 Last edited Tuesday, January 16, 2001

Meaning and Mechanism
(The "Mind/Body" Problem for Machines)

I. Preliminaries

A. Review
1. Last time, I wanted to do two things:
   a. Show how distinct—conceptually (intensionally) & extensionally—the six construals are.
   b. Give an inkling of the conceptual complexity of the underlying intellectual landscape.
2. With respect to the second point (about conceptual richness), it is hard to know what to make of it, on first encounter. With luck you will find it intriguing—but you may also find it daunting, even incoherent or confusing.
3. Don't worry about that now. We will be putting lots of structure onto it, in coming weeks. The intellectual landscape has a more coherent topography than first appearances suggest.

B. Problem set
1. The first problem set is due today.
2. I'll try to have comments and grades posted on problem set #1 by the end of this week.
3. The second problem set has already been posted; you can download it from the web site.

C. Plan
1. This week, I want to start clarifying the structure of the conceptual space, by talking about (what I will call) the **primary dialectic** of computing: the tension, synergy, and interplay between notions of **meaning** and notions of **mechanism**.
2. The very subject of computing, I will argue, involve, at the most fundamental level, an essential interplay between these two notions.
3. For reasons we will see, this can be called the **mind/body problem for machines**.
4. Understanding it—having the right theoretical tools to do comprehend it, and doing justice to the interplay it involves in concrete, real-world systems—will prove to be the most important theoretical mandate that an candidate theory of computation must meet.
5. I might also add this foretaste of where we will end up: in spite of our spectacular success in constructing systems, and in spite of incredibly impressive accomplishments in developing mathematical theories, the conceptual structure of this fundamental interplay has not yet been resolved, to anything like intellectual satisfaction.

II. Introduction

A. History
1. To get at the interplay, let's go back and approach it historically.
2. Consider the synopsis of intellectual history, over the 14th–21st centuries, given in figure 1.

3. In the late Middle Ages, our understanding of material phenomena was permeated with notions of spirit. There were still strong connections to "animistic" theories of nature. Both philosophy and what we now think of as the purview of science were together part of "natural theology." In general, the "mental" and the "physical" were not clearly separated.

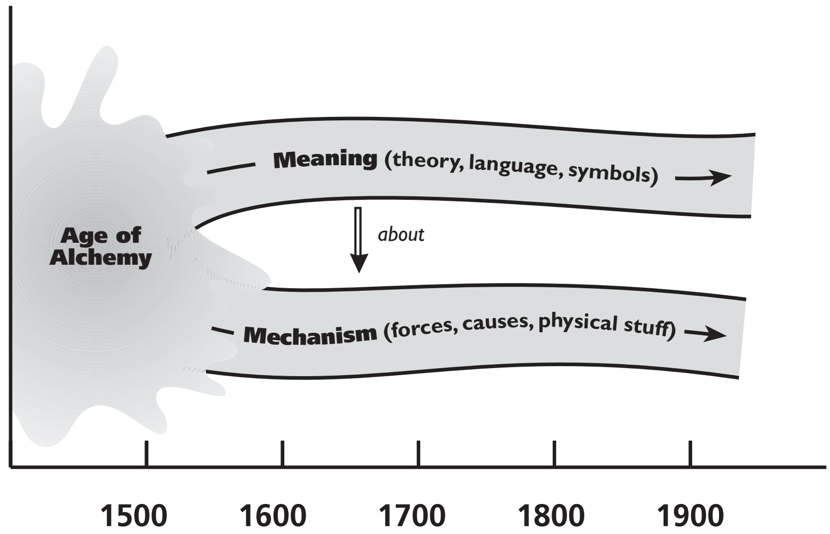

Figure 1 — The Emergence of Science

4. It was the age of **alchemy**: a rag-tag bunch of untheorised practices (most famously: practices that aimed to turn base metals, such as lead, into gold—but in fact a much deeper and larger, even vast, philosophy of nature).

