Towards a Theory of Representation and Modelling

Part I — Categories of Correspondence

Description:	Notes for Representation/Modelling seminar, November 5, 1985
Filed as:	[site]RepMod-Cat.notes
Site:	{phylum}<BrianSmith>papers>|{dsk}<files>papers>|{floppy}
Last edited:	5-Nov-85 09:52:29

Outline

A. Introduction

— Methodological assumptions
  — Start by looking at "categories of correspondence" — relations between structured objects.
  — Things not dealt with here:
    — Metaphysical foundations, *registration* problem, etc.
    — Authenticity, originality
    — Theory of computation based on representation (specifically: *embodied* representational processes, *embedded* in the world)
    — Semantics (direct and indirect), model theory, etc.
    — Consequences of this model of computation for cognitive science (claim that mind is representational). What does it have to do with syntax? with combinatorial structure?
  — We'll come back to all these issues in due course.
— Background
  — Terminology: write '$P \ll Q$' for 'P represents Q'.
  — Standard: strict hierarchy; use/mention distinction; levels of "designation".

    - Cf. quotation and numerals in computer science; then 2-LISP and 3-LISP.
    - Similarly: languages, meta-languages, etc.
  - But also counterexamples:
    - A UTM represents the tape of another Turing machine with numbers on a tape; still say that the UTM computes what the represented Turing machine computes (in fact that is the whole point).
    - Similarly, photograph P of sailboat S, and copy P' of P.
      - Does $P' \ll Q$? or $P' \ll P$? or both?
      - Suppose there is a tear in the original; you suddently "see" P, not just Q. I.e., which of the above you choose depends on what properties you pay attention to. Some go through; some don't. Suggests representation relation is both more complex and more flexible than the standard account would suggest.
    - Another example: If (object) program P_1 is compiled into (target) program P_2, and if P_1 represents anything Q, does $P_2 \ll Q$ or $P_2 \ll P_1$? Can argue for the former. It seems as if one way to characterize what compilers are trying to do is to look *through* P_1 *to* Q (i.e., from $P_2 \ll P_1 \ll Q \Rightarrow P_2 \ll Q$).
    - I.e., we currently take
      - designation or interpretation to be strictly non-transitive, and
      - modelling as transparent and transitive, and therefore ignore it.

      Both assumptions are untenably rigid.
    - So focus on representation for now: will then *define* a model as a particular (and limited) kind of representation.
- Look only at the *relations of structural correspondence*
  - i.e., on the *mode of representation*
  - by no means the only (in fact not even the most) important question wrt representation in general.
  - Assume both structures are already fully registered in advance.
    - Although we are assuming that $P \ll Q$, and that P and Q are fully registered, we don't assume that they need to be fully registered independently. I.e., we allow the possibility that the "integrity of each object as a coherent unity" depends on the other.
    - More on this when we get back to metaphysical foundations.
- Conclusions that will emerge:
    - Since theory of *representation* and theory of *models* are the same subject matter (already noted), it follows that whether one characterizes something *linguistically* or *model-theoretically* is not as different as I used to suppose.
    - cf. Goguen, situtation theory in computer science, etc.

- Continuum of relations of *structured correspondence* between complex objects, from isomorphism (even identity) through ... *who knows*? ... through full scale language (and beyond?).
- Some compose; some don't.
- So use/mention distinction is often a confusion. (Cf. Nunberg, Perry, and others on mention in language; my overly zealous adherence to the use/mention distinction in describing 2-LISP; etc.

B. Paradigmatic Examples

1. Real world (robust physical)

a. balsa models $\preccurlyeq$ airplanes

b. photographs $\preccurlyeq$ physical situations

c. numbers (and units?) $\preccurlyeq$ physical reality

- "F=MA": classical mechanics, etc.

d. Thermostats $\preccurlyeq$ temperature.

e. Musical notation $\preccurlyeq$ music.

f. Blueprints $\preccurlyeq$ buildings

g. Contour maps

2. Computer science

a. bit-map $\preccurlyeq$ Dandelion display screen

b. sets of quadruples $\preccurlyeq$ Turing machines

c. numerals $\preccurlyeq$ numbers

d. Abstract data types

e. notations (lexical & graphical) $\preccurlyeq$ internal structures

f. Term models for Prolog-like languages

- what are they models *of*?

