TSINGHUA LECTURES ON LOGIC AND NATURAL LANGUAGE

Edward L. Keenan

April 19, 20 2021

LECTURE 1

A LOGICIAN'S LOOK AT NATURAL LANGUAGE

As a work-a-day linguist I find the application of concepts, techniques and notation of modern logic to be enlightening in the study of the structure of natural language (NL). This lecture focuses on these applications. Lecture 2 focuses on ways the expression of logical operators in NL differ from their sources in logic.

A personal note As a graduate student in linguistics (1965 – 1969) I was interested in semantics, but then there were no systematic studies of semantics within linguistics. I turned to Predicate Logic (PL). It was thriving and was what linguists could call a Universal Grammar: it provided explicit definitions of a class of languages (those of Elementary Arithmetic, Euclidean Geometry, Set Theory,...) and showed us how to semantically interpret them *compositionally*. The rigor and clarity of this work was breathtaking.

And compositionality was linguistically crucial. It treats the semantic interpretation of a derived expression as a function of those it is derived from. This directly addressed the linguistic goal (Chomsky 1957, 1965) of explaining how we interpret novel utterances: we know what their parts mean and how the entire expression takes its interpretation as a function of that of its parts.

We turn now to a few specific linguistic applications of PL. They are noted in arabic numerals, as **1.** below. Bolded headings in bracketed numerals, **[n]**, are speculative observations, sometimes highly speculative.

1. Entailment

Some linguists and philosophers have sought to define THE MEANING OF (φ), φ an expression, treating "meanings" as if they were abstract objects floating in Platonic space that we sneak up on and "capture". This does not suggest specific questions to ask or experiments to run. PL was more fruitful. It asks *relational* questions: Does *φ entail ψ*? That is, is ψ true in all the situations (models) in which φ is true? We <u>can</u> say (1a) **differs in meaning** from (1b) in that it entails (1c), but (1b) does not:

(1) a. Ted can both speak and read Chinese
    b. Ted can either speak or read Chinese
    c. Ted can read Chinese

Since whether a statement entails another depends on their meaning it is clear that sentences (Ss) with different entailments have different meanings even if we don't know just what they are. So we can make progress in studying the meaning of Ss without having to say just what meanings are. Generalizing (outrageously!) there is possibly an epistemological lesson here:

[1] **Whether two objects stand in some relation is easier to verify/refute than whether a single object has a given property**.

The relational statement *Ted is taller than Ned* is easier to verify/refute than *Ted is tall*. And many apparent properties are covertly relations: whether 5 has the *propert*y of being prime depends on whether 5 stands in the *divisibility* relation to certain numbers. When I was in school, the claim that my desk was one meter wide said it stood in the *has the same length as* relation to a specific object in some temperature controlled vault outside Paris, France.

2. Generalize! Generalize! Generalize!

In any research field, once we discover a generalization we naturally want to apply it to new cases explaining novel phenomena. Still, I find cases in the logical and linguistic work of interest here where an acute thirst for generalization has not been slaked. We study two examples: the ubiquity of boolean operations, and the application of generalized quantifiers to n+1-ary predicates.

2.1 Ubiquity of Boolean Compounds: Syntactic problems: Semantic solutions

In English most categories of expression permit the formation of *boolean compounds* using (*both*)...*and*, (*either*)...*or*, *not*, *neither...nor* among others. Here are a few representatives drawn from categories other than S:

(2)	She smiled but didn't laugh	P_1s (Verb Phrases)
	An attractive but not very well built house is for sale	Adjective Phrases
	She lives neither in nor near New York city	Prepositions
	Sue either praised or criticized each student	Transitive Verbs
	Kim works rapidly but not carefully	Manner Adverbs
	All doctors and most nurses complain about that	Quantified NPs
	John and either Mary or Sue will arrive early	Quantified NPs
	Most but not all poets are vegetarians	Determiners

I addition English has many ways of expressing boolean operators besides the standard logical vocabulary just cited: *at most half* = not *more than half*; *only*

