More than Mere Mechanism
(≈ GA • 3.5)

1. Introduction (~200)

The question of whether computers can be intelligent, conscious, or responsible—have dignity, be our friends, be religious—depends on the prior, fundamental question of *what a computer is*. Some people think that the answer is *given*: that we know what computing is in advance. They assume computers are (i) mere machines, (ii) nothing but zeros and ones, (iii) pure information processors. But how do they know? Do they do double-blind experiments? I doubt it. No, people think they know what computers are *because we build them*—on the assumption that what we build, we understand. But that is surely false. Lots of human creations defy understanding: like cities, smog, government. And lots of human products aren't machines: dinner, works of art—and children. Why should computers be different?

I worry that a priori ideas about what computers are like—digital, abstract, purely mechanistic—may reflect the pride or prejudice of their creators. We need to set presumption aside, and take the foundational question to be substantive, and empirical.

2. Personal history (~300)

I first got interested in these issues—of what computers are—more than 30 years ago, during my freshman year of college. Like others of my generation, I was riven by the tension between science and the humanities. Though I was majoring in physics, my heart—and my friends—were in politics, philosophy, the arts. In the midst of this turmoil, a computer was delivered into the basement of the physics building. I asked my physics professor for 6 weeks off. I wanted to know, I told him, whether computers—whether the *notion* of computing—could satisfy two simultaneous criteria:

i. Be understood w/ rigour, precision, depth of natural science; and
ii. Do justice to the full complexity of the human condition—i.e., to the full social, ethical, political, erotic, cultural dimensions—without the inhumanity and (it seemed to me) dessicating reductionism of all prior scientific attempts.

That is, I wanted to know whether the advent of computing could help heal C.P.Snow's two-culture gap—the gap that was tearing me apart.

Not a bad question. But I was off on the timing. I never got back to physics. Only now, 30 years later, have I finally figured out what computers are.

I'm not simply going to tell you the answer. That would be too easy. (Maybe we can come back to it in discussion.) Rather, what I want to do is to document the history of our un-

---

They are about the *fundamental nature of computing*—about what Turing discovered, about the 20th century's greatest invention, about a "way of being" that evolution may (or may not) have stumbled on, many centuries earlier.

derstanding of computing: where it came from, where it is going, what it signifies for the issues of this conference. In particular, I want to defend the following strong claim: that the advent of computing represents a transformation, of unprecendented proportion, in our understanding of nature, the world, and science—a transformation so major that the question of the relation between science and religion must be entirely rethought.

3. History (~500)

Go back to the 17th century, to the rise of natural science. As we all know, science valorized rationalism and empiricism, in place of the doctrinal focus of a purely religious society. It also focused on *matter, materials, mechanism*—explaining the *physical* world, in terms of *causal* mechanisms, understood from a detached, "objective," 3rd person stance. (It was also preceded by hundreds years of *alchemy*: a rag-tag bunch of perverse, untheorized, pragmatic practices, emphasizing magic, intuition, and first-person, embodied stances—all issues of relevance to us today. Alchemy was shunned, of course, once Newton, Descartes, and the rest replaced its occult rituals with the rigorous epistemological edifices we know today. But alchemists are now recognized as much more interesting—thicker, more savvy, etc.—than once thought. As we'll came back to, they were a crucial prerequisite to the emergence of science as we know it.)

Some people view computers as the product of this scientific world view. But it is not so! At least not in any simple way. Computers emerged from a subsidiary, cross-cutting strand, having to do not with the concrete physical world, but with the abstract world of language and logic. In the mid-19th century, in particular, Boolos, Peirce, and then Frege attempted to derive a mathematical theory of logic, and account for the laws of thought, by focusing on the central and recalcitrant notion of *symbol*.

The history of this cross-cutting intellectual strand is complex, and ironic. The study of the symbolic split, with one part, referred to as *semiotics*, moving into literature and interpretation; the other leading to the development of formal logic and metamathematics. That part splintered, in turn, into formalism and self-reference, as the 19th-century dream of pure

---

On the meaning side, I mean a lot of stuff, as indicated on the slide: signs, symbols, information, reference, descriptions, codes, blueprints, meaning, truth, objectivity, subjectivity, syntax, semantics, content, data, language, interpretation, memory, prediction, intentionality, and the like. That is: things that are *about* something else. (Information is *about* the rise of interest rates; the mass of Jupiter is not *about* anything at all.)

objective knowledge crashed into logical paradox, quantum uncertainty, and self-referential doubt.