B. Science
1. At the beginning of the Scientific Revolution, with Galileo, Descartes, and Newton, the mental and the physical were pulled apart.
   a. Descartes is most widely studied, these days, in cognitive science, philosophy of mind, and philosophy more generally—for his views about the *mind*.
   b. But of course Descartes was also the person responsible for formulating geometry (and measurement) in arithmetic or algebraic form (cf.. "Cartesian" coordinates, etc.).
   c. More seriously: his work on *both* sides of the divide, by helping to sever the connection between the mental and the physical, unleashed the possibility of studying the physical, unhindered by theoretically recalcitrant issues of spirit, animus, etc.
   d. In other words, even though we broadly indict Cartesian theories of mind, these days, known under the rubric of "dualism," the divide he created may have been an historical necessity, in order to allow the emergence of natural science.[1]
2. Out of the Scientific Revolution, two threads stabilised:
   a. Theory, knowledge, insight, etc.—a realm of **meaning**
      i. I.e., the realm in which scientific knowledge is expressed: the papers, diagrams, symbolic expressions, etc., that are *about* the natural world;
   b. The subject matter of science: matter, materials, physical stuff—a realm of **mechanism**
      i. I.e., the realm of mass and force, atoms and electrons, fermions and bosons: the physical world that scientific theories explain.
3. Note that (in current imagination) the two strands are largely distinct
   a. No one expects *physics* to explain *physicists*

[1] One of my own goals, in my metaphysical work, is to re-unite the two sides of this divide, without compromising the depth and power of science, and without doing an injustice to the richness and subjectivity of mind.

b. The upper thread is *outside the realm of (scientific study)*
4. The two realms are connected by: ⇒ **semantics**
   a. We will have a lot more to say about this in a moment
5. Studying the relation between our knowledge (upper, meaning strand) and what is known (lower, physical or mechanism strand) has been a huge industry ever since, and continues to be so to this day.

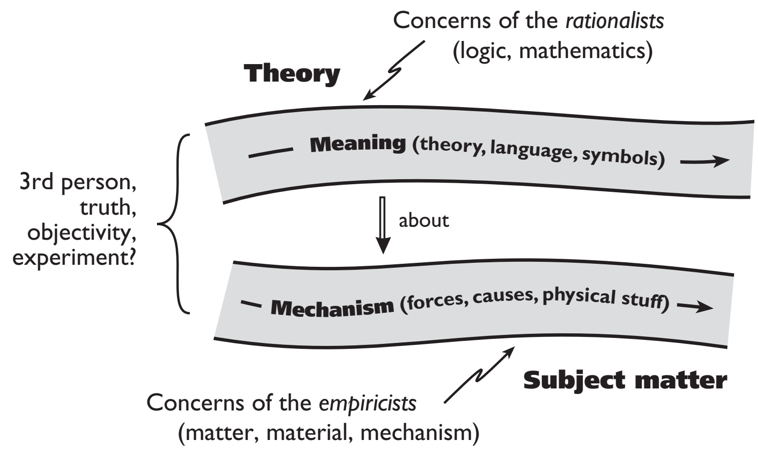

Figure 2 — Relations between Realms

   a. There is lots to be said about the theory–subject matter relation
      i. Is truth the only norm? (science studies)
      ii. Is it entirely referential? Does engineering or construction matter?
      iii. What is the role of experiments?
      iv. What about mathematics? (see below)
      v. What about models?
6. For more on all of this, cf. modern science studies, as well as philosophy of science.

C. Eighteenth and nineteenth centuries
1. Progress on understanding the lower strand (the world of matter, materials, and mechanism) was stupendous; that is the heritage of the sciences.
   a. The study of the lower strand is typically by *empirical methods*; and people who study the mind, or other such general topics, in such terms, are often called **empiricists.**
2. However people also studied the upper strand: the realm of logic, theory, reasoning, etc.
   a. Note that whereas the upper strand represents, in a sense, our understanding of the lower strand, this project—of studying the upper strand itself—is a "meta-level" kind of activity: understanding of understanding.
   b. Cf. Boole, Pierce, etc. (following Plato and Aristotle)
   c. Led to Frege, then Russell and (early) Whitehead, etc.
   d. People who study the mind in such terms are typically called **rationalists**
3. Note that mathematics occupies an uneasy position, between the two threads
   a. Is mathematics *part of nature?*
   b. Or is mathematics a *theoretical tool*—equipment, with which we *understand* nature?
4. (Story of the attempt to derive mathematics from logic alone—Frege).
5. There are also other problematic issues connecting the two realms; cf. figure 2.

D. Twentieth century
1. The two realms remained largely separate, for several hundred years
2. But at the beginning of the 20th century, an amazing intellectual development took place.
3. The physicalist world view was so successful (our understanding of the lower strand was impressive) people began to wonder whether *everything* might be understandable that way.

a. In philosophy, this is called **physicalism**: the (metaphysical) thesis that everything is itself physical, or arises as a product or emerges from physical stuff (including physical arrangements).

