Natural logic and natural language semantics

Chris Potts, Ling 236/Psych 236c: Representations of meaning, Spring 2013

Apr 11

1 Background

van Benthem (2008) on the origins:

> In the 1980s, the idea arose that these observations had a more general thrust, namely, that natural language is not just a medium for saying and communicating things, but that it also has a 'natural logic', viz. a system of modules for ubiquitous forms of reasoning that can operate directly on natural language surface form. This idea was developed in some detail in van Benthem 1986, 1987. The main proposed ingredients were two general modules:
>
> (a) Monotonicity Reasoning, i.e., Predicate Replacement,
>
> (b) Conservativity, i.e., Predicate Restriction, and also
>
> (c) Algebraic Laws for inferential features of specific lexical items.

Lakoff (1973:159):

> CONCLUSION 1: The rules of grammar, which generate the grammatical sentences of English, filtering out the ungrammatical sentences, are not distinct from the rules relating the surface forms of English sentences to their corresponding logical forms.
>
> [...]
>
> At present, the theory of generative semantics is the only theory of grammar that has been proposed that is consistent with conclusion 1.

Dowty (1994) on inference and psychology:

> humans do not carry anything analogous to infinite sets of possible worlds or situations around in their heads, so the study of deduction — inferential relations based on syntactic properties of some kind of "representations" of denotations — are [sic] potentially of relevance to the psychology of language and to computational processing of meaning in a way that model-theoretic semantics alone is not.

Partee (1988:117) on possible worlds and psychology (see also Partee 1980):

> In earlier writings I had raised some worries about the possible psychological reality of possible worlds which in retrospect I think arose partly from misguided concern with trying to fit possible-worlds theory "inside the speaker's head": but we don't try to fit the theory of gravitation "inisde a falling apple". Nevertheless, linguists *do* care about the mechanisms, and if possible-worlds semantics is a reasonable theory *about* the language user's semantic competence, there still have to exist some internal mechansisms partly by virtue of which the external theory is correct.

MacCartney & Manning (2009:140–141):

> For a semanticist, the most obvious approach to NLI relies on full semantic interpretation: first, translate *p* and *h* into some formal meaning representation, such as first-order logic (FOL), and then apply automated reasoning tools to determine inferential validity. While the formal approach can succeed in restricted domains, it struggles with open-domain NLI tasks such as RTE. For example, the FOL-based system of (1) [=Bos & Markert 2005 —CP] was able to find a proof for less than 4% of the problems in the RTE1 test set. The difficulty is plain: truly natural language is fiendishly complex. The formal approach faces countless thorny problems: idioms, ellipsis, paraphrase, ambiguity, vagueness, lexical semantics, the impact of pragmatics, and so on. Consider for a moment the difficulty of fully and accurately translating (1) to a formal meaning representation.
>
> Yet (1) also demonstrates that full semantic interpretation is often not necessary to determining inferential validity.

See also Sánchez-Valencia 1991; Moss 2009, 2008; MacCartney 2009; MacCartney & Manning 2007, 2008; Schubert et al. 2010; Icard 2012

2 Comparisons with possible-worlds models

Properties in common

- Commitment to compositionality
- A rich lexicon and a small stock of composition rules
- Basis in higher-order logic
- Aspirations to direct interpretation of surface forms

Differences

- Emphasis on
  - Possible worlds: Model theory
  - Natural logic: Deduction
- The lexicon
  - Possible worlds: Meanings as truth functions are primary
  - Natural logic: Relations between meanings are primary
- Composition
  - Possible worlds: emphasis on function application between lexical meanings
  - Natural logic: emphasis on function composition between lexical relationships

3 Generalized entailment

Definition 1 (Generalized entailment for meanings)**.** For all domains $\mathbf{D}$ and meanings $a, b \in \mathbf{D}$:

i. If $a, b \in D_e$, then $a \sqsubseteq b$ iff $a = b$
ii. If $a, b \in D_s$, then $a \sqsubseteq b$ iff $a = b$
iii. If $a, b \in D_t$, then $a \sqsubseteq b$ iff $a = \mathtt{F}$ or $b = \mathtt{T}$
iv. If $a, b \in D_{\langle\sigma,\tau\rangle}$, $a \sqsubseteq b$ iff for all $d \in D_\sigma$, $a(d) \sqsubseteq b(d)$

Definition 2 (Generalized entailment for forms)**.** For all models $\mathbf{M}$, assignments g, and expressions α and β, $\alpha \Rightarrow \beta$ iff $[\![\alpha]\!]^{\mathbf{M},g} \sqsubseteq [\![\beta]\!]^{\mathbf{M},g}$.

Definition 3 (Non-entailment, $\not\Rightarrow$)**.** If it's false that $\alpha \Rightarrow \beta$, then $\alpha \not\Rightarrow \beta$

Definition 4 (Synonymy, $\equiv$)**.** If $\alpha \Rightarrow \beta$ and $\beta \Rightarrow \alpha$, then $\alpha \equiv \beta$.

4 Monotonicity

Definition 5 (Upward monotonicity)**.** For $f \in D_{\langle\sigma,\tau\rangle}$, $\Uparrow f$ iff for all $a, b \in D_\sigma$, if $a \sqsubseteq b$, then $f(a) \sqsubseteq f(b)$.

