Formality & Separation

Brian Cantwell Smith
Xerox Palo Alto Research Center
and Stanford University

ABSTRACT

1. PREAMBLE

Just about everyone thinks that computers are *formal* — that they manipulate symbols formally, that programs specify formal procedures, that data structures are a kind of formalism, that computational phenemona are uniquely suited for analysis by formal methods. In fact the computer is often viewed as the crowning achievement of an entire "formal tradition" — an intellectual orientation, reaching back through Gallileo to Plato, that has been epitomised in this century in the logic and metamathematics of Frege, Russell, Whitehead, Carnap, Turing, etc.

This history would suggest that formality is an essential aspect of computation. But I do not believe this is right. For one thing, it is not clear what the allegiance to formality is an allegiance to. It is not as if "formal" is a technical or theory-internal predicate, after all. People may believe that developing an idea means formalising

---

3333 Coyote Hill Road, Palo Alto CA 94304. This paper presents, in very abbreviated form, some of the ideas developed in a forthcoming book, *The Middle Distance*, on the foundations of computation and intentionality.

 bcsmith@parc.xerox.com

it, and that programming languages are formal languages, and that theorem provers operate on formal axioms — but they do not write FORMAL(X) in their daily equations. Moreover, a raft of different meanings and connotations lie just below the surface. Far from hurting, this apparent ambiguity has helped to cement popular consensus. Freed of the need to be strictly defined, formality has been able to serve as a lightning rod for a cluster of ontological assumptions, methodological commitments, and social and historical biases.

Because it is tacit, goes deep, has historical roots, and permeates practice, formality is an ideal foil with which to investigate computation. It has proved especially useful in my own search for a satisfying theory of computation — one that is able (i) to explain and inspire computational practice, and (ii) to serve as a tenable foundation for cognitive science. The second goal has been a special focus. There has been a great deal of debate, over the years, about whether it is appropriate (correct, illuminating, moral, etc.) to understand people in computational terms. I have never known what to make of such discussions, because I have been unable to figure out what computation is. The one thing I am sure of is that all our current theories are inadequate — which makes cognitive discussions premature (it is surely irrelevant to ask whether people are correctly characterised by a notion of computation that is not even true of Microsoft Word). The notion of formality, furthermore, is partly responsible for these conceptual inadequacies.

2. READINGS OF FORMALITY

Almost a dozen different readings of "formal" can be gleaned from informal usage: *precise*, *abstract*, *mathematical*, *a-*

contextual, *digital*, *explicit*, *syntactic*, *non-semantic*, etc. They are alike in foisting recalcitrant theoretical issues onto center stage. Consider explicitness, for example, of the sort that might explain such a sentence as "for theoretical purposes we should lay out our tacit assumptions in a formal representation." Not only has explicitness — and its partner in crime, implicitness — stubbornly resisted theoretical analysis, but both notions are parasitic[2] on something else we do not understand: general representation. Or consider "a-contextual." Where is an overall theory of context in terms of which to understand what it would be to say of something (a logical representation, say) that it was not contextually dependent?

Considerations like this suggest that particular readings of formality can be most helpfully pursued within the context of the general theoretical edifices that have been constructed (more or less explicitly) in their terms. Five such readings are particularly important:

F1. An *antisemantical* reading: the idea that a symbolic structure (representation, language, symbol system, etc.) is formal just

---

. At one stage I asked people what they thought "formal" meant — not just computer scientists, but also mathematicians, physicists, sociologists, etc. It was clear from the replies that the term has very different connotations in different fields. Some mathematicians and logicians, for example, take it to be perjorative, in contrast to the majority of theoretical computer scientists, for whom it has almost the opposite connotation.

. On its own, an eggplant cannot exactly be either formal or explicit, at least not in its ordinary culinary role, since in that role it is not a representation at all. In fact the only way to make sense of calling something non-representational explicit is as short-hand for saying that it is explicitly represented (e.g., calling eggplant an explicit ingredient of moussaka as a way of saying that the recipe for moussaka mentions eggplant explicitly).

in case it is manipulated *independent of its semantics*. Paradigmatic cases include so-called formal logic, in which it is assumed that a theorem (such as MORTAL(SOCRATES)) is derived by an automatic inference regimen without regard to the reference, truth, or even meaning of any of its premises.

F2. A closely-allied grammatical or *syntactic* reading, illustrated in such a sentence as "inference rules are defined in terms of the *formal* properties of expressions." Note that whereas the antisemantical reading is negatively characterised, this one has a positive sense.