For our purposes, I just want to say two things.

First, the way we have traditionally *understood* computing—e.g., in computer science departments—has largely derived from the upper, symbolic tradition. You can see this by noticing the technical vocabulary in which computing is analysed: programming *languages*, *data* bases, *knowledge representation*, *programming language semantics*, etc. You can see it by listening to Mitch Marcus's talk, which reflected the received theoretical view: that computing involves information. You can see it by noting that the stories from the upper tradition are not causal: few logicians, in my experience, are materialists. This history explains why AI and computer science treat computing so *abstractly*—an astonishingly widely-held view, reflected for example in the idea that the internet is virtual. (It has taken AI, computer science 50 years to come to see computing as a *concrete, materially embodied, physical activity*, *physically coupled to the environment*—something that would have been *obvious,* if it had come from the "mechanism" side.)

In spite of this asymmetry of understanding, however—in spite of this asymmetry in how we *talk* about computing—my conclusion, from studying the stuff, is that, in nature, the phenomenon of computing is much more balanced. It is a dialectical subject matter, involving an inexorable interplay of *meaning* and *mechanism*. Without that interplay, the subject would collapse: it wouldn't have started, would never have captured people's imaginations; wouldn't have come into social power—would in fact reduce to vacuity.

4. Intentionality (~1000)

What can we say about symbols, meaning, intentionality? Out of that enormous story, I want to say just four things.

4.a. The non-causality of reference (~ 400)

First, symbols, language, images, descriptions, and the like, paradigmatically *refer* to something else. And the "arrow of reference"—the relation tying a symbol to what it is about—is stunningly non-causal. Imagine having a thought about X, and think about the relation between your mind and X—the relation in virtue of which the former is about the latter. This

---

3 The latter grew out of a romantic 19th-century dream: that physics, queen of the sciences, could provide us with pure, perfect, objective knowledge of the empirical world. There was only one problem: physics required mathematics, and no one was quite sure where mathematics came from. So Gottlob Frege, progenitor of analytic philosophy, and one of the greatest 19th-century philosophical minds, set out on a life-long attempt to derive mathematics from logic alone. It didn't work. Russell showed that Frege had failed; Gödel proved that Frege *had* to fail; and Wittgenstein argued that even if Frege had succeeded, he would have failed. Meanwhile, on the empirical side, the developments of quantum mechanics and relativity had devastated any naive dreams of pure, crystalline, universal certainty. Freud, for his part, had shown that we couldn't blithely trust our own rationality. And so on and so forth. The whole dream of a "perfectly-known, perfect world" came crashing down, in other words. An intellectual panic ensued, unleashing an obsessive concern with formalism, self-reference, and doubt, which quickly spilled over into arts, literature, and much of society.

relation violates all kinds of cannons of physicality. We can refer forwards and backwards in time—backwards to the Pharoahs of Egypt, forward to the first female American President—all without violating physics. We can refer to the sun, without our reference taking 8 minutes for reference to succeed. I can refer to you, without your being able to detect it (not even the NSA could build a meter—something you could wear in your pocket—that would detect "being referred to"). We can refer outside our light cone (necessary in order for the notion "light cone" to make any sense), to designate places to which physics proscribes the propagation of any signal. Moreover, reference possesses astonishing acuity: I can, at this very moment, refer to the 127th tallest currently asleep person in Bombay—and my reference sails 9000 miles southeast, and picks among the multitudes of people to select just exactly the person I have in mind—with no fuss, confusion, or heat.

In sum: being able to *refer to* the world is bloody mysterious. (So mysterious, in fact, that it sometimes gets you in trouble. I was pointing out the above facts to an undergraduate class, recently—trying to get them to realize how stunningly impressive this ordinary human achievement was. They were convinced it was amazing—so convinced, in fact, that they decided it didn't exist. For the rest of the semester, I couldn't convince them it was possible to refer to anything at all!) And yet, of course—in spite of the mystery—reference is not only possible, but can be *done by physical stuff*. The non-causal *relationality* of reference is compatible with physicalism, that is—which is *why* you can't build a reference-detecting meter (if you could, that truly would violate physics). Crucially, however—and this is the point— *reference is no less real, for not being physically effective*. So it is something a theory of meaning or intentionality must explain.