4. So the natural question arises: maybe reason, understanding, symbols, etc., are *also part of the (concrete) physical world*.
5. This leads to the insertion (injection) of mind, meaning, etc. into the physicalist world view (cf. figure 3).
6. That intersection, I will argue, is the "site" of computing (figure 4).
7. It is also the site of cognitive science—something we can talk about off-line.

III. Meaning and Mechanism

A. Introduction

1. On the face of it, the properties of things that are physical or mechanical (concrete objects) and things that are meaningful (thoughts, ideas, equations) are wildly different.
2. Today, because of the rise of computing, their compatibility may seem obvious. But it will help us to understand the notion of computing to spend some time highlighting the differences (apparent and/or real).

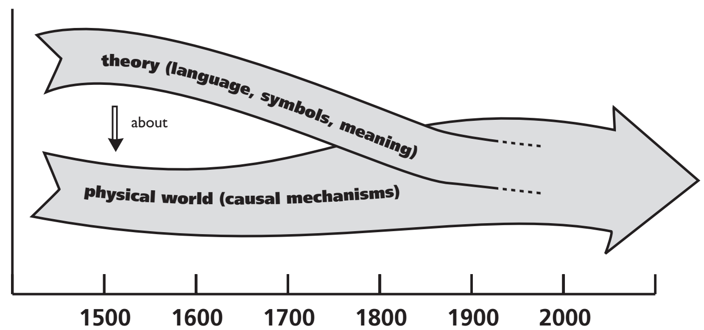

Figure 3 — Inserting Meaning into Mechanism

3. We need to see how genuinely different they are *in order to appreciate how the ideas on which computing is based has attempted to unify them*.
4. Ultimately (as mentioned in the introduction), I will argue that:
   a. The interplay of **meaning and mechanism** is the primary dialectic of computing
      i. Fundamentally, computing involves an intricate interplay of these two notions.
      ii. Dialectical relations between them—the parallels and differences, the ways they push and pull and play off against each other, the requirements, in a single embodied system, to do justice to each, without violating the other—lie at the very heart of computation.
      iii. No other dialectic (though there are several competitors) cuts as deep.
   b. The entire historical intellectual edifice of logic and metamathematics, on which our understandings and theories of computing are based, can only be understood in terms of this primary dialectic.

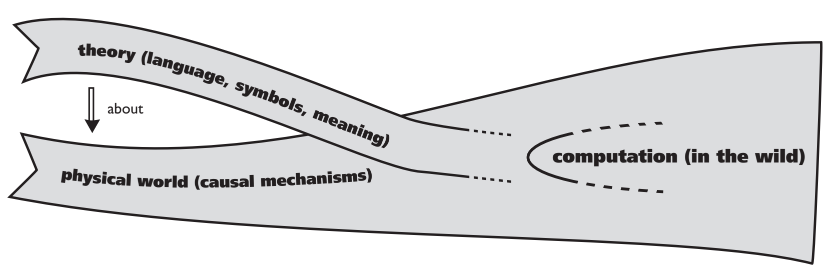

Figure 4 — Computing at the Intersection of Meaning and Mechanism

5. As a result, any candidate comprehensive theory must do justice to it.
6. In fact we can identify this as a fourth (substantive) criterion fourth *substantive* criterion (as opposed to the 3 formal criteria embraced last week—empirical, conceptual, and cognitive):

> **4. Mind/body:** A comprehensive theory of computing must explain the nature of, and dialectical interplay between, the *meaning* and the *mechanism* of computation in the wild.

B. Plan: in order to see what this all comes to, we will adopt the following strategy (this week):
1. First, say something about the nature of the 2 realms (prior to any issue of merging them):
   a. The nature of meaning
   b. The nature of mechanism
2. Second, (very quickly) summarize the development, and conceptual structure, of formal logic—in order to see how it is a solution to this dialectic.
   a. Part of the goal, here, will be to see how we might *generalize* logic's solution, in a way that might be applicable to a wider class of systems; and
3. Third, frame the question of computing against this conceptual and historical background.

