Introduction to Representation

Overview

- We have suggested that serious NLP tasks require representing the meaning of an utterance.
- Example: We hear
  - "Jan picked up a rock. She broke a window."
- A representation should enable us to make clear
  - "She" and "Jan" refer to the same thing.
  - The rock Jan picked up was what she used to break the window.
- I.e., make explicit what is implicit in our understanding.

The Language of Thought

- We will need a *language* in which we can encode what a person is thinking having heard an utterance, or what a person might have been thinking when uttering something.
- The system that the mind presumably uses for thinking, and especially, for encoding the concepts and conceptualizations underlying a natural language utterance, has been called
  - The language of thought
  - *Mentalese,*(Fodor)
  - Internal representation language (IRL).
- Specifying exactly what this language is is the *meaning representation* problem.

Aside

- There is much to debate philosophically about the nature of the IRL, including whether such a notion is sensible at all.
- We will not confront these issues head on for now.

Meaning Representation and Knowledge Representation

- We will need to write down facts about the ideas underlying an utterance, so we have the *knowledge representation* problem too.
- We will assume that these are essentially variants of the same problem.
  - Particular utterances will generally encode more specific facts, but we'll assume no principled difference.

Benefits

- Can give a fairly concrete rendering of what we mean by terms like "understanding":
  - "Understanding" becomes computing an IRL formulation of an utterance.
  - "Production" (or "generation") becomes encoding ideas (in the IRL) into NL utterances.
  - "Thinking" becomes the manipulation of IRL statements.

Properties of an IRL, Generally

- *Declarative*
  - which is really just another way of saying that it is a *language* in which we can write statements or other expressions and then interpret them.
- *Compositionally*
  - The meaning of an expression is determined by its components and rules that say how these go together.
- *Independent of context*
  - Without looking at the context (speaker, hearer, location, time, previous utterances)
- Expressively *adequate.*
  - *epistemological adequacy :* The language must be powerful enough to express in it everything we would want to say.
  - *heuristic adequacy:* It must be usable to perform the desired tasks.
    - weak, strong, and additive (McDermott)
    - *weak* -- system displays behavior which demonstrates that it possesses certain knowledge.
    - *strong* -- A system is given facts in a language and can use them
    - *additive* -- A system can be given facts incrementally, and reconfigure itself to use them.
- *Uniform.* Don't create different languages for different domains, unless some difference in functionality compels us to do so.

So What About

- LISP? (I.e., programming languages)
  - Compositional, context independent, uniform, maybe expressive, but …
- English? (Or Sanskrit, or your favorite NL?)
  - Expressive, but

Artificial Languages

- Logic
- Semantic Networks
- Frames

The Problem

- Two parts:
  - What needs to be represented?
  - How exactly do we represent it?

Let's Start With a Top-Down View

- What kinds of underlying meaning components might there be?
- Linguistically, can look for simpler and more complex language forms:
  - "Jan broke the window."
  - "Shouldn't Jan have quickly broken the window with a hammer yesterday?"
- The intuition is that there are more or less simple core "situations", from which we build up more complicated ones.