F3. A reading meaning something like *determinate* or *well-defined* — i.e., as ruling out all ambiguity and vagueness. This construal turns out to be related to a variant of the computationally familiar notion of digitality or discreteness.

F4. A construal of "formal" as essentially equivalent to *mathematical*.

F5. A reading that cross-cuts F1–F4: formality as applied to analyses or *methods*, perhaps with a derivative ontological implication that some subject matters — such as computation, perhaps? — are uniquely suited to such analytic techniques.

There are many things to say about these five. The first two, for example, are often treated as conceptually equivalent, but to do that is to assume that a system's syntactic and semantic properties[3] are *necessarily disjoint* — which is almost certainly false. The relationship between the third (determinate) reading and digitality does not have so much to do with what Haugeland calls "first order digitality": the ordinary assumption that a system's states can be partitioned into a determinite set, such as that its future behaviour

---

. They may in fact be disjoint (although I do not believe it), but that is an empirical claim, which requires, in order to be coherent, that "syntactic" and "antisemantical" *mean* different things. That is what equating F1 and F2 attempts to deny.

or essence stems solely from membership in one element of that set, without any ambiguity or matter of degree — i.e., the sense parodied in the "clunk, clunk, clunk" of a 1950's cartoon robot. Rather, vagueness and indefiniteness (not simple continuity!) are excluded by a "second-order" form of digitality — digitality at the level of concepts, in the sense of there being a binary "yes/no" fact of the matter about whether any given situation falls under (or is correctly classified in terms of) the given concept.[4] And finally, the fourth view — that to be formal has something to do with being mathematical, or at least with being mathematically characterisable — occupies something of an ontological middle-realm between the subject-matter orientation of F1–F3 and the methodological orientation of F5.

The moral for computer and cognitive science is that *not one of these readings correctly applies to the computational case*. It can never be absolutely proved that computation is not formal, of course,[5] given that the notion of formality is not determinately tied down. But it can be shown that no standard construal of formality, including any of those enumerated at the head of this section, is both substantive and true of extant computational practice. Some readings reduce to vacuity, or to no more than physical realisability; others break down in internal contradiction; others

---

. See Haugeland's "Analog and Analog" [ref].

. Nor does it help to tie it down, artificially, in an attempt to prove (or deny) such a conclusion. It does not help because the entire weight of the formality thesis rests on the full-blooded, social, historical, notion — the cluster of ideas, derived from several centuries of intellectual work, that infects both theory and practice. The best that would be accomplished by *defining* formality would be to shift the burden of argument from whether computation is formal to whether the proposed definition of formality correctly captured what the tradition has meant by the term — which it almost surely would not. *Formality* is simply not a formal notion.

survive the test of being substantial, but are demonstrably false, even of current systems. One is inescapably led to the following conclusion, in all its historical irony: that the computer, darling child of the formal tradition, outstrips the bounds of the very tradition that gave rise to it.

3. SHARP-EDGED BOUNDARIES

Return, though, to formality, and to the claim that a common orientation unites the various readings. There is no single core notion to be found — the commonality is one of a family resemblance — but there are shared themes, ranging from the etymologically indicated sense of *form* or shape, to social issues of norms and authority (e.g., in a "formal" invitation). The theme that runs deepest, however, is one of being *cut* — severed, pulled wholly apart, completely disconnected. The formal world is a crystalline vision, of separated entities, sufficient unto themselves, and utterly dissociated from any distracting influence of history, context, meaning, environment, interpretation, indefiniteness, or use. Formal objects are neatly cut off from anything "other" that might gradually or messily impinge on the purity of their existence. It is the absoluteness and sharpness of the separation, furthermore, that uniquely characterises the formal — and that as a consequence distinguishes it from the real. Formality, that is to say, is almost diametrically opposed to what Haugeland has called the "vagaries and vicissitudes, the noise and drift, of earthly existence."

What emerges from this picture is the recognition that formality is not, ultimately, an independent mode of existence, but rather a *way* of being separated or differentiated. "Formal," i.e., is what philosophers would call a second order property — in this case a claim that some prior distinction or given separation is clean, sharp, and unambiguous. We have seen a number of such distinctions already — between a situation and its context, between

a sign and what it signifies, between a theorist and his or her subject matter. Formality is not choosy about what distinction it applies to, which partly explains why there are so many different readings of the term. Rather, it works as follows: given any candidate differentation, formality comes along with an absolute and unambiguous knife, and cuts what might otherwise be gradual or partial into two entirely separate ontological realms.

