 Last edited Saturday, February 10, 2001

Formal Symbol Manipulation: Conclusion

I. Preliminaries

A. Today: we will complete our analysis of the FSM construal

B. Plan

1. Quickly review the extensional analysis of dependence and independence (from last time)
2. Talk about an *intensional* version of (in)dependence
3. Step back, to review and summarize the results of this first critique.

II. Review

A. Involvement

1. The main positive conclusion we have reached so far is what I have called **involvement**: the claim that computers are *involved in* (i.e., *causally engaged with*) their subject matters.
2. The question we are addressing is how this positive claim affects our understanding of the "independent of semantics" reading of formality.
3. I will argue that involvement defeats ontological formality—even on the (fine-grained) property reading. Note that it also defeated it on the object reading. That will leave us to consider the conceptual reading. We won't have time to go into the conceptual reading much today, but I will suggest that (and point to arguments as to why) it too is false.
4. Together, these conclusions imply that involvement defeats the negative ("independent of semantics") reading of formality altogether.

B. Extensional independence

1. Last time, as summarised in figure 1, we defined an **extensional** reading of independence.
2. The basic idea was that, wrt the issue of dependence of one phenomenon (α) on another (β), there were 3 possibilities:
   a. **Dependence:** α *depends on* β in the sense that fixing β fixes or uniquely determines α;
   b. **Partial Constraint:** α *partially constrains* β in the sense that fixing β narrows the range of possible variation for α, but (at least in general) doesn't entirely fix or determine it down to a unique value;
   c. **Independence:** α is *independent of* β in the sense that fixing β places no

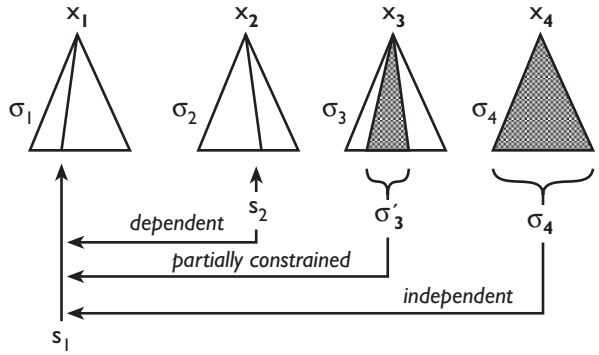

Figure 1—Dependence, independence, constraint

additional[1] constraints on α.

3. That is, there are
   a. Two limit cases
      i. **Independence:** No constraint at all (uncorrelated variation)
      ii. **Dependence:** Complete constraint (no variation at all)
   b. With a (very important) middle region between them:
      α. **Partial constraint:** Some constraint (narrowing of variation)
4. Define: **constraint** ≡ dependence ⊕ partial constraint
5. Theoretical properties

◆ **Symmetry** (if α bears ρ to β, what relation does β bear to α?)		
	a. Independence:	Independence
	b. Partial constraint:	Partial constraint (or stronger)
	c. Dependence:	Partial constraint (or stronger)
◆ **Converse** (if α does *not* bear ρ to β, what relation *does* it bear to β?)		
	a. Independence:	Constraint (i.e., partial constraint or dependence)
	b. Partial constraint	Independence or dependence
	c. Dependence:	Independence or partial constraint

6. That is
   a. Dependence and independence do not constitute a mutually exclusive and exhaustive pair. In particular, both of the following statements are *false*:
      i. The complement of dependence is independence (✘)
      ii. The complement of independence is dependence (✘)
   b. Symmetry
      i. Independence: *symmetrical*[2] (α independent of $\beta \Rightarrow \beta$ independent of α)
      ii. Constraint: *symmetrical*[3] (α constrained by $\beta \Rightarrow \beta$ constrained by α)
      iii. Dependence: a bit more complicated. If α is (extensionally) dependent on β, then β is *at least constrained* by, and *perhaps dependent* on, α.