Core and Elaborations, Crudely

Meaning Component

loosely a reflection of

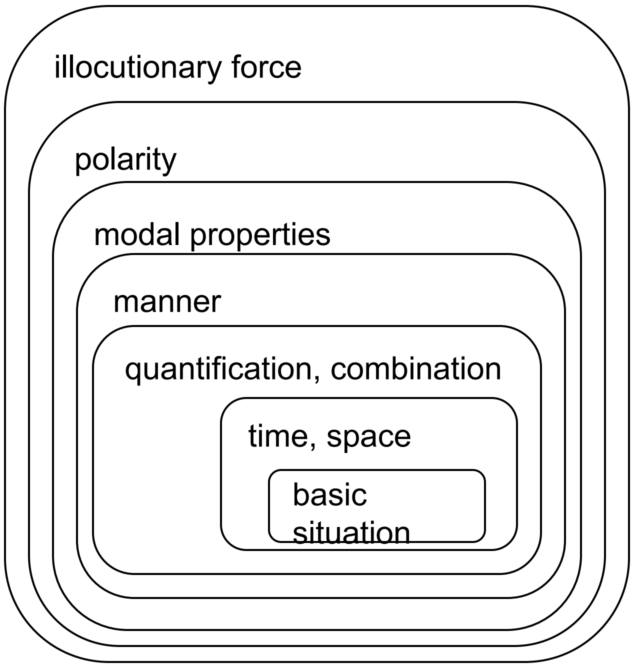

grammatical devices, lexical items

various negation forms

modal auxiliaries

manner adverbials, verb incorporation

quantifier terms, logical connectives

tense, aspect, temporal adverbials; prep. phrases

verbs and their complements and adjuncts

Picture is Really Quite Messy

- Semantic components can get conveyed by a variety of different linguistic objects.
- They can get clumped together in different ways, depending heavily on the language.
- E.g., situation type is typically conveyed by a verb (conflated with some temporal information, etc.), it may equally well come as a noun (e.g., "the debate, "the construction")
  - incidentally, without the temporal information.

Lots of Cross-linguistic Messiness

- English (and, supposedly, all of Indo-European except Romance) and Chinese often incorporate manner and cause into the verb
  - "shout", "whisper" incorporate manner ("vocalize loudly/softly"; "crawl", "drive", "swim" describe manners of motion)
  - "flatten" is "cause to be flat"; "kill" is "cause to die"
- Romance tends to incorporate direction
  - which we see in English, "ascend", "enter", "leave", mostly from the French…
- Some languages (notably Atsugewi and Navajo) incorporate object information into verb
  - which you also see in English with verbs like "spit" (and "drink", as in "For my New Year's resolution, I quit drinking.")

Kinds of Meanings

- Philosophers have distinguished between "sentence meaning" and "speaker meaning".
- *Sentence meaning* is supposed to be:
  - assigned to sentences independent of context, speaker, etc. (There may be multiple potential assignments, in which case the sentence is ambiguous.)
  - compositional (in terms of the lexicon and the grammar)
  - determined by (or equal to) truth conditions
  - equivalent to "literal meaning"
- *Speaker meaning* is supposed to be:
  - what a speaker means by using the sentence
  - subject to various kinds of non-literal interpretation, e.g., metaphor, idioms, speech acts, literary devices, sarcasm, etc.
  - equivalent to the sentence meaning when there is no reason to depart from it.

Examples of Speaker Meaning Departing from Sentence Meaning

- "It's cold in here."
  - In the right context, might be a request that the listener turn on the heat, close a window, etc.
- "Jan kicked the bucket."
  - Intended (non-literal) interpretation may be about dying.
- "That was a brilliant idea."
  - When uttered sarcastically, used to mean opposite of its sentence meaning.
- "No man is an island."
  - Use is no doubt metaphorical (and hence, is a counter-example to assuming non-literal interpretations attempted only when there is some semantic conflict.)
- "I'm out of here."
  - Recent form in which the meaning is false in the present to indicate it coming true in the immediate future.

Components of How Utterances are *Used*

- Austin (an "ordinary language" vs. "ideal language" philosopher) made the following distinctions:
  - *Propositional content*
  - *Illocutionary force*
    - These are the "universal and inevitable consequences of an utterance".
    - E.g., question, assertion, command, request, promise
  - *Perlocutionary force*
    - non-universal effect of the utterance
    - E.g., convince, persuade, insult, inform

Where Do These Fit In?

- "Propositional content" is all but the last oval,
  - which is the "illocutionary" force.
- It's not that we want to ignore perlocutionary force.
- It's just that this is completely inferential, without a linguistic "anchor".

Sentence/Speaker is Probably Fatally Flawed

- A decent account of a language's grammar and lexical should let us find for a sentence
  - conventional, but non-compositional interpretations (e.g., idioms)
  - conventional, but metaphoric interpretations
  - its illocutionary force
    - (was it a question, a command, etc.?)

