Types and Lexical Semantics

Nicholas Asher

CNRS, Institut de Recherche en Informatique de Toulouse, Université Paul Sabatier

Cambridge, October 2013

Some background

- Words, sentences express meanings, concepts, or thoughts (Plato, *inter alia*). What are concepts, thoughts?
- Meanings are support inferences. How does that work?
- Dominance of symbolic and denotational methods in theoretical syntax, semantic and pragmatics. (Frege, Russell, Tarski, Montague).
- Exploited a close connection between meaning and truth.
- Used logic to account for inference.

Moving to computational linguistics

- 1980s and early 90s symbolic methods also dominant in computational linguistics.
- A sea change in computational linguistics in the 1990s. Move to probabilistic and machine learning methods over large corpora.
- Necessitated by the brittleness and lack of coverage of symbolic systems and enabled by machine readable text (web) and more powerful computers.
- How did this sea change affect semantics, pragmatics and the philosophy of language.
- explore an interaction between modern computational methods for for lexical semantics, where traditionally formal semantics has had little to say.. The meaning of 'cat' is $\lambda x \, \mathrm{cat}'(x)$

Talk Outline

- what can formal semantics do for computational linguistics and what can computational linguistics do for formal semantics?
- Kant's slogan: "concepts without data are empty, data without concepts are blind"
- introduction–types
- two problems: meaning composition and lexical content
- two levels of content—internal and external.
- specifying internal content

Introduction to types

For Montague–two types, E and T.

But lots of subtypes of E

count vs. mass, kind vs. individual, abstract vs. concrete (informational object vs. physical object), eventualities vs. objects propositions vs. facts vs. eventualities, locations vs. objects, different subtypes of eventualities (telic vs. non-telic)

1. John swept the kitchen
2. John swept the dust
3. #John swept the dust and the kitchen.

Types and meaning composition

- Types play an important role in reducing ambiguity and in meaning adjustments that result from combining meanings of words in a predication.
- dual aspect nouns: $\lambda y \colon \text{PHYS-OBJ} \bullet \text{INFO-OBJ} \ book(y)$.
- specifiable predicates:
  $\lambda y \colon \text{LOCATION} \vee \text{PORTION-MATTER} \ \lambda x \colon \text{AGENT} \ sweep(x,y)$.
- work on difficult cases of predication: coercion, copredication with predicates involving incompatible types (Asher 2011)
    1. The bottle had a nice label and was yummy.
    2. The book has a nice presentation but is very boring.

Types and meaning composition

- How many types are there?
- What types affect meaning composition?
- How do these types affect meaning composition?

Some answers: Types and selectional restrictions

1. The number 2 is soft.
2. Is the number 2 soft?
3. The number 2 is not soft.
4. If numbers could have textures, the number 2 would be soft.

Working assumptions

Types and presuppositions

predicates presuppose types of their arguments (selectional restrictions); types of arguments must justify these types.

- selectional restrictions are presuppositions about the types of arguments.
- type presuppositions flow from the predicate to the argument
- There are a variety of ways that arguments can meet the selectional restrictions of predicates (presupposition justification, cf. Heim, van der Sandt *inter alia*).

More on type presuppositions

Distinguish between necessary falsity and semantic anomaly (type presupposition failure)

- a Tigers are animals.
- b Tigers are robots.
- c #Tigers are financial institutions.
- d #Tigers are Zermelo-Frankel sets.

- Many philosophers take (a) to be necessarily true and (b) to be necessarily false
- Nevertheless, according to most people's intuitions, a competent speaker could entertain or even believe that tigers are robots.
- Much harder to make sense of a competent speaker's even entertaining that tigers are literally financial institutions, or ZF style sets.

Distinguish between type presuppositions and fine grained types

- type presuppositions of predicates are typically general—INFORMATIONAL OBJECTS, PHYSICAL OBJECTS, EVENTUALITIES, AGENTS
- lexical items have finegrained types subtypes of general types.
- finegrained types yield finegrained shifts in meaning.
  - adjectives: *flat tire, flat country, flat beer*
  - verbal modification: *load the hay on the wagon, load the wagon with hay*

Which way do the presuppositions go in the NP?

- From MG: Adjectives are functors from NP meanings to NP meanings
- *soft number* has a type clash, but which is the predicate and which the argument?
- *flat tire, flat country, flat beer*
  looks like the predicate is the Noun.
- material modification: *stone jar, wooden jar, wooden zebra, stone lion, chocolate lion*.
- handle finegrained shifts via a parametrization of the adjectival meaning.

Adjectival modification of nouns

Preserves the general type of the noun. Adjectives are arguments to nouns and conform to the type presuppositions of the noun.