C. Substantive results
   1. Given this formal model, we turned to substantive questions
   2. We asked: **is semantics independent of syntax?**[4]
   3. We saw that no one thinks this except quick-change artists (such as Searle), who pretty much *assume* it.[5]

---

[1] Beyond those put in place by the background coherence condition **c** (see the notes for lecture 6a).

[2] Modulo some technicalities

[3] Again, modulo some technicalities

[4] Note: this is the inverse of the main question we are interested in, of whether syntax is independent of semantics.

[5] Moreover, he assumes it because of his presumption that computation is formal. Since we are trying to figure out *whether* computation is formal, we clearly cannot just assume it.

4. We agreed to focus on what seems to be the wildly most plausible assumption: that **semantics is partially constrained by syntax** (i.e., by behaviour, potent processes, etc.).
   a. This is true on any conceivable principle of charity
   b. It is also true on any extant proposal for what semantics is like
5. That is:
   a. Semantics: ¬ independent of syntax
   b. Semantics: ¬ dependent on syntax, either
   c. Rather: relation (between semantics and syntax) is one of *partial constraint*
6. But if that is true, it follows from the symmetry of constraint (above) that **syntax (behaviour) is constrained by syntax**, too. (That is why II.C5.c is stated symmetrically.) In sum:

> ◆ On an extensional reading of independence, the property sub-reading (of the negative reading) of formality—i.e., that computational syntax and processing are independent of semantics—is false.

D. Discussion
   1. Some people—*including me*—will be dissatisfied with this result. For what it really says is that the ways in which a computer behaves (the ways in which a symbol are used) play a role in determining the semantics or interpretation of the symbols. In a sense, this is a banal result; virtually everyone (except quick-change artists) will agree to it.
   2. More seriously, the original intuition—that, somehow or other, semantic properties just aren't the sort of thing that can play a role in engendering behaviour—hasn't really been addressed. Nothing yet said seems to imply that this original intuition is false. So it seemingly remains on the table as a possibly viable reading of formality—one that might survive (in the sense of allowing the FSM construal of computing to be true).
   3. To get at it, therefore, we have to dig deeper still: and look at an *intensional* reading of dependence and independence.
   4. So today we will turn to that.
   5. In particular, I will argue three things:
      a. Partial constraint
         i. There is an *enormous region* between (pure) dependence and (pure) independence
         ii. That "middle region" (a recurring theme!) is of great theoretic importance
         iii. A version of it will continue to hold on the intensional reading.
      b. Will preserve two claims:
         i. **Involvement:** Computers are *involved* in their subject matter
         ii. **Ineffectiveness:** Semantical properties are *causally ineffective*
      c. And deny a third (constitutive of FSM)
         i. **independence:** Behavior is *independent* of semantics

III. Intensional (In)dependence

A. Preliminary remarks
   1. What follows isn't necessarily an analysis of a received or empirically warranted reading

2. Rather: an attempt to formulate a coherent view that does justice to (at least some of) the intuitions that (I think) lead people to pledge allegiance to the FSM slogan.
3. Caveat emptor ("your mileage may vary").

B. Two motivating intuitions (conditions that an intensional account of independence should meet):
   1. **Strength**
      a. It must be possible to say
         i. That β (intensionally) depends on α …
         ii. But α does not (intensionally) depend on β …
         iii. *Even if* α *and* β *are (completely) extensionally correlated.*
      b. Examples
         i. Physics: α = decay of a particle; β = emission of a gamma ray[6]
         ii. Physics: α = explosion; β = radial dissemination of a shock wave
         iii. Shadows: α = your position; β = position of your shadow
      c. In all three cases, α and β are completely correlated.[7]
      d. Yet (in each case) we would normally want to say that β is the dependent phenomenon, α the independent one.
   2. Something to do with **causation** (potency, again!)
      a. All three of the examples given above are causal examples
      b. And causation is a (famously) *directed, asymmetrical relation.*
      c. So the analysis of intensional dependence should be directed and asymmetrical as well.