3. Linguistics

a. parse trees $\preccurlyeq$ the (genuine) syntax of language

b. text $\preccurlyeq$ language

4. Logic and Philosophy

a. {0,1} $\preccurlyeq$ polarity in situation theory

b. possible worlds $\preccurlyeq$ modal relations

c. Godel numbers $\preccurlyeq$ expressions in a formal language

d. Functions : possible worlds $\rightarrow$truth-values $\preccurlyeq$ propositions

5. *Mathematics*

   a. sets of ordered pairs $\ll$ functions
   b. trees $\gg$ models of all sorts: indexed pairs, very simple trees with a single sequence, binary trees with integers (so that daughters of node k are 2k and 2k+1; parent is [k/2]; etc.).
   c. Dedekind cuts, etc. $\ll$ real numbers
   d. order pairs of integers $\ll$ rationals
   e. sets of totals functions $\ll$ partial functions

6. *Language and Mind (last but by no means least)*

   a. the general problem of intentionality
   b. how do languages and minds represent their embedding world?

C. Working assumptions

1. Metaphysical assumptions
   — The world is an infinitely rich, variegated, and transcendent place. There is more to any real thing than its exemplification of any number of properties — even if it is (even necessarily) the only thing that exemplifies one or more of those properties.
   — In order to think, talk, or represent part of the world (all one can do), one must, roughly, conceive of it in terms of regularities: roughly, focus on a finite piece of it, carve it into objects, sort it by properties and relations, etc. To conceive of it at all is to allow the possbility that it might have been other than as it is. And to describe anything at all in terms of properties and relations is to ignore (throw away) certain aspects of it. I.e., classification of any sort does a certain amount of violence. Object identity (i.e., where the *boundaries* of an object are) is at least partially relative to the particular properties and relations in terms of which it is classified.
   — Etc. etc.; this is not the place to say more about this. Nevertheless, one's assumptions about such things impinge almost immediately on what one makes of representation. For example, I will maintain distinctions between lots of things I will call isomorphic. Disagreements about such moves may derive from disagreements at this foundational level. If so, will try merely to note that and move on.
2. Objects, properties, and relations:
   — Will assume (standard CSLI wisdom) finite situations consisting of objects exemplifying or instantiating properties and relations.
   — Can therefore use standard CSLI language of objects, properties, and relations.

- Also of types.
- Simple hierarchy of types, properties, and relationships: so the property of being a person is an object that is one level "higher" than the person him/herself.
  - don't really care too much about what this type hierarchy is like: well-founded, etc.
  - Would matter when things get formalized, but we're nowhere near that rigorous yet.

3. Abstraction
   a. the world is regular (has regularities) at all kinds of "levels of abstraction"
      - means roughly the obvious thing: that properties and relations can be approximately grouped into clusters (mass, charge, and spin might cluster; similarly table, chair, person-in-the-sense-of-a-furniture-user; etc.), in terms of which the salient regularities in the world are approximately organized (in fact this is surely what gives the level its coherence or identity).
      - don't care if this isn't an exact partition.
   b. To represent, characterize, or describe something *at a level of abstraction* is to represent, characterize, or describe it in terms of the properties or relations that cluster to form that level of abstraction.
   c. There is more to anything that exists than is captured by characterizing it at any level of abstraction.
      - Also, since the identity of an object depends in part on how it is classified, the identity may depend on the level of abstraction. Objects, in other words, may be (are always?) *functionally defined.*
      - Nonetheless, even if one object P is functionally defined in terms of a finite number of properties and relations $R_1 - R_k$, it doesn't follow that any other object Q that exemplifies exactly the same relations $R_1 - R_k$ in exactly the same way as P can be identified with P.
      - Why? Because there will always be an infinite number of other legitimate properties on which P and Q differ, which are essential to P even though they don't figure in the identification of P as a an object of the given functional type.
   d. Thus, *even if P is functionally defined at a relatively abstract level of abstraction* (pencil, detente, teacher) *it doesn't follow that P is an abstract object.*
   e. Nonetheless, there may be genuinely abstract objects: types, properties, numbers, etc. I don't know what to say about them. I don't know whether anything I have said about real objects holds of them or not. In fact I don't know whether I believe in them or not.