The basic picture

- a rich system of atomic types including E, T and subtypes thereof. Partially ordered under a sub typing relation.
- Complex types that encode instructions for type interaction $\bullet$ types, ε, δ coercion types.
- To pass the presupposition from the noun to the modifier, use a presupposition parameter (π) that acts like a left context parameter in continuation semantics.

Meaning Entries

- tree: $\lambda \mathscr{P} \lambda x \lambda \pi \ \mathscr{P}(\pi * \arg_1^{\text{tree}} : \text{P})(x)(\lambda v \lambda \pi' \text{tree}(v, \pi')$
- heavy: $\lambda P : 1 \ \lambda x : \text{E} \ \lambda \pi'' \ (P(\pi'')(x) \wedge \text{heavy}(x, \pi'' * \text{ARG}^{\text{heavy}} : \text{P}))$

DPs: $\exists X \sqsubseteq \text{E} \ (X \Rightarrow (\Pi \Rightarrow \text{T}) \Rightarrow (\Pi \Rightarrow \text{T}))$.

Determiners in English encode mass and count type presuppositions:

- $\lambda P : 1 \lambda Q : 1 \lambda \pi \ \exists x (P(\pi * \text{ARG}_1^P : \text{COUNT})(x) \wedge Q(\pi)(x))$

When building up a λ term for a DP, we will typically get a sequence of type presuppositions on the variable bound by a determiner: those given by the verb will have to be justified jointly with those given by the head NP and by the determiner.

Basic rules

Definition

Basic Type Justification
If the type presupposition of a predicate *F* and that of its argument *a*
are compatible, then the type of *a* in the predicational context of *F*
is the meet of the two types. If they are not compatible, the type of *a*
in *F* is undefined.

I'll have a Chardonnay.

The type presupposition of the determiner is COUNT, but *Chardonnay* is
neither MASS nor COUNT
We get a modification of the noun by the determiner.

Predictions

Adjectival modifications involving the adjective *heavy*:

- Suppose it combines with *tree*. Types match and λ reduction works as desired.
- Now combine *heavy* with *number*.
- There are two incompatible typings on the same variable. We get an irresolvable type clash.
- The derivation crashes and no well-formed lambda term or interpretation results.

Extensions

- Some predicates will license operations that allow us to "wrap" the argument with a functor that supplies the right type to the predicate (coercion).
- Some nouns can supply different aspects to suit the demands of a predicate (dual aspect nouns).

A difficult case: material modifiers

Material adjectives like *wooden* and nouns like *glass*, *stone*, *metal*, etc. supply the material constitution of objects that satisfy the nouns these expressions combine with:

- glass (wooden, stone, metal, tin, steel, copper) bowl

Material modification can affect the typing of the head noun.

1. stone lion (vs. actual lion)
2. paper airplane
3. sand castle
4. wooden nutmeg

When the constitution of the object is given by an adjective whose denotation is not a possible type of constitution for the type of object denoted by the head noun, we get a shift in the type of the head noun.

Why a shift in type?

It supports different sorts of inferences.

1. A stone lion is not a lion (a real lion), but it looks like one.
2. A stone jar is a jar

Previous approaches

Davidsonian extensionalism

Adj Noun → Adj ∧ Noun

Gets things wrong with *stone lion* or *wooden nutmeg*/

Kamp-Montague intensionalism

Adj Noun → a new type of NP *Adj*(*Noun*).

Very incomplete. Stone jars aren't necessarily jars.

Polymorphic types

Change Davidson's logical form slightly

Have the type presupposition of the adjective affect the type of the noun.

. .

1. $\lambda P \, \lambda x \lambda \pi \, (P(\pi * \text{MADE OF}(\text{HD}(\text{MAT}), \text{HD}(P)))(x)$
   $\wedge \exists u(\text{mat}(u) \wedge \text{made-of}(u,x)))$

For instance, applying the adjective *paper* to *airplane* converts the type from simply AIRPLANE to an object of the type MADE-OF(PAPER, AIRPLANE). *Paper airplane* would thus yield the following logical form:

1. $\lambda x \lambda \pi \, (\text{airplane}(x, \pi * \text{MADE OF}(\text{PAPER}, \text{AIRPLANE})) \wedge \exists u \, (\text{paper}(u) \wedge$
   $\text{made-of}(u,x,\pi)))$

The object is made essentially out of MAT (Kripke)

Making the value of the polymorphic type more precise

> Use the type hierarchy

- STONE $\sqsubseteq$ MAT(JAR)
- EARTHENWARE $\sqsubseteq$ MAT(JAR) ...

