

Hume and Fodor on Complex Ideas

Peter Eichman

May 2007

In his *Hume Variations*, Jerry Fodor takes a critical look at Hume's Theory of Ideas. While he thinks there is much to recommend some parts of Hume's Theory, he takes issue with other parts. In particular, Fodor thinks that while Hume is correct in construing the mind as representational (he takes the Theory of Ideas to be an early Representational Theory of Mind), he finds fault with Hume's associationism and empiricism.

In this paper I outline Hume's theory of complex ideas, and the relations that hold between the simple parts of these ideas. Fodor thinks that these are the *only* relations that Hume uses in his associationism, and because of the problems this presents, Fodor believes that we should abandon this portion of Hume's theory. This may be too hasty, though, and I briefly sketch some of the responses open to the Humean.

1 Complex Ideas

An important component of Hume's Theory of Ideas is his notion of complex ideas. In Hume's theory of mind, ideas are roughly analogous to concepts. The other important category in his theory of mind are impressions, which are roughly analogous to perceptions. Both of these types of mental objects come in both simple and complex forms. Here I will be focusing primarily on the nature of ideas or concepts, although I will make reference to the nature of impressions as needed.

Complex ideas are composed of simple ideas. The property of being simple is, according to Fodor, "just not having semantically evaluable parts" (Fodor, 2003, p.56). In other words, the sort of simplicity that Hume is aiming for is a sort of conceptual atomism. There are just some ideas whose contents cannot be unpacked into anything further.

It is also worth noting that Humean simple ideas do not even unpack into “use conditions” or any other sort of functional definition. Rather, Fodor argues, the Humean Theory of Ideas follows in the Cartesian tradition that maintains that concept possession is the primitive notion, and the knowledge of how to use a concept is something that follows the acquisition of a concept, not something that is constitutive of acquiring that concept. That is, Fodor maintains, Hume’s Theory of Ideas is the predecessor to modern representational theories of mind.

Another distinction that Fodor notes is the difference between structurally or syntactically simple and semantically simple. There is no particular reason that we should think that the structural simplicity of a concept should entail that it is semantically simple. By syntactically simple, Fodor appears to be drawing a parallel to monomorphemic words in natural language. For although a word like ‘god’ or a concept like GOD¹ do not have any internal syntactic structure, it would be surprising if they had no semantic complexities hidden within them.

This example of a structurally simple yet semantically complex idea can be contrasted with an idea like RED BALL, which is both syntactically and semantically complex. It is structurally complex because it can be split into two constituent parts (namely, RED and BALL). It is semantically complex because the meaning of the entire construction RED BALL can be decomposed into the semantic parts RED and BALL.

We notice here that there are cases where a complex idea is both structurally and semantically complex, the (simple) syntactic and semantic constituents it breaks into are the same. However, there are also cases where the syntactically simple constituents are not semantically simple constituents: e.g., complex concepts like HARSH GOD or GOOD PERSON. So we should be cautious about assuming that syntactically simple means semantically simple².

¹I am following the convention of delimiting words with single quotes (‘cat’) and spelling ideas or concepts in small capitals (CAT).

²We should also be on the lookout for syntactically complex ideas that are, nonetheless, semantically simple. However, these examples are harder to come by.

2 Relations Between Ideas

Hume envisions at least three sorts of associative relations that may hold between two ideas: resemblance, contiguity, and causation. These form a subset of Hume's seven philosophical relations. Bennett (2001) further divides contiguity into two relations, spatial contiguity and temporal contiguity. The abstract schema that all four of these implement is given by Bennett as:

When you have an idea I_1 in your mind, this is likely to cause you to move on to having a second idea I_2 such that $R(\dots I_1 \dots I_2 \dots)$ (Bennett, 2001, p.234).

One of the key aspects of these relations for Hume is the fact that they are “not to be consider'd as an inseparable connexion”, nor are these the only ways in which the mind (and specifically, I think, the imagination) can combine ideas (*THN* 1.1.4, p.10). Rather, these relations should be thought of as “a gentle force” which brings together ideas in their most natural configurations.

Of these four relations, resemblance is the most free. It is therefore, according to Hume, also the most prone to error. It also both constrains other philosophical relations, as well as having a sort of generality or abstractness constraint on it. “Tho' resemblance be necessary to all philosophical relation, [...] when a quality becomes very general, and is common to a great many individuals, it leads not the mind directly to any one of them” (*THN* 1.1.5, p.14). In other words, resemblance is only an effective sort of relation when the property that we are focusing on to determine resemblance of two objects is not too general. For example, it is easier for the mind to focus on one particular thing that resembles a three-inch red ball than on any one particular thing that resembles a large object. The latter is too general of a criterion of resemblance, and causes the mind to try and connect to too many things (planes, elephants, the Empire State Building, etc.). Thus there is no particular idea (or even small handful of ideas) that stands out before the mind³.