Sentence/Speaker Flaws

- On the other hand, just the grammar+lexicon may produce something that is too abstract to serve as a meaning.
  - Consider "The cat is on the mat."
    - which means "lying", but this may require an inference from lexical entry for "on".
  - Noun-noun compounds
    - "grass seed" vs. "bird seed", "coast road"
  - Adjectival modification
    - "red pen"

An Alternative: Primal Versus Actual Content

- Primal/Linguistic Content
  - What one can derive from the lexicon and grammar
  - Both broadly construed to include
    - non-productive forms (e.g., idioms)
    - conventionalized metaphor
    - illocutionary force
  - But with might not be by itself a candidate for the meaning of the utterance.

Primal Versus Actual Content

- The content part of speaker meaning
  - So doesn't include result of indirect speech act inference
  - Is non-compositional, requiring extra-linguistic knowledge and inference to determine that "red pen" means "pen writing red", etc.
- The actual content assigned by the *sensible speaker/hearer* is *ordinary content*.
- It's probably ordinary content that is of primary interest to us.
- But there may be many parts/stages of meaning representation.

We'll Start with Representing Core Situations

- Situations have *participants*.
  - For example, our "breaking" situation involves Jan, a window, and a hammer.
- Participants, as well as the situations to which they belong, are generally *individuals* of some *category*.
  - These are sometimes also called *tokens* and *types*, respectively.
  - "Jan" might be an individual "Person", "a window", an individual "Window", etc.
- Situations involve participants of particular categories in specific ways. E.g., breaking involves something that got broken, which has to be a physical object.
- Participants, and, perhaps a bit less obviously, situations, generally fall into a *hierarchy* of *supertypes* and *subtypes*.
  - E.g., categories for "Person", "Man", "Women", "Physical-object"; "Man" would be a subtype of "Person", etc.

A First Attempt

- In the late '60s, Quillian proposed a *semantic memory* in which
  - nodes represent concepts
  - links represent relations between concepts.

Early Semantic Memory Example

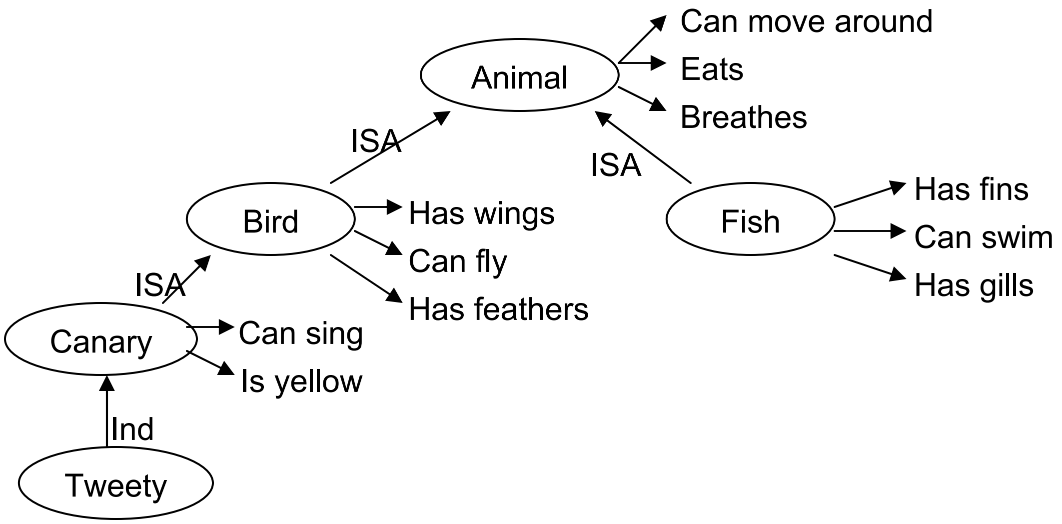

- Introduced *hierarchies* and *inheritance*.
- Motivated by computation economy (store "Can fly" once, with bird, inherit)
- Psychological reality? Did it take longer for humans to answer "Do canaries have wings?" than "Do birds have wings?"
  - Initial results suggestive, complicated by "penguins".