John = John and no one but John, exactly ten = at least ten and not more than ten, *between five and ten = at least five and not more than ten*. Note too that a sentence may host many boolean compounds simultaneously:

(3) *Not one teacher either praised or criticized each student and each dean*

Even when iterated application of boolean operators is cumbersome speakers usually can decode their meaning. The linguist might say we are *competent* to understand boolean compounds even if they present *performance* difficulties. This tells us that the sets in which expressions of these diverse categories denote are ones that support a *boolean* structure – whence we interpret boolean compounds compositionally, as *boolean* functions of the denotations of the expressions they are formed from.

The ubiquity of boolean compounding suggests that it is not specific to any particular category but rather represents ways we think about things rather than properties of things themselves. Maybe ...

[2] Boolean operations are properties of mind ("laws of thought" Boole 1854)

Somewhat embarrassingly linguists initially represented boolean compounds taking their lead from sentential logic where *and*, *or*, and *not* combine with sentences to form sentences. Cases in which we coordinated or negated non-Ss were derived from Ss by "Conjunction Reduction" rules (named but not defined) that deleted repeated material leaving the *and*s, *or*s and *not*s lying between non-Ss. The hope was that the derived Ss had the same meaning as what they were derived from. This works for cases like (4a,b) but not in general, e.g. (5a,b).

(4) a. Ben laughed and Ben cried b. Ben laughed and cried

(5) a. Some boy laughed and some boy cried b. Some boy laughed and cried

(4a,b) are logically equivalent, each entails the other. But clearly (5a) does not entail (5b) as (5a) can be true if some boys laughed and some cried but none did both, making (5b) false. Similarly *More than half the boys laughed and more than half the boys cried* does not entail *More than half the boys laughed and cried.*

The *syntactic* generalization is easy: replace 'S' by names of other categories in the rules we use to form boolean compounds – coordinations and negations. But *semantically* we must say how to interpret boolean compounds of non-Ss.

2.2 A (brief) boolean moment: *Boolean structures* are usually given as boolean algebras – a set B with various functions satisfying certain axioms: two binary functions $\wedge$ and $\vee$ (called *meet* and *join*) that provide denotations for *and* and *or*. Both functions are required by axiom to be *commutative*: $x \wedge y = x \wedge y$ and similarly for $\vee$. Further each *distributes* over the other, e.g. $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ and dually for joins. In addition we have a one place *complement* function $\neg$, (negation) and distinct elements 0 and 1, satisfying $(x \wedge \neg x) = 0$ and $(x \vee \neg x) = 1$. We define the boolean *relation*, $\leq$, by: $x \leq y$ iff $(x \wedge y) = x$ (equivalently $(x \vee y) = y$).

Here we define boolean structures B by taking the boolean relation $\leq$ as primitive, stipulate axioms it satisfies (it behaves like the subset relation $\subseteq$), then define the boolean functions in terms of it. So here **and** is a (finite) greatest lower bound (glb) operator and **all** is an unbounded one. x is a *lower bound* for a subset K of B iff $x \leq y$ for all y in K. And x is *greatest* if for every lower bound z for K, $z \leq x$. We write $x \wedge y$ for the glb for $\{x,y\}$, and for an arbitrary subset K, $\bigwedge K$ is its glb. Dually an *upper bound* for K is a y with $x \leq y$, all x in K. y is *least* if $y \leq z$, all upper bounds z for K. We write $x \vee y$ for the lub for $\{x,y\}$ and $\bigvee K$ for the lub for K. The 0 element is $\leq$ all elements, and all elements are ≤ 1. As before $x \wedge \neg x = 0$ and $x \vee \neg x = 1$. **End boolean moment**[1]. And another lesson from logic:

[3] If you can't say something two ways you can't say it

The important properties of objects of study are those that remain invariant under changes of descriptively comparable notation. So describing something in distinct if isomorphic terms helps us distinguish important properties from ones that are artifacts of notation. We'll see other cases shortly. Linguists have sometimes hurt themselves by insisting on one notation over another rather than being more flexible. Had we insisted on roman numerals over those invasive arabic ones we might never have learned to multiply and divide!