3. Matters of agency
   a. Often want to say that $P \ll Q$ *for an agent A.*

- I.e., to claim that the representation relationship doesn't obtain in virtue of any intrinsic properties of either P or Q.
- But I won't consider the agent A for three reasons:
  - i. Don't believe this is always the case.
    - It had better not be, if we are ever to define a notion of computation on it that will serve as the foundation for cognitive science.
    - Doesn't mean we can't say of some representation relationships that they only hold for an external agent A. Point is only that our theoretic framework won't assume it.
  - ii. Can view what we do today as a study as of how agents take representations to correspond to what they represent
  - iii. Won't assume anything very much about *how* the structure of P is interpreted as corresponding to the structure of Q: only *what* structure of Q the structure of P corresponds to.
    - What matters particularly here is when P is viewed as a *rule* or *procedure* that tells A what the structure of Q is. In such a case the correspondence relationship between P and Q can rely in part on the state of A
    - As I understand it, this is exactly the difference between Press and Interpress representations of marks on a two-dimensional surface: Press files, roughly speaking, are *(effective) descriptions* of the resulting marked page; Interpress files are *(effective) procedures for marking the page* — the major difference being that Interpress files depend on a notion of state in the interpreter (effector), whereas Press files don't (though their ingredient expressions are context-relative to the whole of the embedding Press file itself; that's a different notion).

b. *dynamism*: different question has to do with P's be a process or activity, not a static or passive structure (Q, similarly, can be dynamic, but P is the problem case).
   - This is more important to deal with in getting a theory of embedded computation off the ground; however I will ignore it here mostly because I don't know what to say about it. All I can do here is to try not to bias the account of P against its being ultimately active: i.e., to describe it at a level of abstraction that could be used for active and static structures.

c. *Original (authentic) vs. derivative representations:* The standard distinction obviously holds in this subject matter. I will simply ignore it for the moment: I think it can be dealt with only once we have a better grip on processes.

D. General facts about Representation

1. Representation is not generally ~~symetric~~ reflexive or transitive
   - I.e., $\neg [A \lll B \Rightarrow A = B]$
   - I.e., $\neg [A \lll B \land B \lll C \Rightarrow A \lll C]$
2. A representation is not what is represented
   - as opposed to implementation
3. Representation is relative to a *level of abstraction*
   - or relative to conceptual scheme or "parse", if you prefer more mentalist language; relative to how the world is *registered.*
   - "accurate as far as it goes": this is how that intuition works.
   - i.e., to say that $P \lll Q$ at a given level of abstraction is to identify a cluster of properties, relations, etc., for P (call this $\mathfrak{A}_P$), and a (quite possibly distinct) cluster of properties, relations, etc., for Q ($\mathfrak{A}_Q$).
   - P's possession or exemplification of its properties and relations will in some sense correspond to Q's possession or exemplification of (possibly other) properties and relations.
   - This is meant to get at the idea that representation has to do with *structured correspondence.* There may be other relationships of significance (A's being a symbol of B?) which don't fit this model. Not at this point sure.
   - Terminology: If $P \lll Q$ at a given level of abstraction (of each), say that *P represents Q wrt* $\mathfrak{A}_P$, and that *Q is represented wrt* $\mathfrak{A}_Q$.
   - I.e., there are lots of different relations that we signify with the symbol '$\lll$'.
4. Representations are only *partially significant*
   - It follows (from original claims) that for any objects P and Q such that $P \lll Q$, there are always properties of P that are not important (play no role) in the representation relation holding between P and Q. I.e., *no representation is* (what we will call) *fully representational.*
   - This clearly holds for non-abstract (real) objects, but I am inclined to think it holds for abstract objects as well. If nothing else, P will always exemplify the property of being a representation. But it is at least usually true of other cases: imagine representing a tree with ordered pairs of nodes and appropriate indices; the *length* of the ordered pair corresponds to no property of the tree represented.
5. Representation is with *respect to type*
   - First, there obviously isn't just one representation relation $\lll$, but rather a whole family of them, with different properties.
   - They (at least potentially) relate whole domains, not just one particular. Two (for now; more complex later):

- — The *domain of representations* (might call this the *signal* domain)
- — The *domain of things represented* (might call it a *value* domain, but perhaps this is misleading)

- Important. There are two components to the relation $\ll$ between P and B (when P$\ll$B) that play quite different roles:
  1. The *type-level relation* between $\mathfrak{A}_P$ and $\mathfrak{A}_Q$; and
  2. The *instance-level relation* between (a particular) P and (a particular) Q.