C. Another example (from high school)
   1. Problem: A ferry sets out at velocity **v** across a river that is flowing at velocity **v'** > **v** (take **v** and **v'** to be constants). At what angle θ should the ferry aim, so as to have been swept a minimal distance **d** down-river when it arrives at the other side?
   2. In this example, θ and **d** are normally treated differently
      a. θ is called the *independent* variable[8]
      b. **d** is called the *dependent* variable.
   3. From an extensional point of view:
      a. θ and **d** are perfectly correlated
      b. Not only no (extensional) independence; not even any partial (extensional) constraint.
      c. *Pure dependence*—in both directions!
      d. One can as easily "solve" for θ, given **d**, as for **d**, given θ.
      e. Yet in spite of this perfect correlation, one still, intuitively, takes the "dependence" to be directed.
      f. In the real world of boats, etc., θ is what can be *controlled*; **d** is the *result*

---

[6] I know no physics. Is this plausible, Bob?

[7] Assuming stable background conditions. This may be an unrealistic assumption, though.

[8] In high school we call the *angle* and the *distance* "variables"; once we get to college, we go syntactic, and call 'θ' and '**d**' variables, syntactic items that *denote* the angle and the distance. Yes, there are reasons for having engaged in this semantic ascent—but there are prices to be paid for doing so.

D. Strategy
   1. That suggests a strategy
   2. Maybe we can
      a. Recruit causation (complete with its directedness, asymmetry, etc.),
      b. Define a notion of (intensional) dependence in terms of it; and then
      c. Define a notion of (intensional) independence in opposition
      d. Given the result, check whether a reading of FSM formulated in terms of it, holds of computation-in-the-wild
   3. Remark
      a. If this strategy succeeds, we will ultimately need a theory of causation, to back it up (cf. the conceptual criterion, way back at the beginning)
      b. For now, let's borrow an intuitive version; see if the resulting FSM thesis holds water
      c. Only if it does will it be worth trying to pay for a real copy.

E. Call this (proposal for an intensional version) **causal dependence**.. That is: causal dependence will be (at least initially) be taken to be a *species* of intensional dependence.

IV. Causal (in)dependence

A. Alas! Two kinds of causation must be discriminated
   1. *Temporal:* relation between something happening and its "subsequent consequences in time"
   2. *A-temporal:* relation between something at a lower level of description, and *the same thing* at a higher level
      a. I.e., implementation, realisation, constitution
      b. E.g., RMS velocity of molecules, and heat
      c. E.g., C++ code and Scheme interpreter
      d. E.g., voltages and currents in transistors, and "being a Java applet"
      e. Our old friend: abstraction or implementation boundaries
   3. Vocabulary
      a. Some people may not want to use the word 'cause' for implementation, preferring:
         i. 'Cause' for the temporal case
         ii. 'Constitution' for the a-temporal case (or 'implementation')
      b. I am sympathetic to such inclinations. But it won't matter; here.
      c. For now it is just important to track the two cases separately.
      d. So I will just call them the "temporal case" and "a-temporal" case.

B. Dependence
   1. Intuitively, what do we want to say about (intensional) dependence in these two cases?
   2. Temporal case is relatively easy
      a. Effects *depend* on the cause (✔)
      b. Cause proceeds *independently* of the effect (~✔)
   3. A-temporal case is worse
      a. Implemented *depends* on the underlying implementation (~✔)
      b. Implementation is *independent* of what is implemented (?✘)

4. In general (or so I would argue), it is bizarre to say:
   a. If β is implemented on (or constituted out of) α, then …
   b. α is *independent* of β.
5. That is not to say that α is *dependent* on β, either.
6. Neither category applies!
7. Example:
   a. The way your brain works is independent of your mind (✘)
   b. The way your brain works is dependent on your mind (✘)
   c. C++ code implementing Scheme works independently of the Scheme code that it implements (✘)
   d. C++ code implementing Scheme depends on the Scheme code that it implements (✘)