Now, to define the boolean structures associated with different categories C of expression we have to say what the *domain* of C is (the set in which expressions of category C denote) and what its $\leq$ relation is. E.g. the domain of S (in PL) is the set $\{0,1\}$ of truth values, and the $\leq$ relation is given by: $x \leq y$ iff $x = 1$ and $y = 1$, or $x = 0$ and y is either 0 or 1. So if $x \leq y$ then the conditional sentence *if p then q* is true, 1, when p has value x and q has value y, and conversely.

The boolean relation for many (not all) other categories is built uniformly from that for Ss by iterated **pointwise lifting**.

Consider P_1s, like *laughed*, *laughed and cried*, *laughed loudly*, etc. whose domain is the set of functions from E, the set of objects we're talking about, into the set $\{0,1\}$ of truth values. And for f and g such functions we define $f \leq g$ iff for all b in E, $f(b) \leq g(b)$. That is, if $f(b) = 1$ then $g(b) = 1$. For example **laughed loudly** $\leq$ **laughed** since if (**laughed loudly**)(john) = 1, i.e. John laughed loudly, then **laughed**(john) = 1, i.e. John laughed. From this it follows that meets and joins behave pointwise: (**laughed and cried**)(b) = **laughed**(b) $\wedge$ **cried**(b): Ed laughed and cried iff Ed laughed and Ed cried, and analogously for join. We note as well that **laughed and cried** $\leq$ **laughed** $\leq$ **laughed or cried**.

Quantified NPs (QNPs) like *all students*, *most poets*, as well as *John and some student*, combine with P_1s to form Ss (P_0s, zero place predicates). Here it is natural to interpret QNPs as functions, called Generalized Quantifiers (GQs), mapping P_1 functions to truth values. Its boolean relation $\leq$ is, again, lifted pointwise, this time from the P_1 relation. So $(F \wedge G)(P) = F(P) \wedge G(P)$ and $(\neg F)(P) = \neg(F(P))$. That (6a,b) are mutually entailing as are (7a,b) shows this to be correct.

(6) a. Every student and some teacher arrived early
b. Every student arrived early and some teacher arrived early

(7) a. [Not [every student]] arrived early
b. It is not the case that every student arrived early

We interpret CNPs (Common Noun Phrases) *student*, *enthusiastic student*, etc. as subsets of the universe E. The boolean relation is just the subset relation, $\subseteq$. For f a P_1 function write f* for the set of b in E that f maps to 1. Then $f \leq g$ iff $f^* \subseteq g^*$, so the boolean structure of CNPs is isomorphic to that of P_1.

Lastly, Determiners combine with CNPs to form QNPs and their boolean structure is also inherited pointwise (writing $\equiv$ for *logically equivalent* i.e. interpreted the same):

(8)	(most but not all)(poets)	$\equiv$	most(poets) and (not all)(poets)
		$\equiv$	most(poets) and not(all(poets))

2.3 Solving Syntactic Problems Semantically

Thus in a large variety of categories boolean compounds are interpreted by pointwise lifting, thus providing a uniform semantic solution to the problems posed by "Conjunction Reduction" and greatly simplifying the syntax. We independently generate boolean compounds in diverse categories and interpret them by pointwise lifting from the boolean structure for S.

We must note though that accounting for the equivalence in (4) and the non-equivalence in (5) involves saying what kind of GQs proper nouns denote and just what functions are denoted by various Determiners (e.g. *some* in (5)).