Type constraint for MADE-OF:

a MADE-OF$(\alpha, \beta) \rightarrow (\alpha \sqsubseteq$ MAT$(\beta) \leftrightarrow$ MADE-OF$(\alpha, \beta) \sqsubseteq \beta)$

We can properly predicate *jar* of an object if in fact its polymorphic type is consistent with its being a jar.

b $\lambda x \lambda \pi$ (jar$(x, \pi *$ MADE-OF(STONE, JAR)) $\wedge \exists u$ (stone$(u, \pi) \wedge$
   jar$(u, \pi) \wedge$ made-of(u, x, π))))

Loose talk

- What happens when the polymorphic type is not a subtype of the noun's type (Borschev and Partee 2004)..
  1. (Pointing to a shape that a child has drawn) You've drawn a circle.
  2. stone lion.
- We call things circles or circular when they only approximately resemble mathematical circles.

Loose interpretation with respect to a set of alternatives

the object drawn has a shape that is closer to that of a mathematical circle than any of the relevant alternatives—simple geometric shapes like that of a triangle or a square. *Mutatis mutandis* for stone lions

alternatives

Which set of alternatives is at issue depends on the predicate that is to be interpreted loosely. (a matter of the internal semantics of the predicate, not its external denotation).

We look to the lowest proper supertype in the type hierarchy to find the relevant alternatives. For LION suppose it's *animal*

The alternatives are given by the other types just under this supertype—-LIONS, GIRAFFES, ELEPHANTS, and so on.

Metrics

- The metric for similarity/closeness depends on features associated with the predicate that make up its internal semantics.
- In the general case, it is superficial criteria, rather than the actual extension of the predicate, that define the metric.
- I can judge whether something's a stone lion, even though I have no idea really what the species identifying criterion for lions are.

Features associated with types:

- generalizations using $\sqsubseteq$. LION $\sqsubseteq$ ANIMAL.
- Type specification logic also contains $>$ for determining underspecified types. $>$ encodes generic truths. These also specify features. *male lions have manes* $\mapsto$ LION $>$ HAS-MANE, *giraffes have long necks* GIRAFFE $>$ HAS-LONG-NECK.

Using alternatives to give the semantics of flat tires and flat beer

- Material modifiers affect the type of the head noun. With standard adjectives, head nouns affect the type of the adjective (and its denotation):

  1. $\lambda x \lambda \pi \ (\text{flat}(x, \pi * \text{Applies-to}(\text{BEER}))) \wedge \text{beer}(x, \pi * \text{P}))$

- *flat beer* vs. *flat tire*. *flat beer* denotes beer that is flat compared with the other alternatives (bubbly). *Flat tire* denotes tires that are flat with respect to the alternatives (fully inflated and round).

- Or perhaps these really are coercions as compared to flat water, flat country, flat surface.?

- Similar treatment for adverbial modifiers:

  1. paint a miniature with a brush
  2. scrub the floor with a brush

General picture

- Types handle interactions between word meaning and the predicational context (and discourse context)
- Simple lexical entries for words, enrichments come from the type system.
- Types affect meaning composition in different ways.
- two level semantics (Asher 2011)
  - an internal one for types and the construction of logical form .
  - an external one for intensions, truth and standard entailment.

Why external semantics

- types form an "internal" semantics, information that's used in meaning assembly.
- External semantics: what many sentences are about, the external world.
- Why do internal/external semantics come apart? Because the world isn't always the way we think it is.
- Contextually sensitive expressions are anchored to a particular real world situation. Kripke's Pierre puzzle, Putnam Kripke thought experiments on natural kind terms

Why external content

- Denotations are the content of indexicals, demonstratives, natural kinds, proper names and other so called directly referential expressions.
- E.g., *you* is associated with rules for determining who the audience is in a particular context. But that's not the contribution of *you* to the content of a clause in which it occurs. Its semantic content is the audience itself.
- The behavior of directly referential terms in modal contexts provides compelling evidence that their meaning is not in general determined by "what is in the head" of a competent speaker.
- External content links to truth and a time tested notion of inference, makes testable predictions about inferential relations between sentences.

Why external semantics is not enough: the subtyping problem

- A rich system of types causes problems for a set theoretic model of types.
- Consider first order properties and first order physical properties.
- e.g., If $P \sqsubset E$, then $P \to T$ and $E \to T$ have no common inhabitants.
- types are concepts not identified with sets of their inhabitants but rather something like proof objects, proof theoretic sub typing relation.

Why internal semantics?

- proof theoretic approach gives a natural analysis of sub typing.
- construction of logical form is an internal matter, speaker competence.
- analytic entailments.
- Internal semantics (types) not part of truth conditional at-issue content.
- types are concepts not identified with sets of their inhabitants but rather rules of application, can explore the abstract structure of types category theoretically or using modern type theory (Asher 2011, Luo 2011, 2012).