C. Mereology
   1. The moral is quite general: it applies to other forms of atemporal constitution as well.
   2. Also applies when something is a *part* of something else
   3. Example · 1
      a. Two possibilities
         i. Your body works independently of your arm (✘)
         ii. Your arm works independently of your body (✘)
      b. The only way to interpret either is gruesome
   4. Example · 2
      a. Assume that seven stars, $\mathbf{x}_1$ through $\mathbf{x}_7$, are in orbiting around each other, with some resulting overall center of gravity
      b. Again, neither possibility works:
         i. The orbit of the third star, $\mathbf{x}_3$, *depends* on the center of gravity (✘)
         ii. The orbit of the third star, $\mathbf{x}_3$, is *independent* of the center of gravity (definitely ✘)

D. Summarize this as a lemma

Constitutional Dependence

If β is *made out of* α , in either a mereological (part-whole) or implementational (constitutive) sense, then

1. There is a sense in which β depends on α
2. But it's wrong (category error?) to say that α is *independent* of β.
3. It is also false to say that α is *dependent* on β
4. That is: α is neither independent of β, nor dependent on β
5. α *implements* or *(partially) constitutes* β
6. That's all there is to say!

E. Knowledge
   1. Could say: that one can *give an account* of α independent of an account of β.
   2. This goes back to the conceptual reading: that an *account* of how a computer system works

can be given *independently of an account* of its semantics.

3. Some kind of *logical* independence? Which cuts more finely (i.e., allows a broader set of possibilities) than mere nomological independence? or metaphysical independence?

V. Computation

A. Back to the relation between syntax and semantics

B. Constitution

1. We saw, in the analysis of extensional (in)dependence, that what we are calling "syntax,"—i.e., the overall behaviour of a computational system, at least in the general case—can be **partially constitutive** of its semantic value.
2. That is: the "constraints" liking behavior and interpretation are *constitutive* (atemporal), not *temporally consequential* (temporal)
3. Referring to Pluto is not a *temporal consequence* of a certain pattern of inferential use.
4. That implies that it is the *atemporal forms of dependence* that are relevant (why we've spent so much time on them)

C. Given all of this, let's go back to original question (one last time!). Finally, we have some answers:

1. Is syntax (behaviour) intensionally dependent on semantics? **No!**
2. Is syntax (behaviour) intensionally independent of semantics? **No, not that either!**

D. In sum

> ◆ Syntax (and effective behaviour) are neither intensionally dependent on, nor intensionally independent of, semantics, because it (syntax, behavior) is *partially constitutive* of semantics

E. So the "participatory moral" of the last few weeks has a role to play here, as well (which is why we spent so much time on *it*)

F. Q.E.D.

VI. Causal efficacy

A. Status

1. Are we done? Of course not!
2. There are still two things to be done:
   a. We still haven't made good on what is *right* about the FSM construal
   b. We haven't assessed the adequacy of the conceptual reading
3. Look at the former question first (get to the substantive heart of the FSM construal)

B. Have already admitted the main issue:

> ◆ **Remaining Question:** Are semantic properties effective?

C. Discussion

1. Background
   a. Implicit in the FSM view that semantic properties are not effective.
   b. We have been exploring a participatory view
   c. So you might expect an argument that they *are*.
2. But not so fast.
   a. Paradigmatic examples of FSM ("Plato and Pluto") are *real computation*
   b. Throughout, what we've argued is that they are not *paradigmatic* of computation (i.e., aren't examples on which the whole theoretical edifice should rest)
   c. What is distinctive about them is the great separation they exemplify between the computational realm and the task domain (realm of the subject matter)
   d. In arguing for a participatory view, we have tried to bring the "syntactic and semantic" realms back together.
   e. But we don't want to bring them *too far back together!*
   f. Remember: we already argued (several weeks ago) against semantic implementation (which is effectively the idea that they are brought *completely back*)
   g. Instead, the goal is to bring the semantic realm *partly back*
   h. Yet another "middle ground" conclusion![9]
3. Sometimes ineffective
   a. If Plato/Pluto cases are computational, semantic properties can't *always* be effective.
   b. I.e., it is not in the nature of "being semantic" to be effective
   c. That is guaranteed by disconnection and semantic reach
   d. So part of the answer to the question (whether semantic properties are effective) is …
   e. … *not always!*