4.b. Normativity (~200)

Second, language, symbols, description, meaning, and the like, are *normative* phenomena: they involve *values*, *importance*, *worth*. Unlike physical phenomena, such as the rings of Saturn, which just *are what they are*, intentional phenomena—including thought, language, and computing—are subject to *evaluation*. Machines can *work* or be *broken*, descriptions be *true* or *false*; pictures be *beautiful or ugly*, agents be *good* or *bad*—all in ways that wouldn't apply to "purely physical" events.

Normative considerations permeate computing, as *governing constraints*. As in the case of reference, that does't mean that you can open up the hood of a computer and find the normativity "in there," as if it were a causal ingredient. But norms nevertheless govern the phenomenon—shaping what architectures we build, dictating which of all possible mechanisms *count as computational*, informing the notion of what it is to compute.

Many different kinds of norm are employed in current computing: (i) having a *function* (from biology); (ii) meeting a specification; (iii) being true (from logic); (iv) *maintaining* truth (as in an inference system); (v) making a rational choice (as in decision theory); (vi) maximizing some utility (as in economics); and—by far the most important, these days—*sustaining evolutionary survival* (as in biology). What should be evident, however—especially

given this morning's discussion—is that ultimately, inexorably, norms on dynamic activity will expand to include *ethics*: how to live!

4.c. Directedness (~100)

The third consideration is among the most interesting. Perhaps the most penetrating analysis of what is involved in the core notion underlying the whole realm of meaning and intentionality is that of Brentano (Husserl's teacher), who emphasized the fact that intentional phenemona—processes and entities—are *directed* or *oriented* towards something else. To think about something, to refer to something, to remember or predict or consider something, to plan or speculate or model, involves being *oriented towards* that object or state of affairs.

4.d. Ontology (~300)

Fourth and finally: intentional phenomena—symbols, language, information, etc.—are *ontological:* they deal with the world at a certain level of abstraction and description. If I describe something, I describe it *as a brain*, or *as a mind*, or *as a 7 lb paperweight*. It follows that a theory of meaning or intentionality must provide—or at least supply the wherewithal for its users to provide—a theory of ontology: of what an object is; what makes one thing one, and two things two; what the difference is between abstract and concrete, particular and universal, property and type.

This is very different from other parts of science. Consider an ordinary chair—something we can reach out and touch. *No current science can—or needs to—explain what it is to be a chair*. But consider the dictum "Let there be one, and only one, chair!" To give a theory of that sentence would require saying what would be required for that command to be *satisfied*, which would in turn have to deal with what it was to be a chair—and how one chair differed from two.

In passing, note that people who rely on modern science to sustain the *familiar world of material objects* are in for a rude shock. Physics (at best) provides a picture of world-extensive fields: of mass and density, if you are classical, or of quantum probabilities, if you are not. That's all! *The ordinary world of macroscopic objects can't be explained on the physical, causal, "material" side of the house; they are intrinsically semantical or intentional phenomena*. Many are even normative: chairs, for example, can *break*.

5. The Age of Significance (~1500)

OK, well that's a bit of a picture of the realm of the symbolic. The realm of the material or mechanistic I assume is already familiar. What about computing?

Well, I have spent these 30 years looking for a theory of computing. The net result? I *failed.* Not only that, I *had to fail.* We have not had an adequate theory of computing in the past. We do not have an adequate theory of computing now—I have nothing better to offer. And we will also never have an adequate theory of computing, in the future. We won't ever have a proper theory of computing because *there is nothing there to have a theory of.* There is

no distinct, intellectually autonomous or delineable subject matter, worthy of theoretical study, to have a satisfying theory of. Computers, it turns out, are rather like cars: objects of inestimable social, political, and economic importance, but not entities that are destined, *per se*, *qua computers*, to be the subject of substantive intellectual theory.