C. An (historical) terminological note will help clarify the ensuing discussion
1. As I've already intimated, the *theoretical terminology* of (especially formal) computer science, in terms of which we understand computing, largely derives from the *upper* strand: of language, *meaning*, reference, truth, etc.
2. However: the *content* of present-day understandings of computing has primarily to do with (a somewhat abstract) version of the lower, *mechanism* side.
3. Over the course of the semester, we'll see many ways in which this concrete focus is so, and how it came about. That is:

> ◆ Current theoretical computer science (theories of computability & complexity, mathematical (operational & denotational) semantics of programming languages, etc.
> i. Uses logic-based (*upper-strand*) vocabulary and ideas …
> ii. To talk about mechanical (*lower-strand*) phenomena.

4. As it happens, this is not so innocent.
5. Two primary conclusions we will reach, in this course, are that:
   a. Vocabulary from the upper strand *isn't entirely adequate* to express that lower-strand (mechanistic) aspects of CITW; and
   b. Content (substantive issues) from the upper strand *has not yet been adequately dealt with.*
6. The latter conclusion is a fact to which we are a bit blinded, exactly because we tend to use upper-strand vocabulary—a practice that seduces people into thinking that intrinsically upper-strand (meaning) issues are being adequately dealt with.
7. In the meantime—especially if you are a computer scientist—you should realise that
   a. You will find the vocabulary we use (semantics, denotation, reference, etc.) familiar.
   b. But what I will *say*, using this vocabulary, about the fundamental nature of semantics, denotation, reference, and so on, will be *strange*.

8. In due course, however, especially once we have wrestled with the third (EC) construal, the structure of the territory will start to be more familiar again.

IV. Intentionality (the "Mind" Problem) • Introduction

A. OK, let's take on the first of our three tasks, and look at the top strand: the realm of theories, symbols, language, representation, meaning, etc.

B. Terminology
   1. In the name of the dialectic, I am using the term 'meaning' to indicate this whole realm.
   2. In philosophical discourse, however, it is more traditional to refer to it (following Brentano) as the *realm of the* **intentional** (or sometimes: the realm of the **semantic**).
      a. Note: this is "intentional-with-a-t"
      b. Not: "intensional-with-an-s" (i.e., as opposed to extensional)
      c. It also isn't what I sometimes call "intentional-with-a-d"—i.e., the English noun from the ordinary word "intend" (i.e., in the sense in which, if someone steps on your toe, you might wonder whether they had done that *intentionally*."
   3. For purposes of this course, you can simply take 'intentional' and 'intentionality' to be technical terms, to name the whole suite of phenomena that have to do with meaning, semantics, interpretation (in the logicians' sense![2]), etc.

C. Phenomenon
   1. A list of intentional terms is given in figure 5 (on the next page).
   2. Informally, you can say that intentional phenomena are phenomena characterised by something like "about-ness"—in the sense that, for an intentional entity (such as a symbol, model, representation, simulation, etc.) you can always ask *what they are about*.
   3. So what we need to do is to lay out some identifying characteristics of intentional phenomena (i.e, phenomena that are about other things).

D. Orientation
   1. Perhaps the most basic characteristic of intentional phenomena is that they **orient** towards the world (or allow their users to).
      a. Cf. Brentano
   2. Think of this in term of *reference*:: a way of *pointing* towards something.
   3. Or of a sign or symbol, that *indicates* something (typically something *else*)
   4. In a way, the characteristics of the intentional, which we are going to lay out, are characteristics that are somehow constitutive of what it is to orient in this way.

E. Two notes
   1. Computing
      a. Before we start looking at characteristics, note that there are a spate of computational terms in figure 5 (*program*, *data structure*, *symbol*, *knowledge representation*, etc.).
      b. This fits, because we are going to take computation to be an intentional phenomenon.
      c. Million dollar question: is computation a *unique* kind or type of intentional phenomenon?
      d. Or does the word 'computation' or 'computing' refer (intensionally and/or extensionally) to the whole lot?

---

[2]Cf. Problem Set #3

2. Information
   a. Note that information (as I mentioned last week) passed the "about-ness" test: information (at least in the lay sense) is information *about* something (else).
   b. So for purposes of this course, we will take information to be an intrinsically intentional notion.