Early Semantic Memory (con't)

- No clear idea what "properties" should be or how to represent them.
- However, while "situations" seemed more complicated, they had more or less statable sets of participants.
- Led to the extension of concepts having "roles".

Example: Adding Roles

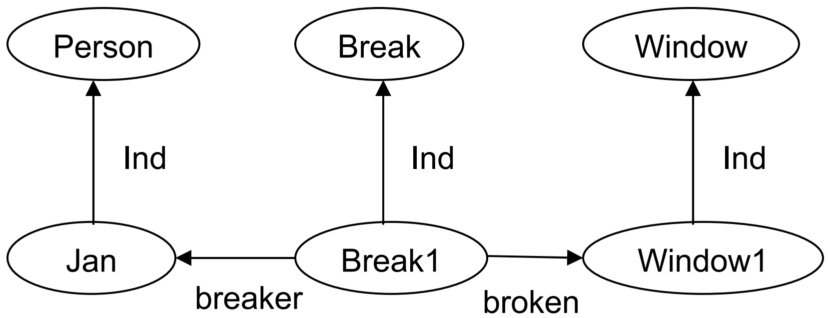

- This elaboration is called a *semantic network*.
- A *network* is a set of nodes connected with directed arcs.
- In semantics networks,
  - node correspond to word senses
  - arcs designate semantic relationships between a verb and its surrounding nominal concepts

In Parallel, Logic

- A logic gives us
  - A language
    - With rules of composition
  - Semantics
  - Inference procedure
- We usually begin with propositional calculus and then FOPC.
  - These have officially nice inference procedures.

FOPC

- Object symbols
  - e.g., Jan1, x.
- Predicate symbols
  - e.g., Mother, Tall, >, Cost
  - which has an fixed "arity"
- (optionally) Function symbols
  - e.g., Mother-of
  - syntactically look like predicate symbols, but act like object symbols

FOPC (con't)

- *Terms*: either object symbols or *function applications*
  - e.g., Mother-of(Jan1)
- *Sentences*:
  - *atoms*, a predicate applied to some terms
    - Mother(Jan1,Lynn1)
    - Tall(Father-of(Jan1))
    - Taller(Jan1,Lynn1)
  - *logical sentences*, i.e., sentences connected with logical connectives, generally, $\neg$, $\vee$, $\wedge$, $\rightarrow$.
  - *quantified sentences*, generally using the quantifiers $\forall$ and $\exists$.
    - $\exists$ x Mother(x,Lynn1)

How to Apply to Situations, Etc.?

- There isn't a single obvious way to FOPC notations to core situation/linguistic needs.
- We'll come back to this in a while.

Formal Limitations

- In FOPC, we can't quantify over predicates (or functions).
- For example, it is reasonable to want to say that two functions are equivalent if they have the same value for all their values, i.e.,
  - $\forall$ f,g (f=g) $\Leftrightarrow$($\forall$ x f(x)=g(x))

  But you can't: This is formally strictly verbotten.
- We do this is NL all the time:
  - "Everything that was true of Russia under czars was true under the communists."
  - "Pat has all of Lynn's fine qualities, plus more".
- For such utterances, we can say them naturally using a higher-order logic, or try to find some way to finagle them into FOPC.
- We know that higher-order logics are strictly more expressive, and that, e.g., there is no effective reasoning procedure.

Propositional Attitudes

- In general, words like
  - "believes", "wants", "thinks", "dreams", "knows", "should", "can", "ought"
  - take *propositions* as their objects.
- One can't map these to predicates, because predicates can't take propositions as objects.
- Instead, one can introduce special purpose "operators" for each of these, and produce to develop various kinds of "modal logics".
- One then tries to write down axioms for these operators, but this has proven hard to do.