  The point is that the first sets up the basic correspondence between properties and relations. Then, given a particular P, the particular way in which P presents under these cluster $\mathfrak{A}_P$ can be used to determine how Q presents under the cluster $\mathfrak{A}_Q$. I.e., once the type-level correspondence is established:
  - — *what is particular to P represents what is particular to Q.*
- The specific type-level relation is a relation type: we'll call it $\ll_x$ for some appropriate x.
- Example: Turing machines with quadruples.
- Example: 0 and 1 used to represent polarity in Situation Theory.
- Example: Suppose ordered pairs of real numbers $\langle r,\Theta\rangle$ are used to represent points on a plane. Call this representation relation $\ll_{pt}$. First we need to establish $\mathfrak{A}_{pair}$ and $\mathfrak{A}_{point}$ (there is of course some latitude in how this is done; more on that later):

Properties of Whole Representation Relations

There are lots of global properties that representation may or may not have. And lots of questions to be asked about each. For example there are various notions having to do with intuitive notions of completeness.

Comprehensiveness

- if P captures everything there is to say about Q (at the appropriate level of abstractions $\mathfrak{A}_P$ and $\mathfrak{A}_Q$)

Totality

— if every element in the represented domain can be represented.

Ambiguity

— P is unambiguous if, for every element (object, property, or relation) of Q there is a single element of P.

- — Note on formalization: Might say that $\ll$ is *unambiguous* just in case:

  $$\forall a,b_1,b_2\ [[a\ll b_1] \wedge [a\ll b_2]] \Rightarrow [b_1=b_2]$$

  But this is strong: it says not only that the representation carries complete information at the level of abstraction at which it is defined, but in fact enough information to determine identity in the represented domain. More appropriately, we might say that $\ll$ is unambiguous just in case:

$\forall a_1, a_2, b_1, b_2$

$[[a_1 \ll b_1] \wedge [a_2 \ll b_2] \wedge$

$[a_1$ same as a_2 in all properties or relations in $\mathfrak{A}_a]]$

$\Rightarrow [b_1$ same as b_2 in all properties or relations in $\mathfrak{A}_b]$

But these aren't necessarily to be trusted. They certainly don't deal with completeness. For example, the first says only that the identity of object in the representing domain is as fine-grained as the identity of object in the represented domain. More appropriate axioms should be developed.

4. *Reduction*
   - $\ll_1$ reduces to $\ll_2$ just in case, for every P and Q such that $P \ll_1 Q$, there is an R such that $R \ll_2 B$ with as much information (or something like that).
5. *Supervenience*
   - between two levels of abstraction $\mathfrak{A}_P$ and $\mathfrak{A}_Q$: no difference in $\mathfrak{A}_Q$ without a corresponding difference in $\mathfrak{A}_P$.
   - i.e., $[P_1 \ll Q_1 \wedge P_2 \ll Q_2]$ wrt $\mathfrak{A}_P$ and $\mathfrak{A}_Q \Rightarrow [\, Q_1 \neq Q_2 \supset P_1 \neq P_2]$
   - May want to talk about P reducing to P' if $P \ll_1 Q$ wrt $\mathfrak{A}_P$, and $P' \ll_2 Q$ wrt $\mathfrak{A}_{P'}$, and if $\mathfrak{A}_P$ supervenes on $\mathfrak{A}_P$', and ...
6. *Effectiveness*
   - If $A \ll B$, then B can be effectively computed, created, accessed, whatever, given A and $\ll$.
7. *Typological*
   - A representation relation is typological if the entire representational burden is carried at the level of the corresponding types. (i.e., no distinguished individuals).
   - Example: representing distances as real-numbers is not typological, since a particular distance (the unit distance) must be specified. Similarly, representing points as $\langle r, \Theta \rangle$ pairs, for same reason.

— Note that many of these hold at the type level (i.e., for the '$\ll$' in question), and also for given representations.

F. Categories of Correspondence

— Goal is to identify and categorize properties of the '≪' relation.

1. Iconicity

- ≪ is *iconic* when every object, property, and relation in the represented domain is represented by an object, property, or relation in the representation domain, respectively.
  - — Example: {x, {x,y}} ≪ ⟨x,y⟩? Not quite.
  - — Example: (idealized) numeral arithmetic ≪ arithmetic.