V. Fundamental Characteristic of Intentional (Semantic) Phenomena

A. Directedness (we've already talked about this)

1. Most important property: aboutness or directedness.

- *General semantic (intentional) phenomena*
  - Signs, symbols, names, pictures, icons, images, schemata, scores, codes, blueprints, …
  - Words, phraes, sentences, expressions, utterances, statements, assertions, questions…
  - Representations, information, descriptions, depictions, specifications, illustrations, …
  - Analyses, simulations, accounts, theories, stories, explanations, models, …
  - Pointing, showing, indicating, directing, interpreting, …
  - Languages, communication, discourse, conversation, interaction, …
  - Proofs, arguments, demonstrations, rebuttals, …
- *Cognitive correlates*
  - Thought, knowledge, recognition, realisation, deliberation, …
  - Belief, desire, concepts, ideas, decisions, reasons, …
  - Action, behavior, perception, (sensation), …
- *Computational versions*
  - Programming languages, specification languages, …
  - Symbol manipulation, information processing, data analysis, …
  - Data bases, knowledge representation systems, meta-level architectures, …
  - Programs, instructions, codes, memory, data structures, …
  - Pointers, identifiers, names, variables, …
  - Call by value, call by name, referential transparency,
  - Syntax, semantics, interpretation, models,
  - Correctness, …
- *Theoretical notions used to analyze semantic phenomena*
  - Syntax, semantics, meaning, content, significance, …
  - Reference, interpretation, designation, denotation, indication, …
  - Truth, falsity, certainty, likelihood, intension, extension, …
  - Inference, consequence, entailment, …
  - Induction, deduction, abduction, …
  - Logic, consistency, soundness, completeness, validity, …

Figure 5 — The Realm of the Intentional

2. Intentional things are those things that **orient to the world** (outside themselves)
3. For any phenomenon α, that points to or signifies β, there is a relation between α and β
4. For purposes of this class, I will in general indicate this directedness or aboutness relation with a **double-tailed arrow**: "$\alpha \Rightarrow \beta$".
5. This directedness, and the relation that something bears to what it is about, is (in general) studied under the banner of **semantics**.

B. Precision

1. Cf. "The left top corner of the ignition switch on the car currently registered in Morocco that has the greatest number of miles on its odometer."
2. This kind of reference is astoundingly precise—and unerring!
3. This acuity (and all these properties of language, symbols, etc.) may seem obvious. But they are stupefyingly impressive achievements, if you think about it (i.e., if you try to make such capacities out of mere lumps of clay—or silicon).

C. Reach (non-locality)

1. Semantic orientation has unlimited span
2. It can reach across time, space—even possibility
   a. *Long ago* (when Nova Scotia went to Europe to pick up Cape Breton)
   b. *Far away* (star nurseries; to where Pluto crosses Neptune's orbit)
   c. *Future* (the day when the first female US President is elected)
   d. *Don't exist* (Bill Clinton's long-hoped-for son)
   e. *Can't exist* (Santa Claus, a world without violence)
3. Referential connection—**intentional reach**—amazing glue, intentional act to distal situation

D. Non-causality

1. Reference or "aboutness" are reminiscent of AT&T's "reach out and touch someone".
2. Except that intentional reach differs from (literal or telephonic) "reaching out and touching someone" in (at least) one crucial respect: *the touching is not, in general, causal*.[3]
3. *Nothing happens, chez you*, when someone else refers to you, or mentions you, or is thinking about you (or even points at you)
4. I.e., nothing *happens*, in any physically detectable sense, at target (β) of the directedness relation, when an intentional event (α) denotes or refers to or otherwise points toward it.
5. Not detectable
   a. Not even the NSA could build a meter to detect object as target of intentional act
   b. Is it true that *nothing happens*, chez the referent?
      i. No! Something *does* happen! The object *is referred to*!
      ii. This is entirely real. Cf. someone's asking "Was I referred to the corporate meeting about down-sizing?" This is a valid question, with a perfectly true or false answer.
      iii. However, nothing *detectable* (*causal* or *potent*) happens, at the referent.

> ◆ The referent or fragment of the world, in virtue of being that towards which the intentional act is directed, is *not thereby bathed in a flux of discriminable energy*.

[3] *Sometimes* it is (or at least seems to be) causal: situations of perception and action. We'll talk about this later.

6. Thus this is the first mystery about semantic reach: it outstrips causal bounds.
7. Doesn't take *time*
   a. Refer to the sun: doesn't take 8 minutes to get there
   b. Can refer outside the light-cone!
8. Note
   a. This non-causal nature of intentional directedness is necessary to its ability to span time, space, possibility in ways that outstrip the (immediate causal) strictures of physical law.
   b. Otherwise it *would* be in violation of physics!