D. Remaining question: are semantic properties *ever* effective?
   1. We can't say for sure, yet, of course, because we *don't yet have a theory of semantics*
   2. Many people would say *yes*
      a. Most proposed semantic theories in philosophy of mind suggest so.
      b. E.g. Dretske's "Putting Content to Work": if content didn't make a difference, a *causal* difference, then it wouldn't be naturalised (i.e., would be epiphenomenal, or something like that).
      c. Many people think that perception is a case of counter-factually supporting causal coupling, in virtue of which the content of the perception is established (which is a strong enough to imply that in that case, semantic property is effective).
      d. Some people seem to have an almost a priori commitment to a causal analysis of content (cf. Fodor's asymmetric dependence theory).
   3. Note, however, that
      a. *Nothing we have said so far* (about the participatory nature of computing, about the cross-cutting boundaries, about the falsehood of the ontological readings of the FSM claim, etc.) *implies that semantic properties* are *effective*

---

[9]These "middle ground" conclusions are going to figure very importantly in the ultimate metaphysical view to which this whole analysis leads.

b. My own view is that they are *never* effective.
c. I.e., I subscribe to:

> ◆ **Impotence:** Semantical properties are *never* causally effective

d. But that has not yet been argued.[10]

VII. The Conceptual Reading

A. Intro
   1. One more thing to do: assess the conceptual reading: whether an *account* of a computational system's syntactic (causal, behavioural) aspect can be given independent of an *account* of its semantics.
   2. Assuming physicalism is true, there is one sense in which this must be true: That is: given any particular computer system, it (at least arguably) must be possible to give an account of *it* at the level of quarks, fermions, etc.
   3. On the other hand, that is not what is being asked. The question is not whether a *lower-level* account of a computer can be given, independent of the system's semantics, but whether one can give an account of it *as computational* independent of its semantics.

B. There are two reasons to suppose that the answer is no.
   1. Internal functionalism
      a. Remember that computations are normally individuated *functionally*. The idea is that one does not specify the precise physical or material make-up of the system (which could vary), but rather the sorts of functional and causal role that its operations play.
      b. Take that functionalist stance together with the morals extracted in the "unbundling transducer" section, where we saw that there were lots of *internally exemplified semantic relations*(such as counting).
      c. If there are genuinely internally exemplified semantic behaviours, then functional characterisation will *include semantic characterisation* (e.g., in the case of counting).
      d. It may not be characterised *as semantic*. But if it is not, then it is not being specified at the requisite level of abstraction.
   2. Ontology
      a. Another reason to suspect an implication of semantics in the functional characterisation of a computing system's effective or behavioural roles has to do with *ontology*.
      b. There is a question as to how the individuation criteria are established for internal structures (e.g., what it is to be an instance of one or another data structure).
      c. In the full AOS series, I will eventually want to argue that semantic considerations play a role in individuating the internal structures (such as data structures, types, etc.)
      d. We can't not ready to do that here, though, because we are not yet equipped to take on the warrant for ontology.

C. At a minimum, it seems, the buck is passed back to someone who wants to argue that semantic considerations do *not* apply in the individuation of internal functional types.

---

[10]For an argument to that effect—or a proposal according to which they are not—*On the Origin of Objects* (1996)

