 Last edited Saturday, January 20, 2001

# Norms, Constraints, and Logic

## I. Intro

A. Problem sets
   1. As I mentioned on Monday, Problem Set #2 will now not be due until January 30.
   2. The problems we were having with Annotate should now be (pretty much) fixed:
      a. People who didn't have access to Annotate: now should
      b. The problem with double carriage returns on Unix not producing paragraph boundaries (in Annotate) should now be fixed.
      c. If anyone has had troubled resubmitting, this means that someone has posted (or tried to post) a comment. If this is plaguing you, send a note to Ruth Eberle[1] and ask her to "clear" your entry, so that you can resubmit

B. Plan
   1. Finish up discussion of semantics, by talking about norms and constaints
   2. Then illustrate what we've assembled so far, by going over the conceptual structure of logic

## II. Review

A. So far, we've identified 4 things (in general systems involving intentional behaviour)
   1. Two *realms*:
      a. A (**syntactic**) **realm** of symbols or representation—realm of symbols or vehicles
      b. A (**semantic**) **realm** that the symbols are about—**realm of reference**
   2. Two kinds of *relation*
      a. *Causal* relations ('→')
         i. In and among symbols, between symbols and reference, in and among referents,
      b. *Semantic* relations ('⇒')
         i. Intentional "directedness"
         ii. From realm of symbols to realm of reference.

B. The realm of symbols
   1. This realm is also sometimes called the **syntactic realm**.
   2. As we will see (when we talk about the "mechanism" side of the primary dialectic), it is something of a *causal* nexus—involving a located, effective structure/process, such as a program or data structure, or a word, sentence, or thought.
   3. Location
      a. Paradigmatically (especially when our primary focus is cognition!) the syntactic realm is in the head—i.e.., is *internal*. (I will often draw it that way, in pictures.)

---

[1] Ruth Eberle <reberle@indiana.edu>

   b. That is, it will in general (in discussing these systems) be **proximal**.
   c. In cases of language (and literature), however, it isn't in the head, but is *external*—on paper, on the internet, in "air waves" during conversations, etc.
   d. So "location" isn't what the syntactic realm is defined in terms of.
4. Definition
   a. Rather, what is distinctive about it (what matters most—what determines, of a given entity or state of affairs, whether it belongs in this realm) is that it is at the *tail end of double arrows* ('⇒').
   b. I.e., it is the *source* end of the relation of intentional directedness (which in this part of the course we are taking to be the "mark of the mental").

C. The realm that the symbols are about
1. This realm is sometimes called the **semantic realm**, or the **realm of reference**.
2. Some other (typically distal) situation or state of affairs that phenomenon *refers to*, *denotes*, *points towards*, or is in some other ways *about*
3. As we've said, if $\alpha$ is the proximal event or phenomenon, and $\beta$ the distal state of affairs that $\alpha$ points towards, then the "pointing" relation between $\alpha$ and $\beta$ will (in this class[2]) be indicated with a double-tailed arrow: $\alpha \Rightarrow \beta$.
4. Location
   a. Paradigmatically, the semantic realm is out in the world—i.e., is *external* (a party next week, the chair you see across the room, a situation in Bosnia described in the morning papers).
   b. That is, the semantic realm is paradigmatically **distal**.
   c. But again, the semantic realm is not *necessarily* external (outside the head).
   d. In introspection, for example, you think about your own thoughts. In that case the realm you are referring to involves structures inside your own head—i.e., *internal* structures.[3]
5. The study of the relation between $\alpha$ and $\beta$ is usually called **semantics**.
6. In general, relations (and the study of the relations!) between various kinds of sign and signified are also called **semantic**, though more particular versions has particular names:
   a. If $\alpha$ is a symbol, name, or term, one often says that $\alpha$ **denotes** or **refers to** $\beta$.
   b. If $\alpha$ is a description, one says that $\alpha$ **describes** $\beta$.
   c. If $\alpha$ is a representation, one says that $\alpha$ **represents** $\beta$.
   d. If $\alpha$ is a sign—or sometimes just in general—one says that $\alpha$ **signifies** $\beta$.
7. Definition
   a. What is distinctive about the semantic realm (the realm of reference) is that it is at the *arrow end of double arrows* ('⇒').
   b. I.e., it is the *target* end of the relation of intentional directedness.

---

[2]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.

[3]In this course, of course, as in any cognitive science inquiry, our subject matter—the situations that we stand in a directed intentional relation to (which is to say, the situations that our thoughts and words bear double arrows towards)—includes what is inside heads.

D. Relations
   1. The semantic relation ('⇒') is something we talked about a lot last time. In general, it is *not a causal relation*. Rather, it is something that (for whatever reason) *holds*, or simply *obtains*.
   2. However a signs, symbol, program, data structure, word, thought, etc. is also (as we will see in detail, next time, when we talk more about mechanism) typically a causal nexus.
   3. That is, the proximal state of affairs—arising out of the activity of some substrate (such as the brain or the circuits of a computer)—is subject to **causal** laws.
      a. Will be caused by some prior events
      b. Will in turn itself cause some subsequent events
      c. All subject to the laws of physics.
   4. Causal relations between and among aspects of the proximal mental state will be indicated (in this class) with *single-tailed arrows* ('→').

E. Summary
   1. In sum: **semantic relations** ('⇒') **determine the structure of the territory**
   2. The causal arrows ('→') don't play that kind of structuring role. Rather, they lead wherever they lead: wholly within the syntactic realm, wholly within the semantic realm, or between the two (in either direction).
   3. This asymmetric responsibility for structuring intentional territory reveals a general pattern:

      > ◆ It is the *semantic relations* (double arrows) that matter;
      > It is the *causal relations* (single arrows) that do the work.