E. Non-locality

1. It follows, from all these things we are saying, that semantic phenomena (such as language or mind) are intrinsically (at least partially) **non-local**.
   a. Pirate's map: determine truth; don't use a microscope
   b. *Truth:* not local or intrinsic property

> ◆ Qua intentional systems, languages, minds—and perhaps computers?—are not entirely *local* phenomena.

2. What something refers to, whether it is true, etc., that is, are not, in general, properties of an intentional event that take place, or are "settled by," proximal facts. They depend on facts in the distal state of affairs towards which the original phenomenon points.
3. The study of semantical phenomena, therefore—such as cognitive science, and the study of the mind, at least, and perhaps (that's the implication we are pursuing) the study of computers, if indeed computers are intentional objects—*cannot simply be a study of the local facts.*
4. Rather
   a. If computers are intentional, and involves symbols, models, information data structures, etc., that refer to or model or denote or carry information about the world, …
   b. It follows that a comprehensive theory of computing (meeting our originally three, but now four, criteria) will have to study computers in relation to the outside world—machines are they are *located in, and as they relate to, their environments.*

F. Pause

1. I want to stop, for a moment, in giving this list of properties, to make a general observation.
2. Reference, denotation—intentional directedness—is an *incredibly impressive achievement.*
3. Indeed, it is *so* mysterious that my undergraduates, after I present them with this list of properties, and explain what it actually involves, are rather in the habit of denying that it happens. I.e., they doubt the reality of semantics, reference, reach, etc..
4. But on reflection, these facts are hard to deny
5. Unless semantic reach did work, the ahistoricity of physical law would prevent us from referring to the past, to the future, to things far away, to non-actual or non-possible events—indeed, to much of anything at all.!
6. In fact without genuinely non-causal semantic reach, there would be no languages, no symbols, no thinking, no information, no minds, and no computers.
7. It follows, therefore, that:

> ◆ Among other things, a theory of semantics or intentionality—and, therefore, a comprehensive theory of computing, if computing is a genuinely intentional subject matter—must explain how the extraordinary time-and-space-and-possibility-transcending semantic reach can possibly work.

G. Separateness

1. In spite of all these relational facts, there is an equally important *local* and *causal* structure to intentional phenomena, including computing.
2. This stands as something of a counterpoint to the long-distance theme (or it to this).
3. In the small, intentional acts—symbol manipulation, information processing, thinking, etc.—are remarkably private and confined, arising out of the local arrangements of very small portions of reality.
   a. Complex brain states, subtle machine configurations, patterns of marks on a paper—in particular situations, each are sufficient to establish a determinate content, seemingly without the immediate participation of anything larger or external.
   b. Thus we can think about things in ways that no one else knows about, without (again, at least in the small) those thoughts affecting or depending on anyone or anything else.
   c. Similarly, we can imagine and worry and refer and think, hypothesize and wrangle and dream, about situations in ways that are unconstrained by any actual (causal) facts about the subject matter—even whether that subject matter exists.
   d. This is (sometimes) a problem—but it is an indescribably good thing, too.
   e. For example, the same moral holds of computers: we can construct machines to simulate and model, explore counter-factual (non-actual) hypotheses, simulate disasters and other non-hoped for outcomes., all without worrying about whether their "referents" exist (in fact depending on the fact that their referents *don't* exist).
4. In sum:

> ◆ The *separateness* of the intentional (computational) event from its directed-at target is as fundamental as the long-distance semantics that binds the two together.

H. Disconnection

1. This separateness or partial independence of an intentional event from that towards which it is directed—this **disconnection** of (for example) a symbol from what it symbolizes—is absolutely critical to the very notion of what it is to be intentional, what it is to be a symbol.
2. Among other things, it underwrites the elemental fact that intentional entities can be *wrong*.
3. That is, *disconnection is the metaphysical reason for the possibility of* **error**.
4. But error isn't disconnection's only consequence. There are positive consequences as well.
5. If there were no gap or room between one's thought and the situation one was thinking about, no gap between a computational model and what it is a model of, we would be locked into a kind of "vice-grip of veridicality"—something that might occasionally be wonderful, but would mostly be horrible.
   a. Modeling, simulating, hypothesizing, planning, and all other activities that involve non-actual representational content require this disconnection

   b. Modeling tectonic plate movement would require dragging the plates along with you.
6. You can appreciate these morals in our cognitive experience. Without disconnection:
   a. Fantasy lives would be metaphysically banned.
   b. Thinking would also be dangerous.
   c. As someone (who?) said:

> ◆ We think so that our hypotheses can die in our stead.