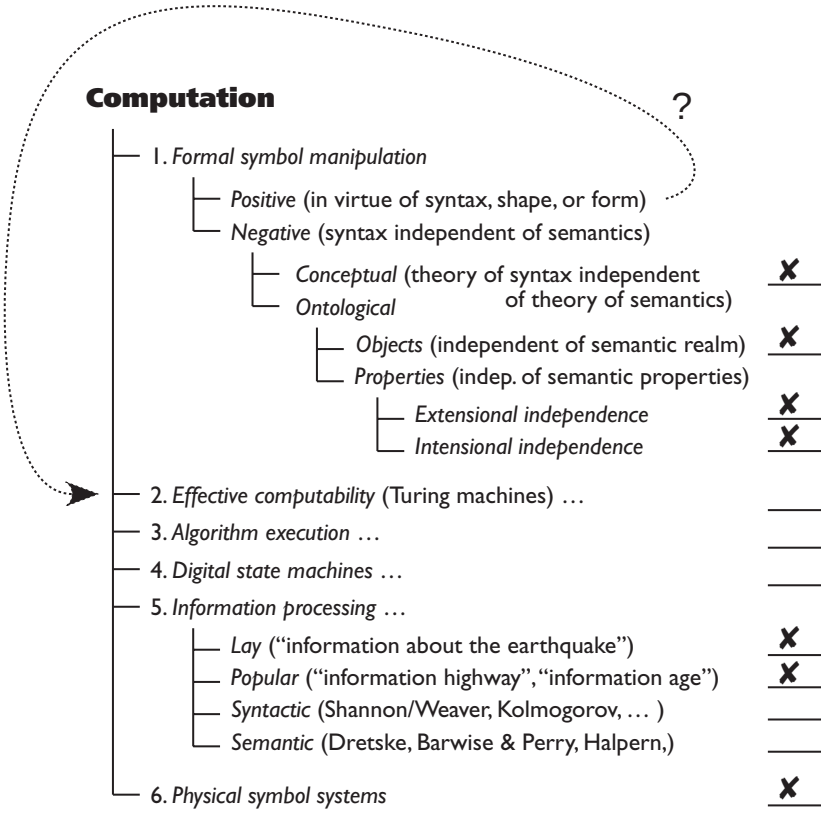

Figure 2 — Investigation Situation at the end of the First Critique

VIII. Conclusion

A. We are done with the first construal!
   1. The situation isn't as complete as we'd like, given:
      a. The lack of an account of semantics, in terms of which to formulate the "independent of semantics" readings, and
      b. Our inability to spend much time on the conceptual reading
   2. Still, the basic picture is that we have not found a negative (independent of semantics) reading of formal that is true of computation in the wild. So the current situation is as depicted in figure 2, above.

B. Draw out some morals
   1. Participation or **involvement**
      a. Computing is a participatory (engaged) intentional (semantic) phenomenon
      b. Cf. half a dozen cross-cutting boundaries (semantic, physical, abstractness, etc.)
   2. The Partial Constitution of Semantics by Behaviour
      a. No (negative) reading of the formal symbol manipulation (FSM) construal could be found that met our criteria
      b. Because of computation's participatory nature, syntax and behavior are not independent of interpretation (semantics)
      c. They are not *dependent* on interpretation, either; that much remains true.

d. Dependence and independence are *extremes*, neither of which apply
e. As with many other things in this story, the right answer is **in the middle**—a thick, textured, middle ground
f. Behaviour *partially constitutes, and partially constrains, semantics*

3. Semantics' Inefficacy

a. No reason yet been seen to argue that semantic properties are effective
   i. Must *sometimes* be ineffective (cf. Pluto and Plato)
   ii. May *always* be ineffective (for all we know)
b. There are good reasons to believe, however, that—in spite of the falsehood of all the interesting independence claims—that, nevertheless, semantics is always ineffective.

4. Effectiveness

a. Effectiveness (potency) matters. It is where all the work gets done.
b. The effective dimension of computing is absolutely central.
c. But we have no theory of it yet. Cf. all the confusion around *syntax*, *shape*, *form*, *cause*, *effective*, *intrinsic*, *local*, *physical*, *grammatical*, etc.—i.e., around the potency predicates.

5. Special?

a. The "formality" of computing was thought to indicate something special about computational processes—something implying that they were "easier" to understand, or at least more evidently naturalisable.
b. We have found no such thing!
c. All the results (participation, cross-cutting boundaries, partial constitution of interpretation by behavior, etc.) seem evidently true of people. *(This is a very important point.)*
d. So what makes computers special? Does anything?
e. We certainly haven't found anything yet.

C. Time to turn to the second construal.

—— *end of file* ——