While resemblance is the rather free subjective relation between particular ideas, the second relation of spatial contiguity is supposed to be an objective relation that can hold between either

³This appears to be related, albeit orthogonally, to the problem of selecting appropriate categories for information. If the categories are too big, they are useless, as it is hard to find anything within them. And if they are too small (down to the level of individuation, say) there are equally useless.

particular or general ideas (Bennett, 2001, p.236). The subjective/objective⁴ difference is supposed to capture the fact that while resemblance is focused on the idea, contiguity is focused on the object that the idea is supposed to represent. When I think of sandpipers and then think of ocean waves, it is not because my idea of sandpipers is in any way spatially contiguous or near my idea of ocean waves. Rather, it is that I usually observe (actual, objective) sandpipers to be near (actual, objective) ocean waves.

One of the reasons that Bennett separates out temporal contiguity from spatial contiguity is because temporal contiguity, like resemblance, is a subjective relation. However, it is also a general relation, since it concerns types of ideas rather than specific ideas. Two ideas are temporally contiguous if they tend to come before your mind at the same time. So if this relation holds between I_1 -like and I_2 -like ideas, then having an I_1 -like idea will bring to mind some I_2 -like idea.

The fourth relation, and for Hume the most important, is that of causation. Causation, like spatial contiguity, is objective; it is really a relation that involves the objects of thought, not the thoughts themselves. And causation can either be particular or general.

So these are Hume's associative relations between ideas. Fodor furthermore claims that these are the relations that glue together complex ideas for Hume. As we shall now see, this presents something of a problem for the individuation of complex ideas.

2.1 Race Condition Problem

One problem that Fodor sees with the Theory of Ideas is that as far as Hume's description of the associative relations between ideas is concerned, there is no way to distinguish between two simple ideas that are part of a complex idea, and the same two simple ideas when one simply follows the other in coming before the mind. Fodor brings this out in his discussion of simple predicate expressions like MRJAMES BITES; here I am moving beyond Fodor's original examples in order to make a more complex point.

⁴In actuality, there may be less to this distinction than Bennett thinks. In Hume's later skeptical doctrines, he treats (subjective) ideas in the same way as objects (see for example *THN* 1.4.2). In particular, he thinks of ideas and impressions *themselves* standing in spatial and temporal relations, not just the objects they are supposed to represent. So it seems that all of these relations may come to be subjective for Hume in the end. This may also be why Bennett separates them, but Hume does not; even in Part 1.1 of the *Treatise*, Hume may be foreshadowing his later doctrines.

The problem is that we want to say that that expression has some sort of unity beyond a simple progression of ideas in our mind. Say there were two things in your environment, Mr. James (a cat) and Big Fido (a dog). You recognize Mr. James, so the idea MRJAMES is now before your mind; assume, also, that Mr. James is docile and not prone to biting. But you also recognize Big Fido, and also the fact that Big Fido is known to bite. So you also have before your mind the ideas BIGFIDO and BITES.

Everything is working out fine up to this point, but there seems to be a subtle race condition⁵ hidden in here. It is conceivable that the timing of these ideas coming before your mind could work out so that you think of them in the following order: BIGFIDO, MRJAMES, BITES. Now suppose that this ordering of ideas occurs on a regular basis (Mr. James and Big Fido live in the same house, say, so you usually see them together). It is true that there is a relation of causation between BIGFIDO and BITES (which is probably what we want), but it also seems that we could say there is now a relation of temporal contiguity between MRJAMES and BITES. On Fodor's interpretation of Hume's criteria it seems we now have a complex idea of MRJAMES BITES, which is definitely not a result we want.

One immediate response Hume could give to this would be to say that causation is a stronger or more binding relation than temporal contiguity. Thus even though MRJAMES is temporally contiguous with BITES, it is the causal connection between BIGFIDO and BITES that wins out. Of course, it also seems that you would need an additional rule that would only allow ideas to be constituents of a single complex idea at any given moment. It is not obvious whether or not we would want to countenance such a restriction.