2. Absorption

- An object, property, or relation is said to be absorbed by a given ≪ if it is used to represent itself.
  - — Example: photographic copies
  - — Example: linear order in linguistic representations of (the syntactic structure of) sentences.
  - — Example: left-to-right order in context-free grammars: S → NP VP.
  - — Example: if integers are represented by their doubles, then addition and subtraction are absorbed, though multiplication and division are not (quite).
- — Subtlety:
  - — When a property or relation is absorbed, there is a question about whether the objects filling the roles must be absorbed, or whether objects that model the represented role fillers are to be used. I.e., if R is an absorbed two-place relation (i.e., R≪R), and if A≪A' and B≪B', then do we have R(A,B) ≪ R(A',B'), or must R(A,B) ≪ R(A,B)? In the arithmetic case just described, the former is allowed (+(double(x),double(y)) ≪ +(x,y)). Similarly, when length is absorbed in a scale model of an airplane, the metric unit is changed.
  - — Suggestion: still call this absorption.

3. Objectification

- ≪ objectifies a property or relation in the represented domain when that property or relation is represented by an object in the representing domain.
  - — Example:

```
expression       ::= numeral | expression  op  expression
op               ::= '+' | '-' | '*' | '/'
numeral          ::= '0' | ['-']  non-zero-digit [ digit]*
digit            ::= '0' | non-zero-digit
non-zero-digit   ::= '1' | '2' | ... | '9'
```

In '3+14' the '+' sign, which is an object in the representation, represents a (two-place) relation in the represented domain.

- — The '+' sign does participate in relations with other objects in the represented domain (most saliently, in the "is to the immediate left of" and "is the the immediate right of" relations).
- — Note that in this example I absorbed left-to-right adjacency.

4. *Implicit/Explicit*
   - ☐ $\leqslant$ represents an object, property, or relation explicitly if it objectifies it (represents it with an object), implicitly if it represents it with a property or relationship
   - ☐ Thus a representation inevitably represents implicitly any property or relation it absorbs.
     - — Example: arithmetic language above represents addition explicitly.
     - — Example: Run length encodings represent the elements of binary arrays implicitly.
     - — Example: FOPC represents predicates explicitly.
5. *Polarity*
   - ☐ When an object's presence in the representation signifies the absence of a corresponding object, or an absence signifies a presence.
     - — Example: a scheme to represent table settings in a restaurant. Represent each place setting with a sequence of seven binary digits, for the plate, knife, dinner fork, salad fork, spoon, napkin, and glass, respectively. 0 means there isn't one; 1 means there is. Thus ⟨0,0,1,1,0,1,1⟩ might represent a place setting missing a napkin, wineglass, and dinner plate.
     - — Example: Numeral '0'
     - — Example: keys in mail slots of hotels.
     - — Example: LISP's representation of lists as dotted pairs. A list $e_1, e_2, \ldots, e_k$ is represented as the ordered pairs $(r_1 \,.\, (r_2 \,.\, \ldots\, (r_k \,.\, \text{nil})\ldots))$. This is a polar, recursive, componential representation, because nil represents "no more elements". I.e., $(r_i \,.\, r_j) \leqslant$ the list L if $r_i \leqslant$ the first element of L, and $r_j \leqslant$ the remaining elements of L, *except that nil $\leqslant$ the absence of any more elements.*
     - — Example: eof to represent there not being any more elements in an input stream.
6. *Componentiality*
   - — Intuition: there is a natural notion of "part" for complex objects in the representing domain, and there are objects in the represented domain such that, if A' is a part of A, and A$\leqslant$B, then ... various things; it's complicated.
   - — Basic idea of a domain of components or ingredients, put together by some method of combination (a kind of glue). Adjacency and parens in many formal languages; function composition in many standard abstract models; etc. It is not clear, however, whether there are natural notions of glue in many standard represented domains.