Modal Logic Example

- For example, one often sees axioms like the following (which I'll express informally, without quantifiers and as implications):

  $K(a,p \rightarrow q) \wedge K(a,p) \rightarrow K(a,q)$

- But this is tantamount to saying that one knows all the implications of what one knows, which seems unlikely.

Model Logic (con't)

- The situation is worse for operators like "Believes".
- One is tempted to write similar axioms, but it is hard to interpret such "axioms" as anything other than prescriptive, which seems to undercut the claim that we are in fact representing the commonly used concept.
  - People are capable of believing all sorts of sentences together.

Quantifiers

- NL has an indefinite number of quantifiers. E.g., consider:
  - hardly anybody
  - a few
  - few
  - very very few
  - several
  - most
  - many
- There are various proposals about how to deal with these.
- Note, though, that no solution that is purely truth-conditional seems adequate.

Degrees of Truth

- Classical logics do not admit decrees of truth, so sentences like
  - "What you say is mostly right."

  are hard to represent with them.
- Again, various attempts to deal with this (e.g., fuzzy logic), but none have met with wide acceptance, at least for NLP.

What To Do?

- We generally won't be interested in “the right reasoning procedure”, but we will need the expressive power.
- So we can use higher-order logics and try to find what reasoning capabilities we need.

Semantics: The Standard Story

- We form an *interpretation* by mapping
  - object symbols to objects "in the world"
  - predicate symbols to relations "in the world"
  - function symbols to functions "in the world"
- Then sentences may or may not be *true relative to an interpretation*, or *satisfied*.
- If sentences have no variables, we can recursively define satisfaction this way:
  - Atomic sentences are true if the relation holds "in the world".
    - (If we have the special predicate called "equality", then this is true if the two arguments designate the same entity "in the world".)
  - Logical sentences are true as a truth-theoretical function of their parts.
- For quantified sentences:
  - A universally quantified sentence is satisfied if the enclosed sentence is true for all variable assignments.
  - An existentially quantified sentence is satisfied if the enclosed sentence is true for some variable assignment.

Semantics?

- We have avoided saying what sentences *mean*, only that they may be true or false.
- Logicians typically claim for a close relation between the two.
- For example, the meaning of a sentence is usually identified with its truth value.
- This is probably a terrible idea.
  - Indeed, Putnam has showed that it renders truth-conditional semantics self-contradictory.
- What we really should say is that "Mother(Jan1,Lynn1)" means that "(Jan1,Lynn1)" is in the "Mother" relation, etc.)

Semantics For Modal Logics

- Modal operators have no interpretation in the straightforward (i.e., Tarskian) semantics we have been using.
- Attempts to provide semantics for same lead to *possible world semantics*, which, although we won't provide an analysis here, seems to make assumptions of logical consistency, etc., which are quite dubious.

Enter Frames

- In 1974, Minsky published a very influential paper, "A Framework for Representing Knowledge".
- Minsky argued:
  - There is a need to organize knowledge into *larger chunks*, as opposed to lots of separate PC formulas, say.
    - » E.g., our knowledge of rooms, birthday parties, comprise lots of coherent information.
  - Many researchers had arrived at similar conclusions in their respective fields,
    - » in particular Schank, Abelson and Norman. But actually, the idea is much older, going back to at least Barlett's notion of *schema*, and (arguably, to Kant.)
  - There was a super-general way to think about this notion of larger chunks of knowledge, namely, *frames*.

A Frame Is

- A data structure for representing a stereotyped situation,
  - E.g., being in a certain kind of room, or attending a child's birthday party.
- A network of *nodes* and *relations*.
  - Top-levels of the frame are fixed (representing what is invariantly true),
  - lower levels have *terminals* (i.e., slots) filled in specific instances of the frame.
  - Slots might be restricted to be filled by a certain type of object, e.g., a person or an event, or to having a certain relation to another slot.
  - Slots get to specify *defaults*.
- Frames where arranged in collections called *frame systems*.
  - Different frames in a system share the same terminals, so the frames in effect constitute different points of view of the same object.