7. And these morals are just as urgent in the computational case. Sans disconnection, it would be impossible to use computers to model, simulate, calculate—or do much of anything else.

I. Partiality

1. Similarly, separation or disconnection is responsible for the *partiality* of information.
2. A give and take between representation and represented allows a representation to capture some things about what it represents, and let other things slide away.

VI. Summary

A. Immanence and transcendence
1. Taken together, all these facts about intentionality—its incredible reach, and the separation (disconnection) between the intentional phenomenon itself, and the object towards which it is directed, and so forth—show that all intentional phenomena, including computing, involve an inexorable interplay between:
   a. The local, occurrent, effective (causal), particularities of an intentional action, embodied in a cohesive physical fragment of the world (what I sometimes call the **"immanent"**), &
   b. The long-distance, non-causal, semantic tentacles, on the other, that reach out and gather up an arbitrarily distant or remote target or subject matter (what I sometimes call the **"transcendent"**)
2. This juxtaposition of immanence and transcendence is one of the great mysteries of intentionality and computing—something a comprehensive theory should explain.
3. One reason computers are so successful, I believe, is because, in their actuality, in their existence, they actually manifest a solution to this great as-yet unsolved mystery.
4. It is the fact that, in their actuality, they so far outstrip our understanding, that makes foundational inquiries like this so important.

B. Overall, the picture so far consists of two parts

1. Proximal

   a. A causal nexus: a located, effective structure/process, such as a symbol, sentence, or piece of information; and the (causal) processes defined over it: the manipulation of a formula, the derivation of a (logical) conclusion, the production of a data structure.
   b. This proximal state of affairs, arising out of the activity of some substrate (such as the brain or the circuits of a computer) is—or anyway surely should be—subject to **causal** laws.
   c. In particular, it will be caused by some prior events, and in turn will itself cause some subsequent events—all as subject to the laws of physics.
   d. These causal relations between and among aspects of the proximal mental state will be

indicated (in this class) with *single-tailed arrows* ('→').

2. **Distal**
   a. Some other (typically distal) situation or state of affairs that phenomenon *refers to*, *denotes*, *points towards*, or is in some other ways *about*.
   b. As already mentioned, if α is the proximal event or phenomenon and β the distal state of affairs that α points towards, then the "pointing" relation between α and β will be indicated (in this class[4]) with a double-tailed arrow: α ⇒ β.

Realm of the sign	*Relation*	*Realm of the signified*†
sign	signifies	signified
symbol	denotes	denotation
name	names	named
term	refers to	referent
description	describes	described
vehicle	carries	content‡

Figure 6 — Terminology

†Often, the signified or denoted entity is simply given its usual name. Thus we might say that the name "Jedediah" denotes a *person*, or that the roman numeral XXVII denotes a *number*.

‡This is too simple; we will talk more about content in a future class.

C. Semantics
   1. The study of the relation between α and β is usually called **semantics**.
   2. Depending on the particular type of thing that α is, the relation it bears to β is called by different names (see also figure 6):
      a. If α is a symbol, name, or term, one often says that α **denotes** or **refers to** β.
      b. If α is a description, one says that α **describes** β.
      c. If α is a representation, one says that α **represents** β.
      d. If α is a sign—or sometimes just in general—one says that α **signifies** β.
   3. In general, the relations between the various kinds of sign and various kinds of signified entity are called **semantic relations**.

D. Constraints
   1. So far, in our investigation of intentionality, we have assembled four parts of the puzzle:
      a. Two realms
         i. A realm of *symbols* and other intentional entities that point; and
         ii. A realm of *referents* or other entities that are pointed at
      b. Two relations
         i. Causal relations ('→') between/among things that bump & shove into each other; &
         ii. Semantic (pointing) relations ('⇒') between symbols (and other intentional entities) and the things that they point to.
   2. Next time, we will add the fifth and final ingredient (in our overall picture of intentionality): **constraints** between and among these four elements.

—— *end of file* ——

[4]Note: the use of single-tailed arrows to indicate causal (physical) relations, and double-tailed arrows to represent intentional or semantic directedness, is not a general practice. It's just a convention we are using in this class.