Definition 6 (Downard monotonicity)**.** For $f \in D_{\langle\sigma,\tau\rangle}$, $\Downarrow f$ iff for all $a, b \in D_\sigma$, if $a \sqsubseteq b$, then $f(b) \sqsubseteq f(a)$.

Definition 7 (Non-monotone)**.** For $f \in D_{\langle\sigma,\tau\rangle}$, $\not\Updownarrow f$ iff neither $\Uparrow f$ nor $\Downarrow f$.

For a richer set of monotone functions, see Hoeksema 1983; Ladusaw 1996; van der Wouden 1997; Zwarts 1998; Icard 2012.

(1) Negation is $\Downarrow$
dog $\Rightarrow$ **animal**
(**not animal**) $\Rightarrow$ (**not dog**)

(2) The implicative verb *fail (to)* is $\Downarrow$:
tango $\Rightarrow$ **dance**
(**fail dance**) $\Rightarrow$ (**fail tango**)

(3) *manage* is $\Uparrow$ (ignoring its presupposition!)
(**manage tango**) $\Rightarrow$ (**manage dance**)

(4) Determiner monotonicity

Det	Restriction	Scope
some	$\Uparrow$	$\Uparrow$
no	$\Downarrow$	$\Downarrow$
every	$\Downarrow$	$\Uparrow$
not every	$\Uparrow$	$\Downarrow$
exactly N	$\not\Updownarrow$	$\not\Updownarrow$
most	$\not\Updownarrow$	$\Uparrow$
a minority of	$\not\Updownarrow$	$\Downarrow$
?	$\Uparrow$	$\not\Updownarrow$
	$\Downarrow$	$\not\Updownarrow$

5 Drop the models?

Suppose we ignored def. 1, taking $\Rightarrow$ to be a primitive relation on forms and accepting def. 3 and def. 4. We could similarly define $\Uparrow$ and $\Downarrow$ as classifications of forms, stipulating that they are exclusive and define $\not\Updownarrow$ in terms of them.

6 Semantic relations

MacCartney's natural logic extends the deductive notion of monotonicity to a wider range of set-theoretic relations, all of which can be applied cross-categorically as in definitions 1–7.

symbol: $x \equiv y$
name: equivalence
example: **couch** $\equiv$ **sofa**
set-theoretically: $x = y$

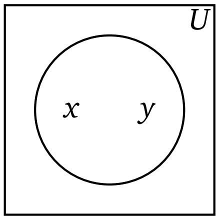

symbol: $x \sqsubset y$
name: forward entailment
example: **crow** $\sqsubset$ **bird**
set-theoretically: $x \subset y$

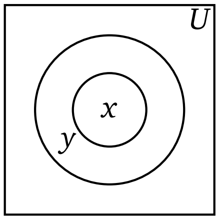

symbol: $x \sqsupset y$
name: reverse entailment
example: **bird** $\sqsupset$ **crow**
set-theoretically: $x \supset y$

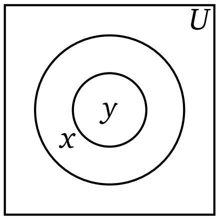

symbol: $x \,^\wedge\, y$
name: negation
example: **human** $^\wedge$ **nonhuman**
set-theoretically: $x \cap y = \emptyset \wedge x \cup y = U$

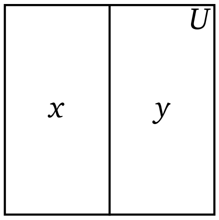

symbol: $x \mid y$
name: alternation
example: **dog** $\mid$ **desk**
set-theoretically: $x \cap y = \emptyset \wedge x \cup y \neq U$

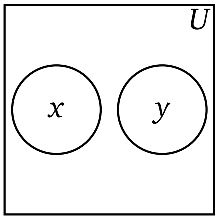

symbol: $x \smile y$
name: cover
example: **animal** $\smile$ **nonmammal**
set-theoretically: $x \cap y \neq \emptyset \wedge x \cup y = U$

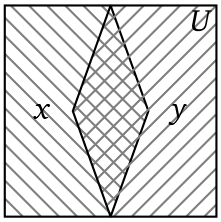

symbol: $x \# y$
name: independence
example: **hungry** # **hippo**
set-theoretically: (all other cases)

Icard (2012) redefines these using the basic set $\{\sqsubseteq, \sqsupseteq, |, \smile\}$, where $x \equiv y$ iff $x \sqsubseteq y$ and $x \sqsupseteq y$, $\smile$ is defined as $x \cup y = U$ and $x \,^\wedge\, y$ iff $x \mid y$ and $x \smile y$.

7 Joining semantic relations via relational composition

Definition 8 (Joins, $\bowtie$).

$$R \bowtie S \overset{def}{=} \{\langle x,z\rangle : \exists y(\langle x,y\rangle \in R \wedge \langle y,z\rangle \in S)\}$$

(5) For all R, $(R \bowtie \equiv) = (\equiv \bowtie R) = R$

(Begin with a set x, and suppose xRy and $y \equiv z$. By the definition of $\equiv$, $y = z$. So we still have R.)

(6) $\hat{} \bowtie \hat{} = \equiv$

(Begin with a set x. There is a unique y such that $x \,\hat{}\, y$, namely, the complement of x. The complement of y is also unique: it is x. Hence $x \,\hat{}\, y$ and $y \,\hat{}\, z$ implies $x = z$.)