- The introduction of componential brings along with it the notion of contextual relativity, because suddenly some properties of the representation other than those that arise from its componential structure can play a role in determining what it represents. Without this, everything is context. You might think that context arises only in representation domains involving a notion of type/token, but I think in fact it is the other way around.
- First, part/whole business needs to be analysed: questions, for example, about whether something that can be a part can in fact be a representation on its own (sometimes no: open formulae, for example).
- Various strengths:
  - for every part A' of A, such that $A \ll B$, then there is a B' in represented domain, such that $A' \ll B'$, and the role that A' plays in representing B could be equally well played by any other representation A" such that $A'' \ll B'$. Stronger: only B' is relevant: need to know nothing about the A" beyond its representing B'.
  - even stronger: the componential reading is required for B too: B' must be a natural *part* of B. I think this is rare.
- *Recursive*: componential structure, such that ...
- Whether there is anything like *compositionality* (which is used *of representation relationships*, not *of representations*) depends on what is taken to exist in the represented domain.
- For example, if satisfaction sets etc. aren't part of the represented domain, then even first-order predicate calculus (which is extensional) isn't compositional.
- Similarly, take arithmetic language above (with '+'). Suppose take the sum relationship to be part of $\mathfrak{A}_A$. Then '3+14' isn't compositional yet, because '+' doesn't represent the sum type (not specific enough). Need it to represent the *function* from (pairs of) numbers to numbers. But that may not have been in the original $\mathfrak{A}_B$. Can extend $\mathfrak{A}_B$ to include functions (abstract objects). In fact I suspect many abstract objects are introduced to play exactly this role. This is why I am cautious about their existence: they may arise in aid of representation, not really as part of what is to be represented.

- *Other properties:*

  - *Aggregation*
    - Intuition: many things in represented domain are represented by a single thing in the representing domain.
  - *Abstraction*
  - *Componential* Structure and Compositionality
  - *Quantification*

- violates strict compositionality
- Introduces a certain kind of *local indeterminacy:* there is an ingredient in a complex representation *without a determinate representation.*
- For example, even if '∃x P(x)' correctly represents some situation in the world, the ingredient expression 'P(x)' doesn't represent anything particular in that situation (or even in the world).
- The simple arithmetic language used early is locally determinate, as is propositional calculus, and parse trees. Predicate calculus, however, is not.
  - What about disjunction? Ahhhhh.....
- Local indeterminacy introduces a certain kind of contextual relativity that we haven't seen before (distinction between meaning and interpretation?).
- You wouldn't, I take it, normally call a locally indeterminate signal a *model*— although local indeterminacy is clearly ok for representational systems (especially for languages).

— *Sentences*
- Double access story (too much information)
- Certainly no assertional force without talking about processes, agents, etc.
- Introduces the notion of truth, for the first time. Requires the possibility of being wrong.

— *Language*
- Does this have to do with structure or with use?
- Perhaps: finite, componential,
- Generated by an inductive set of formation rules?

— *Procedural representations*
- Intuition: story of $\preccurlyeq$ requires state on the part of an agent that establishes the representation relationship in a particular case.

G. Algebraic Basis

- An appealing (though not necessarily attainable) goal would be to identify a basis set of abstract properties and relations that '$\preccurlyeq$' relations could have: say R_1 through R_k. Then suppose we had $P \preccurlyeq_1 Q$ and $Q \preccurlyeq_2 S$, and suppose we could describe each of '$\preccurlyeq_1$' and '$\preccurlyeq_2$' in terms of this basis set. The goal would be, from these descriptions of '$\preccurlyeq_1$' and '$\preccurlyeq_2$', to be able to define a representation relationship '$\preccurlyeq_3$' componentially in terms of the R_1 through R_k such that $P \preccurlyeq_3 S$.
- For example (in terms of the properties described below), suppose that $\preccurlyeq_1$ quantifies and is recursive but not polar; and suppose that $\preccurlyeq_2$ is a weak isomorphism that objectifies one level. Then we could conclude that $\preccurlyeq_3$ is also not polar, but that it quantifies and objectifies, etc.

