and Phenomenological Research*, Vol. 29, No. 4. Jun., 1969: (506–27). Web.
- Sellars, Wilfrid. "Mental Events." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*. Vol. 39, No. 4. May, 1981: (325–345). Web.