(7) $\sqsubset \bowtie \sqsubset = \sqsubset$ (ditto for $\sqsupset$)

(Suppose $x \sqsubset y$ and $y \sqsubset z$. By the transitivity of the subset relation, it holds that $x \sqsubset z$.)

(8) $| \bowtie \hat{} = \sqsubset$

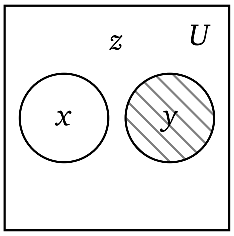

(9) $\hat{} \bowtie | = \sqsupset$

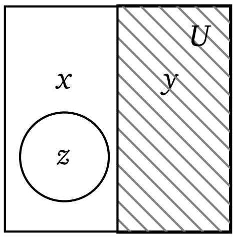

(10) $\sqsubset \bowtie \sqsupset = \#$

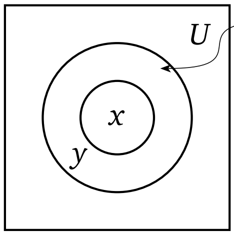

z could be anywhere inside y — disjoint from x, overlapping with x, subsuming or subsumed by x

The following table approximates the set of pairwise join relations. The complete table is given in MacCartney & Manning 2009 (and see their footnote 12).

$\bowtie$	$\equiv$	$\sqsubset$	$\sqsupset$	$\hat{}$	$\|$	$\smile$	#
$\equiv$	$\equiv$	$\sqsubset$	$\sqsupset$	$\hat{}$	$\|$	$\smile$	#
$\sqsubset$	$\sqsubset$	$\sqsubset$	#	$\|$	$\|$	#	#
$\sqsupset$	$\sqsupset$	#	$\sqsupset$	$\smile$	#	$\smile$	#
$\hat{}$	$\hat{}$	$\smile$	$\|$	$\equiv$	$\sqsupset$	$\sqsubset$	#
$\|$	$\|$	#	$\|$	$\sqsubset$	#	$\sqsubset$	#
$\smile$	$\smile$	$\smile$	#	$\sqsupset$	$\sqsupset$	#	#
#	#	#	#	#	#	#	#

Limitations The calculus does not validate de Morgan's laws (its inferences are not strong enough; see MacCartney 2009:§6.5.4 and (17)–(18) below). In addition, because # is a point of no return, some edit orders lead to weaker conclusions than others (MacCartney 2009:§5.6; Icard 2012:§3.3).

8 Important relations

Entailment : $\{\equiv, \sqsubset\}$ **Contradiction** : $\{^\wedge, |\}$ **Compatibility** : $\{\sqsupset, \smile, \#\}$

9 Semantic composition in tabular format

From Bill MacCartney's CS224U teaching materials:[1]

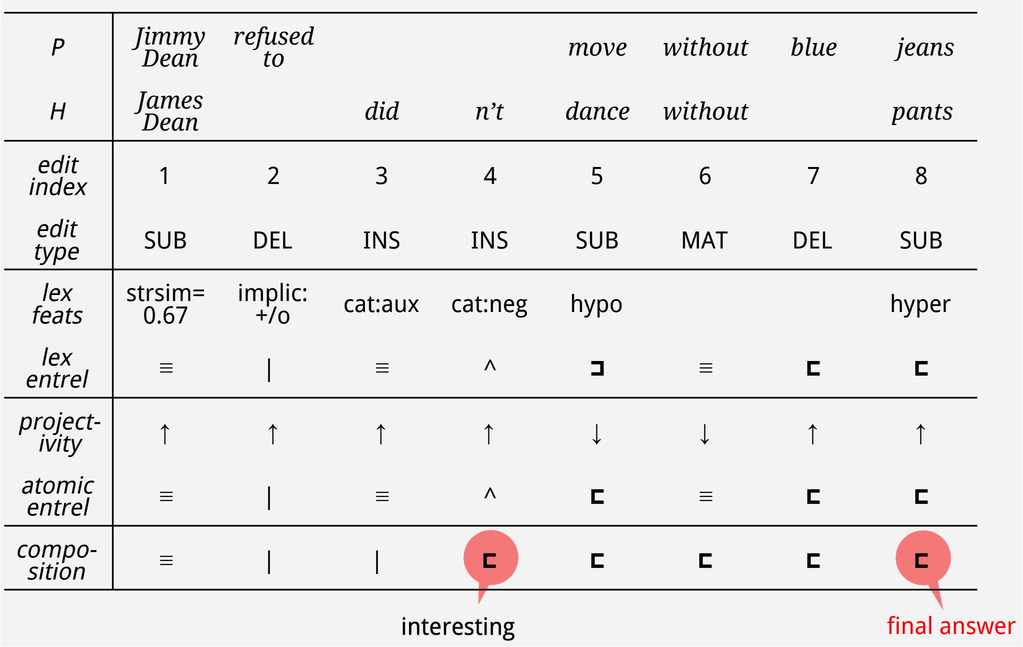

P	*Jimmy Dean*	*refused to*			*move*	*without*	*blue*	*jeans*
H	*James Dean*		*did*	*n't*	*dance*	*without*		*pants*
edit index	1	2	3	4	5	6	7	8
edit type	SUB	DEL	INS	INS	SUB	MAT	DEL	SUB
lex feats	strsim= 0.67	implic: +/o	cat:aux	cat:neg	hypo			hyper
lex entrel	≡	|	≡	^	⊐	≡	⊏	⊏
project-ivity	↑	↑	↑	↑	↓	↓	↑	↑
atomic entrel	≡	|	≡	^	⊏	≡	⊏	⊏
compo-sition	≡	|	|	⊏	⊏	⊏	⊏	⊏