- Example: we ought to be able to predict what properties of the original photograph obtain in the copy; which ones don't; etc. Suppose in particular that the photograph-copy representation relationship $\leqslant_{copy}$ is a strong isomorphism at the level of abtraction at which the more general photograph representation relationship $\leqslant_{photo}$ is defined. Then if P is a photograph of scene S (i.e., $P \leqslant_{photo} S$), and P' is a copy of P ($P' \leqslant_{copy} P$), then the theory would enable us to prove that P is also a photograph of S ($P' \leqslant_{photo} S$).
- Example: suppose R_1 is a linguistic representation of a sentence S that objectifies certain syntactic properties, but absorbs linear order. Similarly, suppose that $R_2 \leqslant R_1$ in a way that also absorbs linear order. Then R_2 absorbs the linear order of S too, whatever else one may want to say about it.
- Example: Suppose $B \leqslant_1 S$, where $\leqslant_1$ is complete and effective way of representing the Dandelion screen (in particular, suppose that B is a bit-map), and suppose that $D \leqslant_2 S$, where $\leqslant_2$ is neither effective nor complete (for example, suppose D is a description of the text displayed on S). Suppose further that someone gives you a procedure P that, given S at the level of abstraction at which $\leqslant_1$ is defined, will give you D. It should follow that an account of $\leqslant_1$ and $\leqslant_2$ should enable you to construct a different procedure P' that would *translate* bitmaps B into descriptions D. (This would obviously be useful in practice because of the fact that the bitmap is entirely connected to the screen. But P and P' should still be distinguished.)

— These algebraic properties are (maybe) part of the long-range goal; we're not there yet.

H. An Extended Example

— Points on a plane. Suppose ordered pairs of real numbers $\langle r,\Theta \rangle$ are used to represent points on a plane. Call this representation relation $\leqslant_{pt}$. First we need to establish $\mathfrak{A}_{pair}$ and $\mathfrak{A}_{point}$ (there is of course some latitude in how this is done; more on that later):

- $\mathfrak{A}_{pair}$ consists of an abstract three-place relation (we'll call it "pair") among: two-element ordered pairs, their first elements, and their second elements.
- $\mathfrak{A}_{point}$ consists of a designated point on the plane (the origin: we'll call it P_0), a designated line emerging from from that point (roughly, an orientation) called A_0, a relational notion of distance between two points, and a similarly relational notion of angle between two lines.
- A distance between two points can in turn be represented as a real number. This is another instance of representation; all the same points arise. In particular, the distance is not itself the real number, so our first axiom is honoured (it would be perfectly rational to be daunted by the distance between you and the nearest source of water, but irrational to be daunted by a real number). So we can characterize this representation relationship as

well. First we have to establish $\mathfrak{A}_{point}$ and $\mathfrak{A}_{real}$. This involves, among other things, designating a *particular* distance to be used as the appropriate unit. Only with respect to that predesignated distance can the real number carry any information about the actual distance. (You may want to say that the *relative magnitudes* of one distance and another are *actually* real numbers; they don't haave to be modelled (represented) by real numbers, so that representation relation may bottom out one level down). We will assume that this [real-number $\preccurlyeq$ distance] story has been appropriately spelled out; we'll call it $\preccurlyeq_d$.

— Similarly, assume we have a [real-number $\preccurlyeq$ angle] relation called $\preccurlyeq_a$ (again, relative to a $\mathfrak{A}_{angle}$ including a designated particular angle to be used as a unit, kind of a "fulcrum" on which to convert between numbers and genuine angles. Also, we need the notion of a line deteremined by two points: we will designate this with line(p_1,p_2)

— So the type-level story is then roughly as follows:

$$
\begin{aligned}
X \preccurlyeq_{pt} P \leftrightarrow \quad & \text{pair}(X) \land \\
& \text{elements}(X,r,\Theta) \land \\
& \text{real-number}(r) \land \\
& \text{real-number}(\Theta) \land \\
& r \preccurlyeq_d \text{distance}(P,P_0) \land \\
& \Theta \preccurlyeq_a \text{angle}(\text{line}(P,P_0),A_0)
\end{aligned}
$$

— What is particular to any given point representation (pair), then, is merely the two real numbers r and Θ. The rest is carried by the type story just given.

— What will we want (in the long run) to say about $\preccurlyeq_{pt}$? Roughly, that it is:

  — complete
  — total
  — objectifies, since $\preccurlyeq_d$ and $\preccurlyeq_a$ objectify by representing a relation (distance and angle) with an obejct (a real number).
  — componential (roughly, parts of the pair directly ...)
  — locally determinate

but that it is not

  — polar
  — recursive
  — quantificational
  — etc.

I. Summary and Conclusion

☐ "Semantic Soup", or the "continuum of correspondence".