Proper Nouns are a subcategory of QNPs, so plausibly the GQs they denote have some properties not shared by other GQs. Consider F below:

(9) a. $F(p \wedge q) = F(p) \wedge F(q)$ b. $F(p \vee q) = F(p) \vee F(q)$ c. $F(\neg p) = \neg(F(p))$

(9) says that F is a *structure preserving function*. (9a) says it maps a meet of P_1s to the meet of what you get when you apply it separately to each of the properties. (9b) is the analogous claim for joins, and (9c) says that F maps the complement of p to the complement of what you get when you apply it to p. A function meeting these conditions is (a boolean) *homomorphism*). So we require:

2.4 Conditions on the Semantic Interpreting Function for English

1. Proper Nouns are interpreted by homomorphisms from P_1s to P_0s (and more generally from P_{n+1}s to P_ns). The equivalence of (4a,b) follows. We shall refer to such homomorphisms as **individuals** and use variables in I to range over them[1].
2. *some* and *all* denote as below, p any property, f* the set of objects f maps to 1

$$\mathbf{some}(p) = \bigvee\{I|I(p) = 1\} \quad \text{and} \quad \mathbf{all}(p) = \bigwedge\{I|I(p) = 1\}. \text{ Thus}$$

$$\begin{aligned}
\mathbf{some}(\mathbf{boy})(\mathbf{cried}) &= (\bigvee\{I|I(\mathbf{boy}) = 1\})(\mathbf{cried}) \\
&= \bigvee\{I(\mathbf{cried})|I(\mathbf{boy}) = 1\} \\
&= 1 \text{ iff for some I such that } I(\mathbf{boy}) = 1, I(\mathbf{cried}) = 1
\end{aligned}$$

$$\begin{aligned}
\mathbf{all}(\mathbf{boy})(\mathbf{cried}) &= (\bigwedge\{I|I(\mathbf{boy}) = 1\})(\mathbf{cried}) \\
&= \bigwedge\{I(\mathbf{cried})|I(\mathbf{boy}) = 1\} \\
&= 1 \text{ iff for each I such that } I(\mathbf{boy}) = 1, I(\mathbf{cried}) = 1
\end{aligned}$$

Given **some** as above the reader can find possible denotations for *boy*, *laughed* and *cried* such that (5a) is true and (5b) is not. And note: we see that **some**, like **or**, is a least upper bound operator and **all**, like **and**, is a greatest lower bound operator. (If memory serves me well, Carnap once tried to think of *all* as syntactically a kind of infinitary *and*, for which Tarski took him to task. It is better said semantically, as here.

A final, and deeper merit, of our focus on boolean structure is:

2.5 Theorem For each E, every GQ over E is a boolean function of individuals

This theorem (see Keenan 2018:78 for a readable proof) states that given any GQ

H, we can combine some individuals in boolean ways ($\wedge$, $\vee$, $\neg$), the result taking the same values at all P_1s as H does. This says, in a sense, that our logic is *extensional*. And it goes some way in accounting for why linguists early on might have taken individual denoting NPs as representative of NPs in general. And another modest lesson:

[4] Variable Binding and Quantification are independent

The semantic interpretation of quantifiers as above is naturally given in the set theoretical format illustrated. No variable binding is required, or natural. The formulation of variable binding operators, especially the lambda operator, has been one of the most significant contributions of logic to linguistic analysis. The lambda operator allows us to represent binding of pronouns, reflexives, even gaps by antecedents, including quantified ones. It is also useful in representing scope ambiguities. But it is not needed for quantification. In (10) we use the lambda operator to bind variables building a P_1 but the QNP is no different than when it occurs in Ss with no variable binding operators:

(10) All children love their mothers (**All child**)(λx(x **love** x**'s mother**))

3. Application of Generalized Quantifiers to n+1-Place Predicates

All occurrences of QNPs discussed so far have been in subject position. But they also occur as non-subjects (regardless of whether the subject is quantified).