- The NLU system determines the 'lex entrel' values based on a variety of methods.
- The 'projectivity' line specifies one of two functions, which map the 'lex entrel' values to the 'atomic enrel' values:

$$\uparrow = \text{identity function} \qquad \downarrow = \begin{bmatrix} \equiv & \longmapsto & \equiv \\ \# & \longmapsto & \# \\ ^\wedge & \longmapsto & ^\wedge \\ \sqsupset & \longleftrightarrow & \sqsubset \\ | & \longleftrightarrow & \smile \end{bmatrix}$$

- The 'composition' line joins its value with the 'atomic entrel' relation to its northeast, to determine the relation to its immediate right.

For more discussion, see MacCartney 2009:§7, especially the notes in §7.3 about the importance of alignment order (see also Icard 2012:§3.3) and the extended example in §7.7

---

[1] http://www.stanford.edu/class/cs224u/slides/cs224u-2013-lec12-nli.pdf

10 Semantic composition on aligned trees

10.1 Lexicon

Lex is defined for expressions of the same semantic type, plus the empty string ε, which could denote a type-polymorphic identity function.

$$\begin{array}{rcl@{\qquad}rcl}
\textsc{Lex}(\textbf{poodle}, \varepsilon) & = & \sqsubset & \textsc{Lex}(\varepsilon, \textbf{poodle}) & = & \sqsubset \\
\textsc{Lex}(\textbf{not}, \varepsilon) & = & \wedge & \textsc{Lex}(\varepsilon, \textbf{not}) & = & \wedge \\
\textsc{Lex}(\textbf{elegantly}, \varepsilon) & = & \sqsubset & \textsc{Lex}(\varepsilon, \textbf{elegantly}) & = & \sqsupset \\
\textsc{Lex}(\textbf{poodle}, \textbf{dog}) & = & \sqsubset & \textsc{Lex}(\textbf{dog}, \textbf{poodle}) & = & \sqsupset \\
\textsc{Lex}(\textbf{dog}, \textbf{mammal}) & = & \sqsubset & \textsc{Lex}(\textbf{mammal}, \textbf{dog}) & = & \sqsupset \\
\textsc{Lex}(\textbf{tango}, \textbf{dance}) & = & \sqsubset & \textsc{Lex}(\textbf{dance}, \textbf{move}) & = & \sqsubset \\
\textsc{Lex}(\textbf{cat}, \textbf{dog}) & = & \mid & \textsc{Lex}(\textbf{cat}, \textbf{noncat}) & = & \wedge \\
\textsc{Lex}(\textbf{great}, \textbf{good}) & = & \sqsubset & \textsc{Lex}(\textbf{awful}, \textbf{bad}) & = & \sqsubset \\
\textsc{Lex}(\textbf{sofa}, \textbf{couch}) & = & \equiv & \textsc{Lex}(\textbf{al}, \textbf{albert}) & = & \equiv \\
\textsc{Lex}(\textbf{every}, \textbf{some}) & = & \sqsubset & \textsc{Lex}(\textbf{every}, \textbf{no}) & = & \mid
\end{array}$$

$$\textsc{Lex}(\textbf{not}, \textbf{not}) = \textsc{Lex}(\textbf{not}, \textbf{never}) = \begin{bmatrix}
\equiv & \longmapsto & \equiv \\
\# & \longmapsto & \# \\
\wedge & \longmapsto & \wedge \\
\sqsupset & \longleftrightarrow & \sqsubset \\
\mid & \longleftrightarrow & \smile
\end{bmatrix}$$

$$\textsc{Lex}(\textbf{every}, \textbf{every}) = \begin{array}{cc}
\textbf{restriction} & \textbf{nuclear scope} \\
\begin{bmatrix}
\equiv & \longmapsto & \equiv \\
\# & \longmapsto & \# \\
\wedge & \longmapsto & \mid \\
\mid & \longmapsto & \# \\
\smile & \longmapsto & \mid \\
\sqsupset & \longleftrightarrow & \sqsubset \\
\sqsubset & \longleftrightarrow & \sqsupset
\end{bmatrix} &
\begin{bmatrix}
\equiv & \longmapsto & \equiv \\
\# & \longmapsto & \# \\
\wedge & \longmapsto & \mid \\
\mid & \longmapsto & \mid \\
\smile & \longmapsto & \# \\
\sqsupset & \longleftrightarrow & \sqsupset \\
\sqsubset & \longleftrightarrow & \sqsubset
\end{bmatrix}
\end{array}$$

$$\textsc{Lex}(\textbf{and}, \textbf{and}) = \begin{array}{cc}
\textbf{left conjunct} & \textbf{right conjunct} \\
\begin{bmatrix}
\equiv & \longmapsto & \equiv \\
\# & \longmapsto & \# \\
\sqsubset & \longmapsto & \sqsubset \\
\sqsupset & \longmapsto & \sqsupset \\
\wedge & \longmapsto & \mid \\
\mid & \longmapsto & \mid \\
\smile & \longmapsto & \#
\end{bmatrix} &
\begin{bmatrix}
\equiv & \longmapsto & \equiv \\
\# & \longmapsto & \# \\
\sqsubset & \longmapsto & \sqsubset \\
\sqsupset & \longmapsto & \sqsupset \\
\wedge & \longmapsto & \mid \\
\mid & \longmapsto & \mid \\
\smile & \longmapsto & \#
\end{bmatrix}
\end{array}$$

10.2 Composition

Definition 9 (Atomic joins).