The resulting sharp edges can be seen in all the readings mentioned above. Sense F3, having to do with concept boundaries, is perhaps the easiest. In this case the background distinction is simply that between what does and what does not fall under a given concept or predicate. What makes a concept formal, on this reading, is the imposition of the additional, second-order assumption: that this boundary, this figure-ground separation inherent in any act of categorisation, is a black and white affair — no matters of ambiguity, or degree, or gradual falling away. So the notion of being an integer would be considered formal in this sense — as opposed, say, to being happy, or being small — since things that are integers are integers absolutely, and things that are not integers are just as absolutely not. Note that mathematical continuity is precise (and hence formal) on this reading; sin(x) may vary smoothly, but *whether* 0.454 is the sin of 27° is a question with a definite yes/no answer.

A similar analysis explains the "abstract" construal, on which a formal object is viewed as completely independent of concrete realisation or earthly existence. Here the background distinction is

---

. Equally clean boundaries separate one formal concept from another. Physics supports continuous measurements, for example, but not continuous concepts — and so we do not encounter things that are mostly a mass, but somewhat a force, with a little acceleration folded in. In being discretely separated in this sense, mechanical notions are presumably different from such psychological ones as self-confidence, pride, egotism, and hutzpah — exactly the sort of concepts one would expect to have difficulty accomodating within a *formal* psychology.

that of abstraction from complete physical detail — a notion whose original sense, as predicted, is gradual. To say of something that it is a room in a house, for example, is to leave aside some implementation details (the material of the walls, the organisation of the paint molecules). But since it does not abstract away from everything, it does not render the kitchen into an immaterial Platonic object. Only when the second-order assumption is added, with its commitment to a total and inviolate causal moat between the abstract and the physical, do we get the formalist distinction between earth and (Platonic) heaven. Abstract formality, that is to say, lets go not only of detail but also of occurrence. It is the ultimate imposition of Cartesian dualism, with the connection between abstract and material existence not just weak or mysterious, but completely sundered.

As a third example, consider the antisemantical reading — the construal that underlies the claim that mathematics and logic are formal disciplines. This case is worth examining in a little detail, since it cuts close to the bone of a reigning view of computation. Note, first, that all intentional processes (thinking, speaking, representing, etc.) proceed in at least relative causal independence of the states of affairs they are about. I can think about Sri Lanka — or about the long-dead Pharoahs of Egypt, or about events outside my light-cone — without any direct causal involvement of or connection to the situations I am referring to. Disconnection is a fact of great practical consequence; it explains, for example, why NSA cannot build a meter, to be placed in the bowels of Cheyenne Mountain, that would register when the control room was the subject of a remote intentional act. On the positive side, disconnection enables fiction, hypotheticals, and fanstasy lives; without it you would not be able to even think about continental drift without dragging the tectonic plates with you. The relative causal independence of an intentional action and its subject matter, in sum, is a precondition to our ability to refer to situations removed from us in time, space, or possibility.

Since intentional disconnection is a familiar aspect of human life, it is not equivalent to antisemantical formality.[7] And this is exactly what we should predict. What distinguishes the (alleged) antisemantical formality of mathematical logic is the severity, the absoluteness, of the separation between sign and signified — a completeness of separation that is foreign to human disconnection. The manipulation of logical axioms, for example, allegedly proceeds in complete independence of what they are about, in a way that human thinking does not. As well as thinking about the coffee cup on my desk, I can reach out, touch, and even drink from it. Not only that; my abilities to interact with the cup in these causal ways interpenetrates my abilities to refer to it (e.g., through indexical or demonstrative reference). In general, human intentionality involves a coordinated blend of more and less engagement with the world, implying that human disconnection is both gradual and partial. It is the contrasting completeness of separation of a set of logical axioms from their semantic interpretation that leads us to brand them as formal.