Using Frames

- Minsky claimed that lots of intelligent tasks involve
  - finding a frame suitable for a given situation (perhaps taking into account one's goals)
  - and *matching* it to the situation.
- Minsky intended frames as an antidote to logic-oriented approaches, emphasizing:
  - Selecting relevant knowledge, i.e., the importance of knowledge about how knowledge is used
  - The importance of making assumptions over proving conditions
  - The non-monotonicity of real reasoning
  - The inappropriateness of logical properties like consistency and completeness

Impact

- Minsky's work on frames had several important impacts:
  - It unified what many researchers were already doing.
  - It gave rise to frame languages, notably:
    - KRL (Bobrow and Winograd)
    - FRL (Goldstein and Roberts)

What Frames Actually Meant for Representation

- *Objects*, corresponding to concepts/situation categories
- *Slots* bearing restrictions associated with objects
- *Hierarchies* and *inheritance*
- *Defaults*
- *Procedural attachment*
- Perhaps some ontological distinctions
- Syntactic sugar

Example of Concepts in a Generic Frame Language (actually, KRL)

(Travel a type of Event with
  (traveler (a Person))
  (destination (a Location))
  (origin (a Location) Default (a Home))
  (mode (a Vehicle)))

(Visit a type of Social-interaction with
  (visitor (a Person))
  (visitee (a Person)))

Now We Can Define An Individual

```
(Event137
    an individual Travel with
        (traveler Jan)
        (destination SanFrancisco)
        (mode Plane)
    an individual Visit with
        (visitor Jan)
        (visitee Lynn))
```

Alternatively,

(an individual Travel with

(traveler Jan)

(destination SanFrancisco)

(mode Plane)

an individual Visit with

(visitor Jan)

(visitee Lynn))

- with token symbols generating internally.

Using Procedural Attachment

- If we know home of individual traveler, we might want to have it automatically inserted in the origin slot of an individual travel event.
- We could do so by
  - attaching a procedure to "origin" slot,
  - Which, whenever instance of event is created, checks to see if home is known, and, if so, inserts it in "origin" slot of new event.

Frames for Persons, etc.

- Would have similar objects for “Person”, “City”, and “Jan”, “Lynn” and “SanFrancisco”.
- For example:

```
(Person a type of Thing with
    (name (a Person-name))
    (address (an Address))
    (age (a Number (Btw 0 120)))) ; or maybe, an Age?

(an individual Person with
    (name (an individual Name with (firstname “Jan”))))
```

Although exactly what to include as slots is not clear…

Frames and Semantic Networks

- Frame systems had frames and roles
  - and defaults, procedural attachment, constraints on role fillers
- Semantic networks had concepts and relations.
- Semantic networks incorporated the new features, mostly by *reifying* roles.
  - Make roles object, say lots of things about them.
- Exemplified by languages like KL-ONE (Brachman et al.)