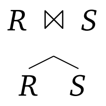

Definition 10 (One-place projection).

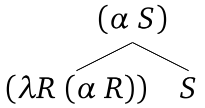

Definition 11 (Two-place projection).

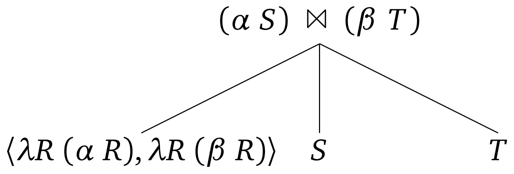

10.3 Examples

(11) *Al didn't move* **entails** *Albert didn't dance*

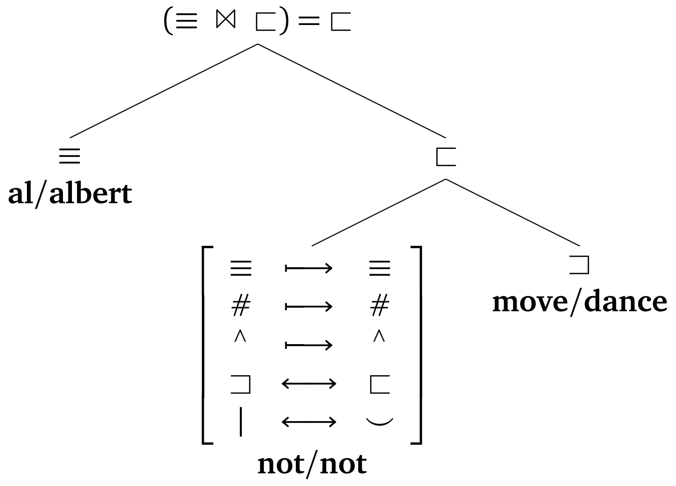

(12) *Albert didn't move* **contradicts** *Albert danced elegantly*

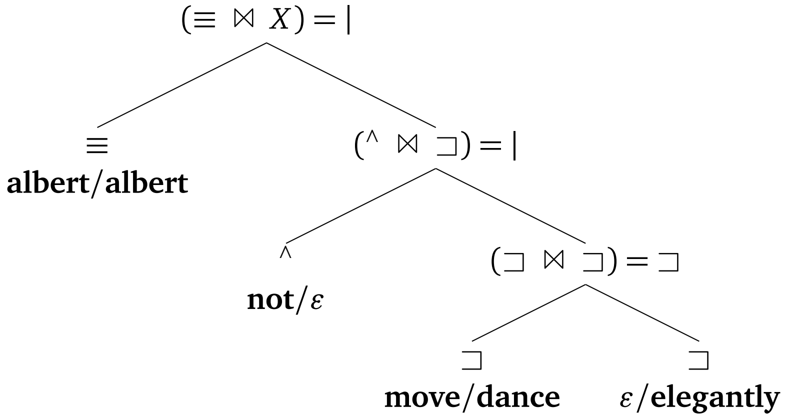

(13) *Albert is not a cat* **is entailed by** *Albert is a sofa*

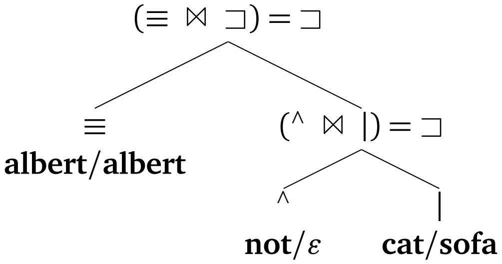

(14) *No dog moved* **contradicts** *Some poodle danced*

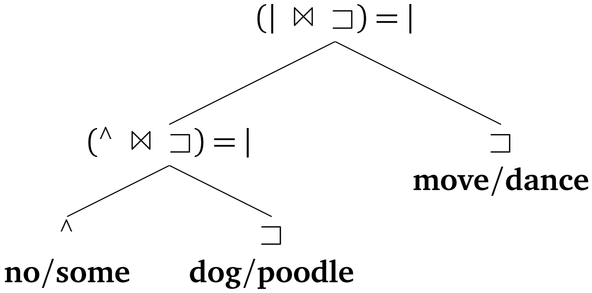

(15) *Every dog danced* **entails** *Every poodle moved*

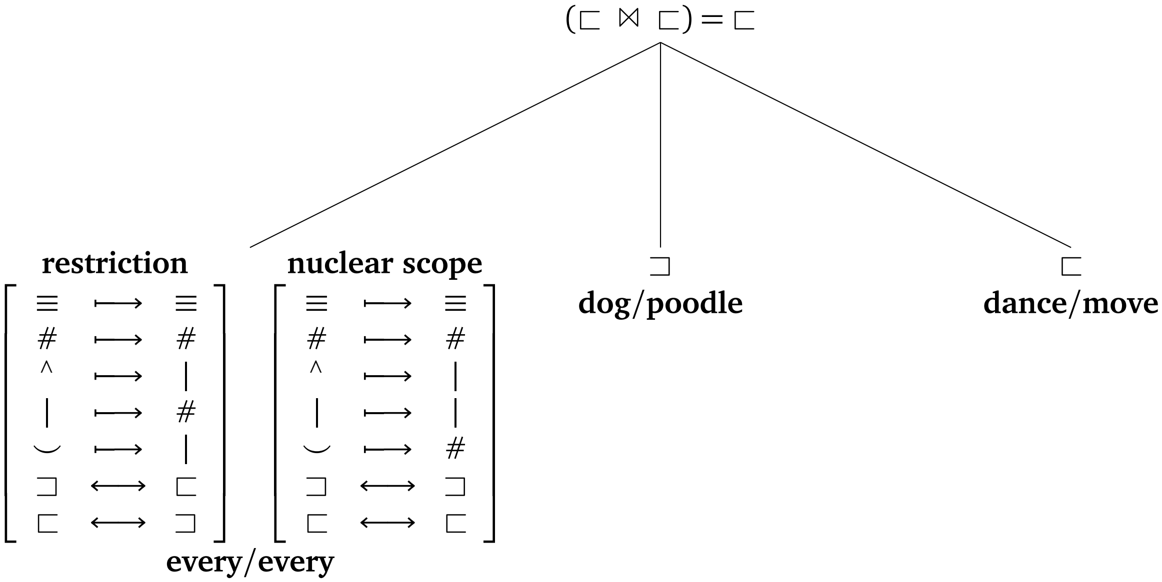

(16) *Every dog moved* **does not entail** *Every poodle danced*

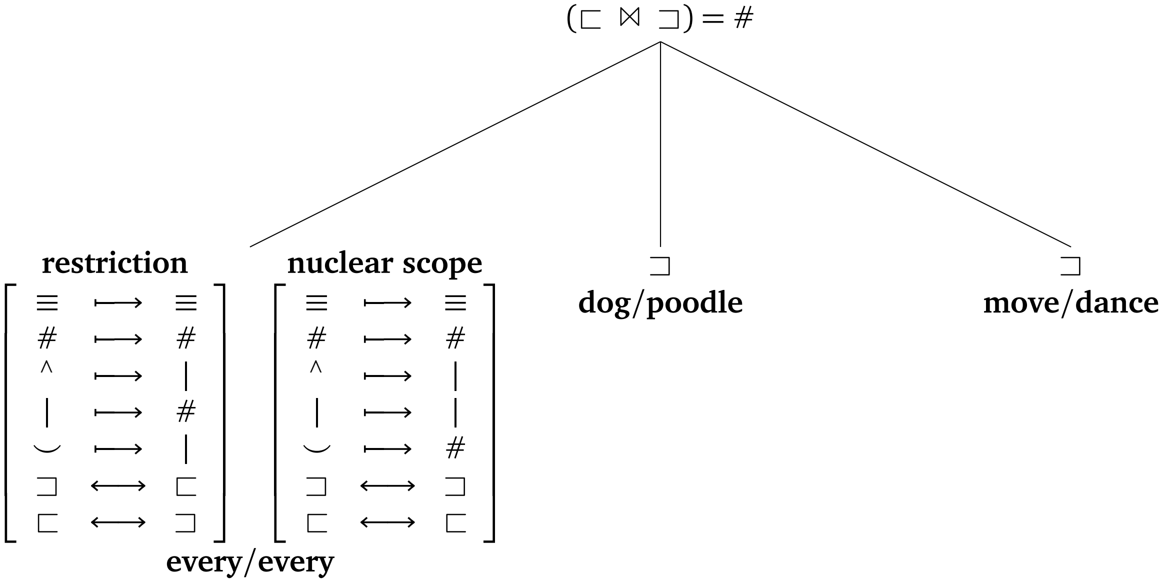

(17) *Some dog doesn't dance* **entails** *Not every dog dances*

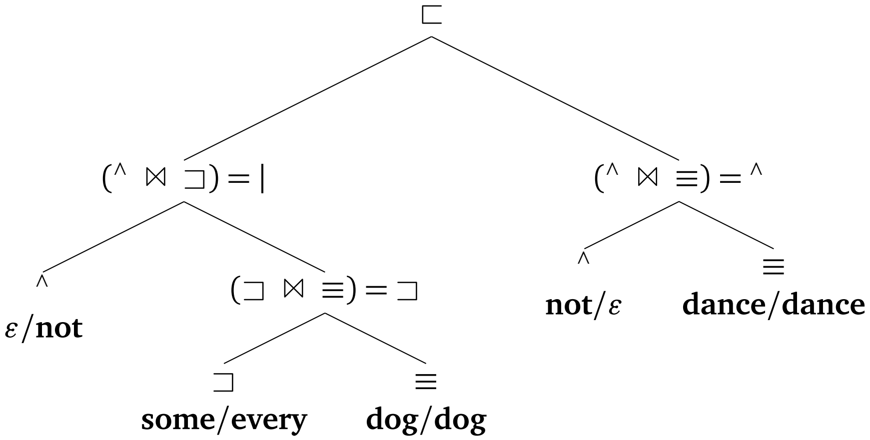

(18) *Not every dog dances* **does not entail** *Some dog doesn't dance* ☹

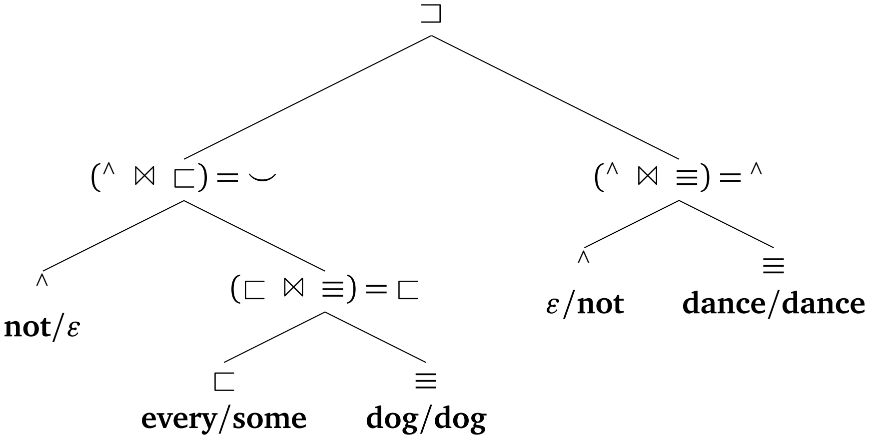

Note: if we reversed the order in def. 9, then we would get this right but get (17) wrong.

11 MacCartney's experiments

The NatLog system

i. Lexical information from WordNet, string similarity measures, etc.
ii. Stanford parser (Klein & Manning 2003)
iii. Projectivity via Tregex patterns (Levy & Andrew 2006)
iv. Alignment between premise and hypothesis.