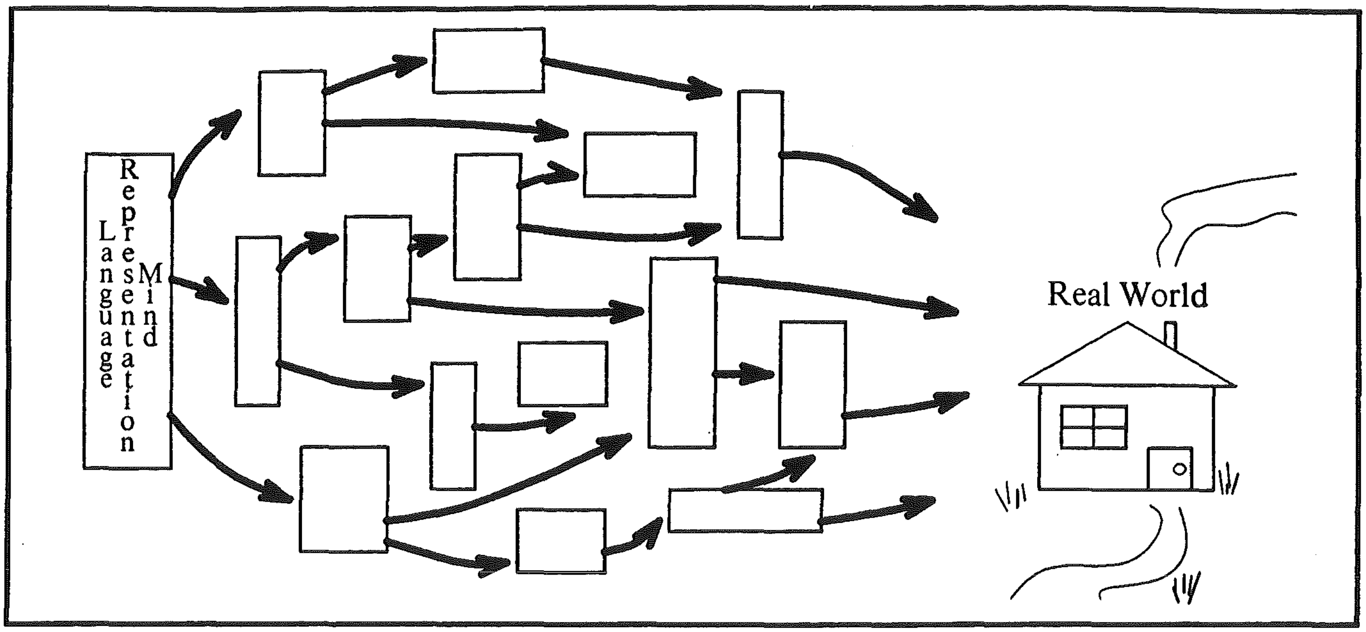

- — Point is that, whether one does things *syntactically* (i.e., using a meta-language) or *model-theoretically* are less different than is normally supposed.
  - — As a consequence of the strong α-relationships pointed out at the beginning.
- — I.e., whether a relation is linguistic or not *isn't a clear question.* All of us (most, at least — determined by introspection) thinnk that the world at the right doesn't comme parsed or registered in advance. So its registration, or *it registered,* is slightly to the left.
- — The simplistic (strict dichotomy) view generates all kinds of debate and heat.
- — On my account, this is *epiphenomal heat.*

☐ Other parts of the problem.

1. Use and inference.
   - — Big subject, not yet treated here.
   - — Intuitively, given that you have one representation R_1 that represents some situation or part of the world S_1, can you perform operations or transformations on R_1 so as to derive R_2, such that R_2 will represent some situation or part of the world S_2? I.e., can you deduce, from how things are in a "local" situaton (viz, for which you have a representation) how they are in a "remote" situation (viz, for which you don't yet have a representation)?
     - — Doing so I will call *inference.*

- Why is it possible?
  - Because of constraints and regularities in the world, to which the representation is attuned (that it represents? to which it is connected? lots of things going on here).
- Kinds of operation
  - Substitution, comparison, unification, etc.
  - Cf. recent work of John Lamping.

2. Connection / Disconnection.
   - What makes a representation a representation (gives it its semantic bite? authority?)

3. Registration.
   - How is the world represented by an independent agent in the first place: "parsed", according to some conceptual scheme?