The reason is simple. Beyond exemplifying the interplay of meaning and mechanism, computers are not sufficiently *special*. In spite of the press, I believe, computers turn out not to be digital, or abstract, or formal—or to exemplify any other property making them a *distinct subspecies* of the genus "meaningful material system." Computers are meaningful physical artifacts, that is—the best we know how to build. Period. The general case. There is nothing more to say.

This might seem a dismal conclusion. But in fact the opposite is true:

> *That there is no theory of computing* per se *is the most optimistic conclusion that anyone involved in the development of computation could possibly have hoped for.*

For instead of being a *subject matter*, warranting its own private theory, computing turns out to be a *site*: an historical occasion on which to see general (unrestricted) issues of meaning and mechanism play out.

What is going on is depicted in the . Go back to the emergence of science. I said at the outset that science as we know it—300 years of "natural science"—is a study of matter, mechanism, materials. What computing, cognitive science, information science, etc., represent, I claim, is a transformation to this conception—a transformation of unparalleled importance, whose character is implicit in everything we have said so far. What computing represents is the end of 300 years of the pure study of matter, materials, and mechanism, serving in its place as a midwife to the emergence of an era of intellectual history of equal importance—something I call an "age of significance," in which normative, intentional issues —meaning and mattering—take their rightful place alongside traditionally scientific notions as matter and mechanism. (Return the "mattering" to matter.)

I pointed out earlier, as indicated in the slide, that the emergence of natural science was preceded by a hundred years of alchemy: a disheveled hand of practically-minded souls trying to turn base metal into gold. As you can see, the new era has the same structure. 20th-century hackers are *semiotic alchemists*, working to turn bits (or web pages) into gold. Once we understand what is going on—once we have our 21st century Newton, and develop a theory of all this stuff—they, too, will probably be shunned, at least temporariliy. But like the alchemists, they are playing an essential precursor role: supplying a wealth of rich, pragmatic expertise about the nature of material symbolism (what Mitch was talking about as "practical kabala"?).

Clearly, we are just entering this new era. We are just inching our way in—making inchoate progress that history will laugh at. But in order to give you a sense of what I think likes in store for us, in the long term—and thereby, to convey some sense of why I think the transformation is so profound—I want to conclude by pointing at just half a dozen properties of this new era.

First, with respect to the matter side of things, there is a fast-growing sense of the importance of seeing intentional systems as *engaged*, *participatory*, *located*, and "*situated*" (context-dependent) agencies. All sorts of issues are relevant here: complexity, emergence, self-organisation, implementation, levels of abstraction, and the like. This is the theory of architecture, of materiality, of body. It is also, ironically, the province of the mathematical theory that goes by the name "the theory of computation," which, on the reconstruction I am recommending, emerges as neither more nor less than a *mathematical theory of causality*.

Second, meaning, language, representation, information, and the like—i.e., the paradigmatic phenomena of the upper, "intentional" strand—are *also* being recognized as engaged, participatory, located, "situated," context-dependent, etc. That is: it will increasingly be recognized that, in order to achieve the kind of integration I am recommending, they, too, can no longer be treated as pure and abstract. This stance towards the symbolic is radically distinct from the reigning conception in the prior era. This is no "view from nowhere," no disconnected, unlocated, a-perspectival stance. Meaningful systems, *qua material systems*, start, first and foremost as embodied, located, and perspectival.

Third, as suggested, it is only when the realms of matter and the realms of meaning are conjoined that we will finally have workable theories of both ontology and epistemology. Weinberg is wrong: it is not physicists who will provide "theories of everything," but the new intentional theorists of the integrated age. Similarly for epistemology: by integrating theories of matter and meaning, we will finally have the resouces to provide what has eluded us for thousands of years: satisfying theories of representation, description, reference, semantics, and the like.

Fourth, as for ontology and epistemology, so for values. This is a fact of unutterable importance. An enormous number of familiar notions that, in the past, have been "extrinsic" or "external" or "extra-curricular"—properties of theorists or observers, only appropriate in "meta-scientific discourse," at philosophy or religion conferences, in the bar—will, in the coming age, be *science-internal:* first-class, fully legitimate notions, part and parcel of what is studied, part of the central subject matter of intellectual inquiry. This list includes: (i) semantic or epistemological notions, such as truth, experience, realism, objectivity, subjectivity, symbol, language, meaning, and consciousness; (ii) ontological notions, such as object, type, property, number, law, regularity, set, etc.; and (iii) normative or evaluative notions, such as purpose, deference, right, worth, value, meaning, concern, responsibility, conscience, and significance. What is being proposed, that is—what is *already happening, even if it is not fully recognized*, as betrayed by computational practice—is that the net of science is being expanded, to include this vast array of substantive notions within its grasp.