v. Semantic composition in (basically) the manner described here.

FraCaS 346 inference problems, devised for testing inference engines. The examples tend to be slightly more complex than what one finds in typical linguistics papers but simpler than one normally finds in corpora. Bill's release: `http://www-nlp.stanford.edu/~wcmac/downloads/`

§	ID	Premise	Hypothesis	Ans
1	38	No delegate finished the report.	Some delegate finished the report on time.	*no*
1	48	At most ten commissioners spend time at home.	At most ten c...s spend a lot of time at home.	*yes*
2	83	Either Smith, Jones or Anderson signed the contract.	Jones signed the contract.	*unk*
5	205	Dumbo is a large animal.	Dumbo is a small animal.	*no*
6	233	ITEL won more orders than APCOM.	ITEL won some orders.	*yes*
9	335	Smith believed that ITEL had won the contract in 1992.	ITEL won the contract in 1992.	*unk*

Table 2: Illustrative examples from the FraCaS test suite

Recognizing Textual Entailment (RTE) More naturalistic than FraCaS. The inferences depend not only on logical and lexical properties, but also on world knowledge and complex default assumptions about evidentiality (Zaenen et al. 2005; Manning 2006; Crouch et al. 2006).

ID	Premise	Hypothesis	Answer
71	As leaders gather in Argentina ahead of this weekends regional talks, Hugo Chávez, Venezuela's populist president is using an energy windfall to win friends and promote his vision of 21st-century socialism.	Hugo Chávez acts as Venezuela's president.	*yes*