A notion of complete separation can also be seen in the positive (syntactic) reading of "formal" — the one with the most obvious direct connection to an underlying notion of *form*. The story is complicated, but the diagnosis is simple: formality, on this syntactic or "effective" view, amounts to neither more nor less than the projection into an abstract realm of the positive constraints of physical embodiment. First, as in the case of the "abstract" reading, there is assumed to be a complete gulf between the relevant linguistic or representational structures (this is the notion

---

. If disconnection did imply (antisemantical) formality, it would be *obvious* that people were formal, which it is not. This is not to say that people are not formal. Cognitive science may teach us, someday, that they are. For example, we will be forced to accept this conclusion if it turns out both that people are computational, and that computation is formal symbol manipulation. But such a result would be substantive, in part, in proportion to its being surprising.

of formal that is most relevant to our conception of a *formalism*) and issues of physical realisation or embodiment. But in fact the totality of this separation is illusory; the influence of the physical world cannot be so easily dismissed. Even in intentional realms, what is ultimately possible is constrained by ineliminable physical facts, whether or not our theoretical edifices try to abstract away from such earthly considerations. And so some property, within the abstract realm, is needed to represent what can and cannot be "done." It is this role of serving as an abstract witness to the physical that (syntactic) formality is asked to play.[8]

Analogous but more complicated stories can be told about the other readings. The mathematical construal, for example, is a compound one, building on several of the others. Not only does it involve systematically setting aside context, embodiment, and semantics (the last on only one of several possible sub-construals — mathematical formality is itself a family of notions), but it also relies on sharp conceptual boundaries of the sort identified with second-order digitality (these are relevant to its interest in precision). Finally, the notion of a formal method introduces yet another boundary, between us as theorists and the subjects of our investigations. What formality's imposition of complete separation adds, in this situation, is exactly what is required in order to support the myth of a perfectly neutral, non-invasive, observer.

4. MAINTAINING SEPARATIONS

One question remains: why the allegiance to formality? What has it all been for? No intellectual stance could have corralled such commitment, or affected so much history, unless it addressed a long-standing and genuine need.

---

[8]. Syntactic formality is formulated at such a level of abstraction as to be intentionally (semantically) coherent, of course. The *subject matter* of syntax is physical embodiment, i.e., but the operative *conceptual scheme* is semantically individuated.

To answer this, we need a better sense of the general role of differentiation. Note, first, that separation is a necessary good. Intellectual life does not consist in assembling the world into an undifferentiated whole. The task, rather, is to take it apart, even if only to say of the resulting pieces that they are inextricably related. Finding regularities, uncovering principles, ascertaining essential structure — all these things require a degree of pulling apart, of distinguishing, of making room, of laying out ingredients or aspects or constitutive pieces. This is why theorising inevitably involves doing a certain amount of violence to the world; the necessary disentangling unavoidably tears up what, in reality, is primordially whole.

It is surprisingly difficult, moreover, to maintain an appropriate amount of separation. Consider just three examples. First, one of science's primary ways of making sense of the world is by distinguishing phenomena of different types — separating phenomena that fall under one concept from those that fall under another, as well as separating what falls under any given concept from what does not. What makes this difficult is the job of finding stable decompositions that can be maintained in the face of bewilderingly complex surface evidence. In fact the very notion of a type rests, foundationally, on this sense of stably partitioning the world into relatively independent aspects. A counter-example can thus be viewed as something that collapses a proposed partition.

As a second example, consider the separation between a sign and what it signifies. This distinction is also harder to maintain than it might seem. When I point at a computer screen, what am I referring to — a pixel, a bit map, a serif, a letter, a word, the file named by that word, the contents of that file? Do all these things represent each other? Or are they all coreferential? Is there even a fact of the matter? Situations like this are so semantically complex that, in practice, people tend to collapse different levels together. Note, for example, how easy it is to think that a real number (as opposed to an isomorphic representation) can be stored in a

computer's memory. Or how readily people will identify a real-world phenomenon (such as a distance) with an abstract mathematical model of it (such as an ordered pair of a number and a unit). And yet the main lesson to be learned from the history of logic, model theory, and semantics, is that it is crucial, especially when things get complicated, to keep distinct representational layers separate.

Or consider, as a third example, how hard it is, when developing theoretical accounts, to leave out what is particular, gratuitous, or inappropriately subjective. It is a general truth, of course, that theories and theorists are both measured by the extent to which they transcend the idiosyncratic details of their own individual histories. This is why we employ techniques ranging from control experiments to peer review; it is all part of an attempt to maintain yet another separation: between what is idiosyncratic and what is universal.

It is because of these and other such difficulties[9] that formality has held such sway. In fact formality, in its sharp-edged sense, is nothing more than a particularly blunt way of addressing the problem of appropriate separation. If you are trying to pull things apart, and do not know any better, it is natural to pull them apart completely, as a way of ensuring that they do not collapse together again behind your back. And it works. The formal tradition has been spectacularly successful, for example in its ability to deal with combinatoric complexity.