Fifth, it follows, from all this, that whereas science has for 300 years been viewed as a detached, 3rd-person, "observing" practice, it is *intrinsic to the nature of the forthcoming age* that its practices will become (in part) inexorably first-person. The foregoing list of semantic and normative notions can't be included in the widened scope of intellectual inquiry, while leaving us outside. On the contrary, the conduct of science—the age-old epistemological

process of attempting to discern and express the nature of the world—fall squarely within the perview of the new inquiry. It is not an option, in other words: the new era will be intrinsically self-referential, with the conduct of science becoming a first-class example of its own subject matter.

This list could be extended in many ways—to talk, for example, about an increased emphasis on the particular and specific, as well as the universal and general. I could talk, too, about my own work in this area: which is essentially to develop a theory of ontology, epistemology, and even consciousness, on top of the concrete substrate of field theory. But instead I want to conclude with a sixth characteristic that, it seems to me, will not only play a profound role in the coming world-view, but is directly relevant to the concerns of this conference.

I said above that one of the most penetrating analyses of intentionality focuses on the *directedness* or *orientation* underlying semantic (intentional) behavior. Strikingly, considerations of *directedness* and *orientation* seem, on reflection, to underwrite a far larger part of the picture than even Brentano imagined. What would it be, in particular, if, as part of a foundational account of this new era, we could develop a theory of orientedness or directedness that could underlie:

a. The *semantic* directedness that underlies *reference*—directedness to *what is the case*;
b. The *psychological* directedness that underlies *attention* and *commitment*;
c. The *epistemological* directedness that underlies *curiosity* and the search for knowledge;
d The *emotional* and *ethical* directedness that underlies love, care, and concernfulness;
e. The intentional directedness that underlies purpose or *telos*—directedness in time; and
e. The *religious* orientation that underlies *awe*, *dedication*, and *reverance*?

It is the prospect that that kind of intellectual unification that keeps me awake at night.

6. Conclusion (~340)

What does this all have to do with the topic of the conference: the classical dichotomy between religion and science, and its implicate distinction between God and the world?[4]

---

[4] One way to say it is by noting an etymological irony. The word "matter" (including in my own writing) has for several hundred years been associated with the physical, mechanical side of our dialectical story. But of course the term 'matter' has another meaning, as well: to be important, to be serious, to be significant. So too for 'material': as in a material object, which is physical, and a material argument, which is an argument that matters. One way we can understand the integrated age that we stand on the brink of, therefore—an age into which computing, ironically, is ushering us—is to understand it as reintegrating the two senses of matter: of bringing significance back into the material world.

Personally, I am not in any explicit sense religious, nor do I find 'God' an especially meaningful term. I do believe, however, that *the world in opposition to which people believe in God*, and *the conception of science in opposition to which people believe in religion*, are woefully inadequate—and undergoing radical transformation. At a minimum, therefore, as we enter the third millenium, any substantive questions about religion and science should be asked, not with respect to the purely physical or mechanical world of causes we are leaving behind, but with respect to the new age of significance we are entering into. It is to this new world that any putative God must be contrasted; with this new transformed understanding that religion must grapple.

For me, I guess, questions of ultimate significance—about what matters, and what gives meaning in life; about the nature of reality, and our place in it, and the like—are too important to leave hostage to the pro- and anti-religion debate. And so my recommendation is not for science and religion to be made compatible, in a quietist way, nor for either to be given prominence over the other, but for the very distinction between them to be transcended. It is towards such a transcendent fusion that it seems to me we are inexorably groping—towards a world of which we are a part, a world that so spectacularly defies description that the very notion of "description" is defined over and against it as a way of watering it down, a world of matter and a world of mattering, a world in whose significance our own significance rests, a world unpredictable and risky and hard to master, a world to fight for and preserve, a world to change, a world to play in, a world to defer to.

It may not be God. But it might be enough.

——— end of file ————