(11) a. John offended <u>every linguist</u>
b. No editor read <u>every paper submitted</u>
c. Several witnesses told <u>each detective</u> <u>more than one lie</u>

To my surprise I find that basic textbooks (e.g. Heim & Kratzer 1998:184 – 189) claim that even simple Ss like (11a), their example, cannot be directly interpreted because of a "type mismatch". They treat QNPs as mapping P_1 functions to truth values and stick with that interpretation in (11a), so they cannot interpret *offended every linguist* by function application since **offend** is not a P_1 function but rather a P_2 one, mapping two individuals (in succession) to a truth value. So **offend** is not in the domain of **every linguist** and Heim and Kratzer do indeed have an interpretation problem. To solve it they expand the syntax of (11a) so that *every linguist* occurs outside the scope of *offend* and takes a (derived) P_1 as argument.

But there is another option, *Generalize!* Just interpret QNPs as maps from n+1 place predicates to n-place ones, all n. This maintains the natural, and simpler, syntax of Ss like (11a) and yields a directly compositional interpretation. For x an object in E and R a binary relation on E of type (e,(e,t)), write xR for the set of y which are such that x stands in the relation R to y. That is, $xR =_{def} \{y|xRy\}$. Now let F map properties of objects to truth values, and add binary relations to the domain of F by setting $F(R) = \{x|F(xR) = 1\}$. So the value of F at relations R is uniquely determined by the values F takes at sets. This is a way of saying that its behavior at P_1s determines its behavior elsewhere. So *offend every linguist* will denote {x|(**every linguist**)(x **offend**) = 1}. That is, *offend every linguist* denotes the set of x such that **every linguist** holds of the set of objects x offended, which is what we desire. The uniform statement extending GQs in this way to P_{n+1}s in general is given in Keenan & Westerståhl 1997 and Keenan 2018).

Generalizing has an additional advantage: namely, it leads us to find many denotable functions from P_2s to P_1s that are novel, not extensions of functions from P_1s to P_0s. Anaphors are one set:

(12) a. All poets admire <u>themselves</u>
b. Some student embarrassed <u>both himself and a teacher</u>
c. John criticized <u>every student but himself</u>

In standard English, reflexive pronouns – *himself*, *herself*, *themselves*, etc. do not occur naturally as subjects and so do not generalize semantically from subject uses to object uses (for proofs see Keenan 1992). And while their interpretation in (12) is straightforward – **self**(R) = {x|xRx} – the functions they denote are not the extensions of any subject GQ function (restricted to P_2s). The functions are new[3]. As items that coordinate are normally semantically comparable the approach in Heim and Kratzer does not expect coordinations such as that in (12b). **A teacher** only takes P_1 denotations as arguments (some derived) while *himself* here only takes P_2 denotations as arguments. So they have nothing semantically in common on their view.

On our generalization view, both **self** and **a teacher** map P_2s to P_1s so it makes sense to coordinate them in such contexts and interpret the coordination pointwise. Note that an attempt to derive (12b) from a coordination of Ss repeats the problem in (5). Equally (12c) is not the boolean compound some have suggested: *John criticized every student but John didn't criticize himself.* This S actually entails that John is not a student, whereas (12c) entails that John is a student).

To close this section I note an objection to generalizing QNP denotations to P_{n+1}s. Namely, aren't QNPs now semantically ambiguous? The answer is NO.

They just have a larger domain than they did when we thought of them as just taking P_1s as arguments. So they assign values to some new things, but thing in the richer domain still just gets one value, so there is no ambiguity.

4. Wrapping Up

We have discussed some points where the conceptual tools of PL (Predicate Logic) have been useful in elucidating semantic properties of NL (natural language). This has led us to consider a first study of the boolean structure of NL. We can go much farther, and lecture 2 will mention some further results in this area. As a more general lesson from this work I would offer:

UNDERSTANDING IS TRANSLATION FROM THE UNKNOWN TO THE KNOWN

If you are to get me to understand something new you must present it to me in terms that I understand. A description in Hittite avails me little. The enormous (unreasonable?) utility of mathematical models is that we understand them – we made them up, defined them rigorously, precisely, gave examples. We know what functions are and what boolean structures are – so anything you can explain to me in such terms can be understood. We cannot of course limit ourselves to boolean models – perhaps next time around it will be Bayesian statistics, or game theory or modal logic (van Benthem 2010) or .... Myself I have a weakness for algebraic models, as algebras are made up to have non-isomorphic models. Had we more time we would have seen that the boolean structures we associate with predicates, quantified NPs, Determiners, and Modifiers (like Adjectives and Adverbs) each has a distinctive property. Learning these algebras is a little like learning different human languages – delightfully diverse, but all cut from the same plan.