   4. You can think of the coordination conditions (between causal and semantic relations) that we will be talking about today as another instance of this general attempt to integrate *what matters* and *how things work*.

F. Plan
   1. These four types of ingredient (two kinds of entity; two kinds of relation) are not enough
   2. Reason
      a. They are too *unconstrained*
      b. Need *constraints*: to tie them all together
      c. Today, talk briefly about them: **constraints** or **norms**
   3. Then, talk about logic
   4. Next time, we will talk about mechanism ("the body problem"), and then finally bring all these morals back to computing.

## III. Constraints and Norms

A. Introduction
   1. This double division of our subject matter into two kinds of things (two kinds of entity, related by two kinds of relationship) generates a whole series of **constraints** that an intentional or cognitive system must honour.
   2. In general, these constraints follow a general form
      a. The *semantic* relations are taken to be established in advance.

   i. How they are established (i.e., how it is that "dog" is our word for dogs, 'Muddy Waters' refers to an old blues player, and 'or' is used to signify disjunction) is very tricky—and not all that well understood.
   ii. By and large, we'll assume that they *are* established (usually it will be clear from context what content or semantic value these symbols or representations have).
  b. The *causal* relations, in contrast, are assumed to be things that we (e.g., as designers of a computer, or as rational beings) can control.
  c. The causal relations are *judged* by the extent that they honour (establish, satisfy) various constraints.
3. In general, we will describe these constraints (which an intentional system must honour) under a general rubric:

> **$C^5$** — Coordination conditions on content and causal connection

4. The basic idea behind the $C^5$ model is simple: your actions (what you do, how you sense the world, how you think—e.g., how you proceed from one stage to another), viewed as *causal connections* (between your mind and the world, and between one state of your mind and the other), must honour certain *conditions* or *constraints* that have to do with the ways they are *coordinated* with the *semantic content* of the thoughts and words and ideas and symbols.
5. That is, intentional systems must honour—through their $C^5$ constraints—something like the following overarching goal:

> **Overarching desideratum** (on an intentional system)**:** to have, or arrange for, the *causal relations and processes honour the semantic relations*

B. Norms
  1. When we talk about the constraints between causal and semantic relations, we are not just talking about conditions that happen to relate them (like the conditions that relate force and mass, which you learn about in physics courses)
  2. Rather, these constraints are of a very special kind, having to do with another essential characteristic of intentional phenomena: its **normative** character.
  3. Motivation
    a. To understand normativity, go back to the *separateness* of signs and symbols that we have talked about, and how this separateness leads to the possibility for error.
    b. When things do part company, it is not just that one is different.
    c. Rather, in the general case, the representation will be **wrong**.
  4. In other words, we not only produce and consume intentional activity—we not only are or are not coordinated, in the ways we are going to talk about
    a. Rather, intentional phenomena are *judged*; some are better than others.
    b. Being *true*, for example, is better than being *wrong*.
  5. Values
    a. This is value talk, which leads into subjects like ethics, which at the outset we identified

as an aspect of the mental that we were not going to consider too much in this course.

b. It is important to realize, however, that truth and falsity—and the *correctness* of representation—all of which are normative notions, are central to how intentionality works.

c. So in this special sense, we will be considering values after all. (More on this later.)

6. Success conditions

a. Given their normative character, we can talk about the coordination between semantics and causation in terms of what are sometimes called **success conditions**: what the conditions are when everything is working correctly—when everything is doing what it *should* be doing.

C. Example

1. To see all this in action, let's turn to logic
2. I assume that all of you are familiar with logic, at least in the sense of having seen the syntax
3. What I want to do here is not to deal with technicalities at all, but instead to lay out the conceptual structure that underlies logic, so as to see:

a. What problem it was designed to solve

b. What general character of solution it illustrates

c. Whether we can bring that framework to bear in our analysis of computation.

## IV. Logic[4]

A. Introduction

1. Setup

a. Predicates: denote ('⇒') properties

b. Relation symbols: denote ('⇒') relations

c. Terms: denote ('⇒') objects

d. Given that, sentences denote ('⇒') truth values

2. Inference

a. So far, just a static systems

b. Define an "effective" or "mechanical" relation of derivability ('|–')

c. But not *just any old effective relation*: that would be *mechanics!*

d. Model *entailment* (not rationality! — big shock, when met logicians)

3. Define *norms* (**success conditions**)

a. Soundness: if provable, then true: *want what you get*

b. Completeness: if true, then provable: *get what you want*

4. Interpretation function, and truth vs. validity (Division of labour)

a. How does the system know what the symbols denote?

b. It doesn't!

c. That is up to the "user"!

---

[4]Note: these notes are very abbreviated. A great deal of discussion came up in class, on a variety of topics. If I have time, I will come back in the future and write details notes incorporating these facts. For the time being, it will be enough to follow the discussion into computing (and also to see the review of logic given at the beginning of the notes next time (3a)); the important points will be brought up again, in the computational context.

   d. In: the *interpretation function*
5. Two facts about interpretation
   a. Not a function in computer scientist's sense: nothing *happens* (◆)
   b. Norms on system (validity, etc.) are true *for any interpretation* (couldn't possibly say whether a given axiom is *true*; somehow "not within the system's ken")

B. Perception and Action
1. If there is time, we will consider the standard analysis of perception and action
   a. **Success conditions**
2. Note: all within the general C5 model

—— *end of file* ——