788	Democrat members of the Ways and Means Committee, where tax bills are written and advanced, do not have strong small business voting records.	Democrat members had strong small business voting records.	*no*

Table 5: Illustrative examples from the RTE3 development set

Results (MacCartney & Manning 2008)

§	Section	#	P %	R %	Acc %
1	Quantifiers	44	95.24	100.00	**97.73**
2	Plurals	24	90.00	64.29	75.00
3	Anaphora	6	100.00	60.00	50.00
4	Ellipsis	25	100.00	5.26	24.00
5	Adjectives	15	71.43	83.33	80.00
6	Comparatives	16	88.89	88.89	81.25
7	Temporal	36	85.71	70.59	58.33
8	Verbs	8	80.00	66.67	62.50
9	Attitudes	9	100.00	83.33	88.89
1, 2, 5, 6, 9		108	90.38	85.45	**87.04**

Table 3: Performance on FraCaS problems (three-way classification). The columns show the number of problems, precision and recall for the *yes* class, and accuracy. Results for NatLog are broken out by section.

System	Data	% Yes	P %	R %	Acc %
Stanford	dev	50.25	68.66	66.99	67.25
	test	50.00	61.75	60.24	60.50
NatLog	dev	22.50	**73.89**	32.38	59.25
	test	26.38	**70.14**	36.10	59.38
Hybrid, bal.	dev	50.00	70.25	68.20	68.75
	test	50.00	65.50	63.90	64.25
Hybrid, opt.	dev	56.00	69.20	75.24	**70.00**
	test	54.50	64.45	68.54	**64.50**

Table 6: Performance of various systems on RTE3 (two-way classification). The columns show the data set used (800 problems each), the proportion of *yes* predictions, precision and recall for the *yes* class, and accuracy.