But it also may not be as severe a problem as Fodor seems to think. One of the assumptions that Fodor's critique rests on is the assertion that it is the four associative relations (resemblance, spatial and temporal contiguity, and causation) that are doing the work of binding simple ideas into complex ideas. If it is not these relations that bind complex ideas together, then Hume is

⁵I am borrowing the term "race condition" from computer programming, where it refers to a class of bug where, in between two instructions of a program, some change of state external to that program causes the second instruction to do something wrong (write to the wrong file, or read from the wrong memory address, for example). In the example at hand, you can consider the perception of Big Fido and Mr. James to be separate programs, and the perception of Mr. James changes the state of your mind between your having the idea BIGFIDO and the idea BITES.

not actually faced with this problem. And to be sure, Hume does not think that these relations, stemming from the faculty of the imagination, are the only way that ideas in the mind can be connected. Speaking of the imagination, he says that “nor yet are we to conclude, that without it the mind cannot join two ideas” (*THN* 1.1.4, p.10).

Still this is only a negative response to what the relations that bind complex ideas are: they are not associative relations. But now Hume needs to give some answer as to what they *are*. This is difficult to find, and this difficulty is probably what leads Fodor to conclude (in the absence of any other good candidates) that the associative relations are the idea-binding relations.

Perhaps this is also a sort of terminological confusion. It is conceivable that both Fodor and myself are guilty of reading more into Hume’s Theory of Ideas by trying to think of complex ideas as encapsulated concepts. We should recall that Hume does say that the “uniting principle among ideas is not [...] an inseparable connexion” and is instead “a gentle force, which commonly prevails” (*THN* 1.1.4, p.10). Thus for Hume, it may be that while there are certain complex ideas that are more common due to the custom and habit of our experiences, it is the fact that these experiences are routine that makes certain complex ideas natural and common. The other way around, of course, would be to say that there are certain natural ways for us to assemble complex ideas. A strong version of this thesis would be to say that we are, in fact, constrained to certain ways of assembling complex ideas, by virtue of some facts about the capabilities of the human mind. This is most certainly not the road down which Hume wanted to go, since this way leads away from empiricism to nativism.

3 Requirements on a Theory of Complex Ideas

Fodor identifies two possible restrictions that one might place on any sort of system that generates complexes from simple inputs. (Hume’s Theory of Ideas is certainly an example of such a system.) These two requirements are *semantic productivity* and *semantic transparency*.

The requirement of semantic productivity is the requirement that the system be able to produce novel complex forms from a (likely finite) set of simple inputs. This seems like a basic requirement if we want the system in question to be able to “surpass its inputs”, so to speak. For if we cannot

produce as output anything that is not already in the set of potential inputs to the system, then there seems to be little need for the system in the first place.

Hume cannot do without the notion of semantic productivity, since he indeed wants to give an account of ideas that outstrip our experience of the world. As he puts it, “all simple ideas may be separated by the imagination and may be united again in what form it pleases” (*THN* 1.1.1, p.57). That statement actually says something further. Not only do the complex ideas produced by the imagination go beyond (in the sense of adding or elaborating on) our experience, they can also be the result of the imagination rearranging and recombining our experience.

The requirement of semantic transparency is the restriction that all of the content of the output of our system be traceable to some one or other of its inputs. That is, the process of assembling the simple inputs together does not “inject” any additional content or meaning into the output complex. If we think of the system as “gluing together” a number of simple inputs, then a semantically transparent system has invisible glue.

Another way of putting the issue is to say that in a semantically transparent system, the content of a complex idea is not just determined by its constituents, it is *exhausted* by the content of its constituents (Fodor, 2003, p.89). Because of his strong empiricist commitments, Hume wants his Theory of Ideas to be semantically transparent in this way.

3.1 Implications of Empiricism

Fodor points to Hume’s commitment to empiricism as what ultimately moves him to require that his system for assembling simple ideas into complex ideas be semantically transparent. But, Fodor is quick to point out, Hume is now in quite a bind, since it is impossible to have both a truly semantically productive system that maintains semantic transparency. To see why, we will examine the following problem case (Fodor, 2003, p.89ff.).

Take the sentence ‘John kissed Sue’, or to be more precise, the complex idea it expresses; namely, JOHN KISSED SUE. Unless we have some other rule like “agent first”, we cannot differentiate which of two possible situations this idea is supposed to (unambiguously) represent. And as soon as we have the agent-first rule, then we have to ask from whence that rule came. Now if we confine

ourselves strictly to a Humean empiricist position, it seems we must say that we learned the agent-first rule from somewhere. That seems an altogether daunting task.

If on the other hand we give up our empiricist strictures (and thus begin a departure from a solely Humean theory of ideas), then we can assume that the rule of agent-first is innately specified. Innateness is, of course, a quite broad notion, so I will try to give a more specific reading here. One plausible way of fleshing out the idea of “innately specified” is to say that there is something (given in our genetic makeup, say) about the way our minds are structured that disposes us to normally interpret “object–action–object” ideas as “agent–action–patient”.