Example: The Concept of an Arch in KL-ONE

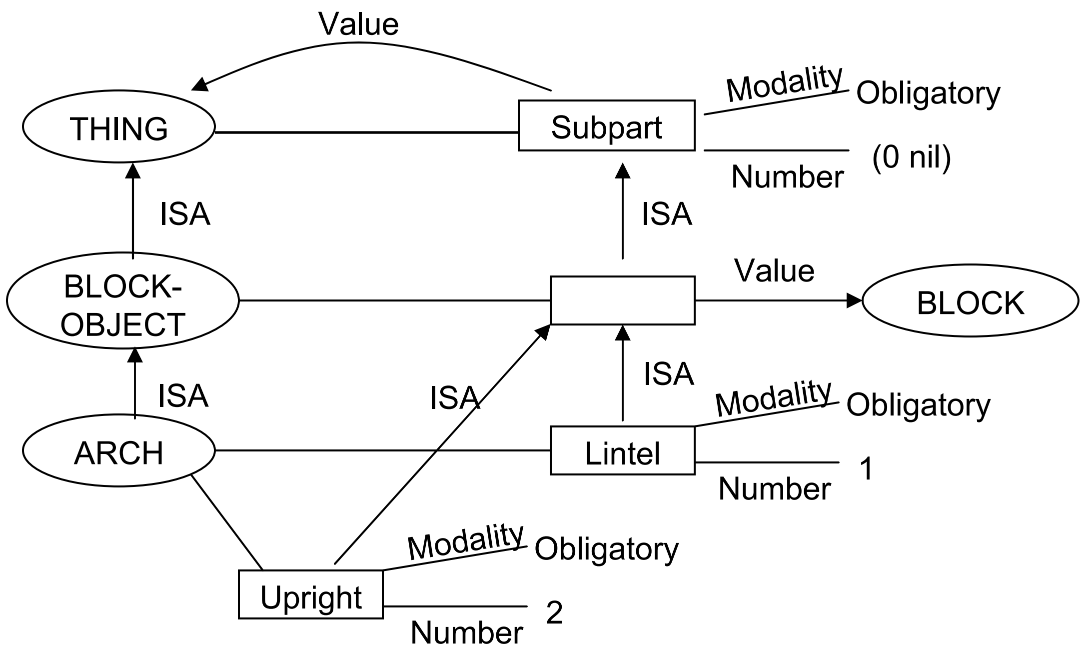

- Ovals represent concepts, boxes represent roles.
- Lines from roles provide information about roles:
  - **Value** states that concept playing role for an individual of concept bearing role is restricted to be of a certain type.
  - **Modality** specifies whether role is optional or required.
  - **Number** states number of such roles that can appear in individuals of concept.

Inheritance Revisited

- The insight here is that inheritance requires a *set* of links including
  - links between roles
  - rather than just a link between concepts,

  to establish the relation of one object in a hierarchy to another.
- This set of links is referred to as a *cable.*
- A similar cable with **Individuate** links would link an individual concept and its component roles to the concept it individuates.
- This idea is referred to as *structured inheritance.*

Limitations

- Some kinds of relational information between roles could be accommodated in this spirit.
  - For example, it is possible to link two roles to show that their fillers must be the same.
- On the other hand, some relational information is not so easily represented.
  - For example, we have not represented the crucial relational fact that, in a arch, the uprights support the lintel.
  - To do so in general may require an escape into predicate calculus.

More Problems

- Questions:
  - What does it mean to be a slot in a frame?
  - What can be a slot in a frame?
  - How do we express certain concepts?

What Does It Mean to Be a Slot in a Frame?

- You don't get to express what relational concepts like "Age-of" "Address-of" mean, other than by how they appear in frames and how they are "used by the system".

What Can Be a Role Where?

- While determining the roles for situation concepts (i.e. "participants") in was more or less clear, the roles for objects, etc., were much less clear, and perhaps, hopelessly ill-defined.
- The intuition is that a concept "has" a slot. But what is the nature of this intuition?
- If we allow "Name" in "Person",
  - We should allow "Mother" and "Father" also.
  - But then, what about "Uncle", or "Mother's-maiden-name" or "Favorite-Robert-De-Niro-Movie"?
- I.e., "has" represents very different relations, and is context dependent.
- Frame aficionados would often say that the presence of any particular slot is an efficiency consideration.
  - But then frames aren't a knowledge representation, then, but just a convenient indexing (and data entry) scheme.

Hard to Express Important Concepts This Way

- E.g., "Age" is usually a slot, and therefore gets defined procedurally.
- But "Age" is a perfectly good concept.
- Indeed, it is easier to define than "Person" (but not just is a frame/semantic network notation).

Today

- It is probably the case that most people regard such formalisms as shorthands for logical notation.
- Logic as the assembly language of representation.
- Let's go on with that.