But there is a problem. On analysis, it turns out that a great many distinctions that have traditionally been analysed formally

---

. It may seem that all the different distinctions being considered (figure-ground, conceptual orthogonality, sign-signified, subject-object, situation-context) are so different that it is either a pun, vacuous, or downright misleading to collect them together under a single word "separation." In fact, however, I believe almost exactly the opposite: not just that they are all related, but that they reflect the same underlying metaphysical variagation.

are, in fact, *essentially gradual*. In each case, it is crucial to the phenomenon in question that various relevant aspects of the world be *partially separated*, but *not pulled apart completely*. By analogy, think of the metaphysical substrate as a kind of clay or silly putty. Each phenomenon requires that the clay be pulled and stretched, in order to create shapes, make room, allow for motion, establish creative tension. The problem with formality is that it pulls the putty apart so much that it actually breaks. The resulting system may be simple, and the disconnection pure, but those gains are achieved at untenable cost. Rather than yielding a productively structured whole, the process breaks the world into a collection of irreconcilable fragments.

All three separations mentioned earlier in this section have this property of being essentially gradual: concept boundaries, semantic relations of interpretation (between sign and signified), and the appropriate stance of a theorist towards a subject matter. As it happens, all three are also implicated in the foundations of computation. Not only does this imply that computation is not formal; it also explains why our methodological allegiance to formality has blocked us from understanding what computation is. The problem, in each case, is that the essence of the phenomenon lies in the "textured middle" of the range of differentiation.

To get the sense of this, consider another analogy, to black and white photographs. Imagine someone arguing that although in practice one encounters all kinds of middle level greys, the proper scientific idealisation should impose a formal binary split, separating pure blackness from pure whiteness. Such a person would claim that the *theory* of photographs should comprise a theory of blackness and a theory of whiteness, with everything in between considered messy, extra-theoretical — perhaps only of engineering concern. But of course this is all absurd; in fact photography deals essentially with middle-level phenomena in their own right — texture, shadows, specular reflection, etc.

A similar moral holds for each of the three distinctions under discussion. Reference, for example, requires that concept boundaries be flexible — in order to make sense, for example, of the term "water" in the question "Is there any water in the refrigerator?."[10] Theoretical inquiry similarly requires a flexible (non-formal) degree of engaged participation in the subject matter. But let me set those two aside, and close with a brief look at the third separation, between sign and signified.

A particular form of separation plays an essential role in intentionality, more general than the difference between syntax and semantics in logic, but more specific than mere difference *per se*. It is crucial — in fact it is partly constitutive — for the semantical to transcend or outstrip the purely effective or causal. In fact that is what semantics and intentionality are for: to enable a located and physically embodied agent to reach beyond its immediate causal surround, in order to be coordinated with at least a fraction of the wider world.

It turns out, on analysis, that this separation is essentially partial. Some causal engagement is required, for example in order for an agent to register its surrounding situation in terms of objects, properties, and relations. The reason is that the act of abstraction that underlies classification and categorisation depends on the maintenance of a kind of "middle distance." The situation is a little like that of visual focus: you cannot see something that is plastered against your eyeball, and you cannot see it if it is infinitely far away. It must be somewhere in between. Similarly, in order for an agent to take an object as an object, some separation is necessary, but only some — enough for the object's constitutive spatio-temporal continuity to decouple from the agent's, but not too much, so that the agent can interact with the object enough to establish the objective relationship. Similar degrees of partial engagement are required in order to coordinate intentional and

---

. This example is from Winograd [ref].

physical activities (so that the agent could know, for example, as soon as it turned around to find something lying behind it, that it was now in front). Overall, this coordination is reminiscent of a dance. If you were glued to your partner, you could not dance with them. If you were to run into the next room and lock the door, you could not dance with them either. You need to be just appropriately apart.

No one of these examples, on its own, would be sufficient to indict formality. One might imagine modelling the notion of middle distance, for example, with a formal, continuous quantity. The ultimate problem, however, is that the gradualism and partiality go all the way down. There simply is no cleanly-chopped formal bedrock on top of which to construct such formal models. As a consequence, we are faced with the following challenge, if we are interested in moving past formality: of learning how to disassemble a whole but partially separable world, replacing formality's over-zealous knife-edge cuts with more appropriate forms of distance, separation, and detachment.

—— *end of file* —— ��