But notice that as soon as we do that, we have not only given up strict empiricism, we have given up strict semantic transparency as well. For the whole of the meaning of JOHN KISSED SUE is now no longer contained in its simple idea constituents. We have an additional rule, agent-first, that contributes to the meaning of the complex idea from the outside.

One empiricist response to this, of course, would be to claim that the disposition to treat structures like JOHN KISSED SUE and ARTHUR HUNTS RATS as “agent–action–patient” merely comes from experience. Most transitive predicates we hear are in fact structured this way, so as children acquiring concepts, language rules, and so forth, one of the rules we learn is “agent-first”.

Another related (and to my mind, better) solution that the empiricist might offer is that part of knowing the concept KISS⁶ involves knowing that it involves a KISSER (the agent) and a KISSEE (the object or patient⁷). It might be reasonable to assume that whenever the idea KISS is before our minds, we expect there to be something that is a KISSER and something that is a KISSEE before our minds as well. But there still seems to be a stumbling block for the Humean approach. For to be sure, we want to distinguish between the complex ideas JOHN KISSED SUE and SUE KISSED JOHN.

There is another route that is open to the Hume in responding to Fodor’s criticism of the use of associationism to actually do the job of binding simple ideas together into complex ideas. Van Cleve (pers. comm.) suggests that Fodor is conflating two questions about complex ideas. The first is descriptive: “What is it to form the idea of X and the idea of Y into a complex idea $[XY]$?”

⁶Here the uninflected form is standing in for the basic concept.

⁷To borrow a term from thematic-role terminology in linguistics.

The second, however, is explanatory: “Why do I more often form the complex idea XY than XZ ?”

Fodor is critical of Hume on the grounds that associationism is not a very good answer to the descriptive question. However, it may not be the case that Hume was trying to use associationism and the various relations discussed above as in account of *what* the glue between the constituents of complex ideas is, but merely as an outline of *how* we come to associate the ideas which we do. This reading leaves open the possibility that there are other relations that serve as the actual glue between ideas and answer Hume’s descriptive question.

3.2 A Nativist Challenge

The argument against the first response is well-rehearsed in the nativist tradition, and especially in connection with linguistics. I will not attempt to give a full-fledged rationalist response to the empiricist on this point, but merely gesture at one of the stronger pieces of evidence that casts doubt on the empiricist position.

The fact is that all people converge on the same (or very similar) rules, regardless of diverse and varied experiences. This is true in language; even though the languages that different people in different communities speak appear very different on the surface, they actually share a surprising number of underlying features. And we tend to think of thought as something much more universal than language. That is, regardless of what language we speak, we have the same mental capabilities; in the particular context of this paper, the capabilities I am concerned with are the ways in which we construct complex ideas. Now, given the wide variation of experience available to different members of our species, the empiricist has a massive explanatory problem: why do we all converge on the same associative relations, or even the same relations between ideas in general?

This is actually a problem for the empiricist program in general, and thus presents difficulties not only for the first response given above, but also the second. Even if we accept (unlike Fodor) that Humean complex ideas are constructed using other relations than the associative ones, there is still the question of how we come to know these relations. And again, it seems that we all come to (more or less) the same “mental toolkit” for creating complex ideas, despite our potentially vastly different experiences.

4 Conclusion

Hume has four relations—resemblance, spatial and temporal contiguity, and causation—that he uses to associate simple ideas and form complex ideas. Fodor takes these relations to not only associate ideas, but to bind them together into discrete complexes. This, however, seems like a bit of a hasty conclusion, since it conflates the questions of what binds ideas together into complexes with the question of why certain complexes come before our minds in the way that they do. Yet I think Fodor is not altogether unwarranted in making the move that he does, since it is difficult to determine what Hume thinks the answer to that first question is.

Furthermore, there is a nativist challenge hanging over Hume's entire program. Even if we can discover what Hume takes the relations that bind together complex ideas to be, and these are distinct from his associative relations, it seems likely that these idea-binding relations will be common across people in very different environments. On strictly empiricist grounds, vastly different experience should logically lead to vastly different knowledge. So if we all converge on the same or very similar idea-binding relations, then the empiricist has some explaining to do.

References

- Bennett, Jonathan. *Learning from Six Philosophers: Volume 2*. Oxford: Oxford University Press, 2001.
- Fodor, Jerry A. *Hume Variations*. Oxford: Oxford University Press, 2003.
- Hume, David. *A Treatise of Human Nature*. Analytical Index by L.A. Selby-Bigge with text revised and notes by P.H. Nidditch. Second edition. Oxford: Oxford University Press, 1739/1978.