References

van Benthem, Johan. 1986. *Essays in logical semantics*. Dodrecht: D. Reidel.

van Benthem, Johan. 1987. Meaning: Interpretation and inference. *Synthese* 73(3). 451–470.

van Benthem, Johan. 2008. A brief history of natural logic. In M Chakraborty, B Löwe, M. Nath Mitra & S. Sarukki (eds.), *Logic, navya-nyaya and applications: Homage to Bimal Matilal*, .

Bos, Johan & Katja Markert. 2005. Recognising textual entailment with logical inference. In *Proceedings of human language technology conference and conference on empirical methods in natural language processing*, 628–635. Vancouver: ACL.

Crouch, Richard, Lauri Karttunen & Annie Zaenen. 2006. Circumscribing is not excluding: A reply to Manning. Ms., Palo Alto Research Center.

Dowty, David. 1994. The role of negative polarity and concord marking in natural language reasoning. In Mandy Harvey & Lynn Santelmann (eds.), *Proceedings of Semantics and Linguistic Theory 4*, 114–144. Ithaca, NY: CLC Publications.

Hoeksema, Jack. 1983. Negative polarity and the comparative. *Natural Language and Linguistic Theory* 1. 403–434.

Icard, Thomas F. 2012. Inclusion and exclusion in natural language. *Studia Logica* 100(4). 705–725.

Klein, Dan & Christopher D. Manning. 2003. Accurate unlexicalized parsing. In *ACL '03: Proceedings of the 41st annual meeting of the Association for Computational Linguistics*, vol. 1, 423–430. ACL.

Ladusaw, William A. 1996. Configurational expression of negation. In Jaap van der Does & Jan van Eijck (eds.), *Quantifiers, logic, and language*, 203–223. Stanford, CA: CSLI.

Lakoff, George. 1973. Linguistics and natural logic. *Synthese* 22(1–2). 151–271.

Levy, Roger & Galen Andrew. 2006. Tregex and Tsurgeon: Tools for querying and manipulating tree data structures. In *Proceedings of the 5th edition of the International Conference on Language Resources and Evaluation*, 2231–2234.

MacCartney, Bill. 2009. *Natural language inference*: Stanford University dissertation.

MacCartney, Bill & Christopher D. Manning. 2007. Natural logic for textual inference. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, 193–200. Prague: Association for Computational Linguistics. `http://www.aclweb.org/anthology/W/W07/W07-1431`.

MacCartney, Bill & Christopher D. Manning. 2008. Modeling semantic containment and exclusion in natural language inference. In *Proceedings of the 22nd international conference on computational linguistics (coling 2008)*, 521–528. Manchester, UK: Coling 2008 Organizing Committee. `http://www.aclweb.org/anthology/C08-1066`.

MacCartney, Bill & Christopher D. Manning. 2009. An extended model of natural logic. In *Proceedings of the eighth international conference on computational semantics*, 140–156. Tilburg, The Netherlands: Association for Computational Linguistics. `http://www.aclweb.org/anthology/W09-3714`.

Manning, Christopher D. 2006. Local textual inference: It's hard to circumscribe, but you know it when you see it – and NLP needs it. Ms., Stanford University.

Moss, Lawrence S. 2008. Completeness theorems for syllogistic fragments. In Fritz Hamm & Stephan Kepser (eds.), *Logics for linguistic structures*, 143–173. Berlin: Mouton de Gruyter.

Moss, Lawrence S. 2009. Natural logic and semantics. In Maria Aloni, Harald Bastiaanse, Tikitu de Jager, Peter van Ormondt & Katrin Schulz (eds.), *Preproceedings of the 17th Amsterdam Colloquium*, 71–80. University of Amsterdam.

Partee, Barbara H. 1980. Semantics – mathematics or psychology? In Egli Bäuerle & Arnim von Stechow (eds.), *Semantics from different points of view*, 1–14. Berlin: Springer-Verlag.

Partee, Barbara H. 1988. Possible world in model-theoretic semantics: A linguistic perspective. In Sture Allen (ed.), *Possible worlds in humanities, arts, and sciences: Proceedings of Nobel symposium 65*, 93–l23. Berlin and New York: Walter de Gruyter.

Sánchez-Valencia, Víctor. 1991. *Studies in natural logic and categorial grammar*: University of Amsterdam dissertation.

Schubert, Lenhart K., Benjamin Van Durme & Marzieh Bazrafshan. 2010. Entailment inference in a natural logic-like general reasoner. In *Commonsense knowledge: Papers from the AAAI Fall symposium*, 90–95. AAAI.

van der Wouden, Ton. 1997. *Negative contexts: Collocation, polarity and multiple negation*. London and New York: Routledge.

Zaenen, Annie, Lauri Karttunen & Richard Crouch. 2005. Local textual inference: Can it be defined or circumscribed? In *Proceedings of the ACL workshop on empirical modeling of semantic equivalence and entailment*, 31–36. Ann Arbor, MI: Association for Computational Linguistics. `http://www.aclweb.org/anthology/W/W05/W05-1206`.

Zwarts, Frans. 1998. Three types of polarity. In Fritz Hamm & Erhard Hinrichs (eds.), *Plurality and quantification*, 177–238. Dordrecht: Kluwer.