Questions at the frontier of formal lexical semantics and computational linguistics

- what types are there? (each semantic predicate has a type, and perhaps each lemmatized word, as well as more general types) (Lewis & Steedman 2012)
- Formal semantics and data about semantic anomalies gives us some hints about general types, but not a detailed picture.
- Distributional methods yield similarity classes. Very interesting project.
- Simple clustering techniques validate type distinctions between AGENT, I and P (van der Cruys 2010), or between MASS and COUNT, telic vs. atelic eventualities (Abrusan & van der Cruys 2013).
- work on subtyping also promising in the distributional paradigm. Gives us some analytic entailments.
    1. A dog is in the garden. → An animal is in the garden.

The content problem

- What is the content of types, the content of internal semantics?
- that "tracks" external content (Asher 2011).

Internal semantics in mathematics

- '2' is a name of an object that can be constructed from the instructions in the numeral—i.e., $2 = S(S(0))$, 0 is an object defined axiomatically in N.
- the type N is inductively defined from the primitive object 0 and successor. To say that $x\colon N$ is to say that x can be proven to be 0 or some successor of 0.
- $\lambda x\colon N\ \mathrm{Prime}(x)$ is a property that for any object x of type N returns a proof that x is divisible only by itself and 1 or a contradiction. (division is defined axiomatically in N).
- $\|\mathrm{Prime}(2)\|$ is the proof that the object that can be constructed from the instructions in the name '2′' has the property of being prime

Internal content for non mathematical expressions?

- Standard formal semantics can provide an internal semantics for closed class words (proof theory for connectives and quantifiers, temporal structure in FO($<, \circ$), and perhaps discourse structure, modality. For the full class of determiners, the proof theory is incomplete.

- For open class terms, formal semantics has little to offer. Lexical semantics shows how to use types to get some inferences (type hierarchy, type disarmbiguation).

- A bigger problem: sentences in non mathematics don't convey proofs (actually there's also a problem for mathematical statements that are false).

Getting content through distributional methods

- Similarity classes over what kind of contexts?
- Similarity over syntactic environments is not right on its own to give the internal content of types. Antonyms are quite similar distributionally,
- yet they don't have the same internal content (e.g., *claim* vs. ₌ deny)
- Can similarity yield a good notion of logical consequence on its own?
    - ?? $\phi \sim \psi \rightarrow (\phi \rightarrow \psi)$. NO
    - ?? $\phi \sim \psi \rightarrow (\phi > \psi)$ NO.
    - For individual words $w\ w'$, $(w \sim w' \wedge X(w): t) \rightarrow (X(w) > X(w'))$ (still doubtful when we look at similarity classes provided by distributional semantics over syntactic environments)
- Internalist conceptions of meaning have a problem of infinite regress/ and or circularity.
- Formal semantics and distributional semantics together might provide a good answer, but how?

Prototypical features

- Frequencies of certain co-occurring expressions in restricted contexts, constructions apt to give us prototypical features, might give us a better handle on meaning than just syntactic distributions.
- In 300 dimensions coming from an nmf model (van der Cruysm pc), with the most salient nouns and the most salient dependencies for those nouns, for each dimension, we often see prototypical features of nouns.
- how do we use prototypical features?

From generics to justification

- Generics: $\forall x(\phi(x) > \psi(x))$
- normal elements of $\|\phi\|$ normally have the property $\|\psi\|$
- such ψ properties are often used as evidence for something's being a ϕ
- the cluster of such properties is used to build *a justification* for ϕ.
- We associate the cluster of such properties with the internal meaning associated with ϕ.
- true generics anchor the internalist semantics to the external one.

Meanings and compositionality

- Meanings are justifications.
- $\|N\|$: functions from inhabitants of any subtype of E into a justification that the individual is of the type N.
- Adjectives $\|N\| \to \|N\| \times \|A_N\|$.
- Det: $\|N\| \to \|N\| \to$ PROP
- TVs: $\|DP\| \to |DP\| \to \|DP_x DP_y V(x,y)\| \sqsubseteq$ PROP
- PROP: a collection of justifications for the truth of propositions/thoughts.

An example

- $\|\text{BOY}\|$: a function that given any individual either provides a justification that the object has the characteristics associated with *boy* or returns $\bot$.
- $\|\text{RUN}\|$: a function that given any individual either provides a justification that the object has the characteristics associated with *run* or returns $\bot$.
- $\|(\text{EVERYBOY})\text{RUNS}\|$ provides the following justification given any object x for which there is a justification of its being a boy, there is a justification of x's running.
- BOYS LIKE CATS a justification that given any object x for which there is a justification of its being a (typical) boy, and for any object y such that there is a justification of its being a cat, there is a justification that x likes y.

Conclusions

- use the best of formal semantics and distributional semantics
- try to capitalize on a subset of contexts for meanings of words, those contexts that serve to justify the use of that word rather than others.