FOOTNOTES

1. Here is a complete relation based definition of *boolean structure*. A *boolean lattice* is a pair $(B,\leq)$ for B a set (the *domain* of the lattice) and $\leq$ is a partial order relation (*reflexive*: $x \leq x$, all x; *transitive*: $(x \leq y$ and $y \leq z) \Rightarrow x \leq z$; and *antisymmetric*: $(x \leq y$ and $y \leq x) \Rightarrow x = y$, so no two different objects stand in the $\leq$ to each other. For all x,y in B, $\{x,y\}$ has a glb and a lub, provably unique and noted $x \wedge x$ and $x \vee y$. A z is a *lower bound* (*lb*)for a subset K of B iff for all x in K, $z \leq x$. z is *greatest* of the lbs if all lbs y for K are $\leq z$. We write $x \wedge y$ for the glb of $\{x,y\}$. Dually z is an *upper bound* (*ub*) for K iff for all x in K, $x \leq z$, and z is least of the upper bounds iff for all ubs y for K, $z \leq y$. Glbs are required to distribute

over lubs and vice versa (per the text). 0 is the glb of B and 1 its lub; $0 \neq 1$. As before, $x \wedge \neg x = 0$ and $x \vee \neg x = 1$. In general we require that all the boolean lattices we have mentioned are complete, meaning that for all subsets K of B, K has a glb and a lub. (All lattices with finite domains are complete in any event).

2. Technically we use *complete* homomorphisms, ones that commute with arbitrary glbs and lubs: $I(\bigwedge\{p_j | j \in J\}) = \bigwedge\{I(p_j) | j \in J\}$, usually noted $\bigwedge_{j \in J} I(p_j)$, and analogously for $\bigvee$.

3. When an extended GQ decides to put some x in the set it associates with a binary relation R, it makes its decision solely by looking at the set xR. It doesn't have to know which individual x is. This means that for x,y objects and R,S binary relations and F a GQ, $x \in F(R)$ iff $y \in F(S)$ if $xR = yS$. E.g. if Bob likes just the people Rosa admires then Bob likes most poets iff Rosa admires most poets. But for anaphors this is not the case. Suppose that the set of people Bob trusts is {Miriam, Rosa, Ted, Zelda, Dana} and that that is exactly the set of people that Ted distrusts. Then *Bob trusts himself* is false but *Ted distrusts himself* is true.

REFERENCES

Boole, George. 1854. *An Investigation of the Laws of Thought*. Reprint. The Open Court Pub. Co.

Chomsky, Noam. 1957. *Syntactic Structures*. Mouton

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. MIT Press

Heim, Irene and Angelica Kratzer. 1998. *Semantics in Generative Grammar* Blackwell Pubs.

Keenan, Edward L. 2018. *Eliminating the Universe*. World Scientific

Keenan, Edward L. 2016. *In Situ interpretation without type mismatches* in: *Journal of Semantics* 33.1:87–106.

Keenan, Edward L. 1992. Beyond the Frege Boundary. *Linguistics and Philosophy* 15:199–221

Keenan, Edward L. and Dag Westerståhl. 1997. Generalized Quantifiers in Linguisitics and logic. In *Handbook of Logic and Language*. Johan van Benthem and Alice ter Meulen (eds). North Holland. Pp. 859 – 910.

van Benthem, Johan. 2010. *Modal Logic for Open Minds.* CSLI Publications