

A BRIEF HISTORY OF NATURAL LOGIC

Johan van Benthem, ILLC Amsterdam & Stanford, <http://staff.science.uva.nl>

Abstract

This paper is a very brief history of natural logic at the interface of logic, linguistics, and nowadays also other disciplines. It merely summarizes some facts that deserve to be common knowledge.

1 From classical to modern logic

For beginning logic students in Amsterdam in the 1960s, the following ‘standard example’ was supposed to show once and for all how the modern 19th century logic of Boole and Frege came to supersede traditional logic. De Morgan’s famous inference runs as follows:

“*All horses are animals. So, all horse tails are animal tails.*”

This is supposed to show the inadequacy of the traditional logic of ‘monadic predicates’, because binary relations are essential to understand the validity of the inference. And the latter are brought out in the standard first-order logical forms that students are trained in:

$$\forall x (Hx \rightarrow Ax) \models \forall x ((Tx \& \exists y (Hy \& Rxy)) \rightarrow (Tx \& \exists y (Ay \& Rxy))).$$

What is more, we can understand the general phenomenon going on here as follows. Semantically, we are replacing a predicate by one with a larger extension, and we have *monotonicity* for the relevant assertion involving that predicate:

$$\text{if } \mathcal{M}, s \models \varphi(P), P^{\mathcal{M}} \subseteq Q^{\mathcal{M}}, \text{ then } \mathcal{M}, s \models \varphi(Q).$$

But this semantic behaviour has a syntactic counterpart, as may be checked with the above first-order formulas. The occurrence of the predicate ‘horse’ in ‘tail of a horse’ is *positive*: which means that it lies in the scope of an even number of negations, or stated differently, that its environment of the occurrence of ‘H’ is created only from these syntax rules:

$$H\text{-free formulas} \mid \& \mid \vee \mid \forall \mid \exists.$$

It is easy to see that syntactic positive occurrence implies semantic monotonicity (a sort of ‘soundness’, if you wish), but it is much less trivial that the converse ‘completeness’ direction holds. But it does, witness a well-known model-theoretic result from the 1950s:

Lyndon’s Theorem A first-order formula $\varphi(P)$ is semantically monotone in P iff $\varphi(P)$ is equivalent to a formula whose only occurrences of P are positive.

Lyndon's Theorem does not hold for all extensions of first-order logic, but positive occurrence does imply monotonicity in many higher-order logics – and natural language. Thus, modern logic provides the right forms for the inference, and it backs these up with important meta-theorems probing the broader extent of the phenomenon.

2 Distribution in traditional logic

But the usual De Morgan story is misleading and historically quite false. Inferences like the one with the horse tail were well within the scope of traditional logic, which was much subtler than its modern critics acknowledge. They blame it for defects it never had – all the way to the invective of Peter Geach, who even saw demoniacal political consequences, in his phrase 'The Kingdom of Darkness'. Indeed, the Aristotelean Syllogistic was the main tool of traditional logic, and this might seem merely a trivial theory of single-quantifier inferences with patterns $Q AB$ on unary predicates A, B . But this view in terms of 'formal systems' is a modern way of looking which does no justice to what the Syllogistic really was: a method for *one-step top-down analysis* of statements of any kind into one layer of quantification. In particular, A, B could be predicates with further structure, of whatever expressive complexity: they are not constrained at all to some fixed formal language.

Later on, in addition to this base system, the Medieval Scholastics (another much-maligned set of subtle minds, guilty of less angel-pinpointing than the average philosopher today) developed the so-called *Doctrine of Distribution*, a general account of contexts $\varphi(P)$ where the statement was either about 'all of P ' (the 'Dictum de Omni'), or about 'none of P ' (the 'Dictum de Nullo'). Again, these contexts could be of any sort of linguistic complexity, where the expression φ might include iterated quantifiers such as "Someone loves every human", or even high-order constructions. The modern first-order/higher-order boundary with its extensive customs facilities seems mainly a modern mathematical 'systems concern' without any clear matching jump in natural reasoning.

Authors like van Van Eijck 1982, van Benthem 1986, Sanchez Valencia 1991, and Hodges 1998 have pointed out in more detail how the distribution principle called the "Dictum de Omni et Nullo", corresponded to admissible inferences of two kinds: *downward monotonic* (substituting stronger predicates for weaker ones), and *upward monotonic* (substituting weaker predicates for stronger ones). Traditional logic investigated these dual phenomena for a wide range of expressions, without any boundary between unary and binary predication – another artefact of wearing predicate-logical glasses as one looks at history.


To be sure, this does not mean that all was clear. To the contrary, traditional logic had a major difficulty: providing a systematic account of complex linguistic constructions from

which to infer, and in particular, despite lots of valid insights, it wrestled with a good general account of *iterations of quantifiers*. Dummett 1973 makes a lot of this, by saying that Frege's compositional treatment in terms of merely explaining single quantifiers and then letting compositionality do all the rest "solved the problem which had baffled traditional logicians for millennia: just by ignoring it". Again, while there is a kernel of truth to this, there is also a good deal of falsehood. Indeed, as the extensive historical study Sanchez 2007 remarks, it seems more fair to say that De Morgan represents a low point in logical history as far as understanding the scope of monotonicity reasoning is concerned. Things got better after him – but as the author says, they also got better and better moving back in time to Leibniz and then on to the Middle Ages...

Perhaps not surprisingly then, traditional logicians felt some unfairness when modern logic arrived, since it attacked a caricature of traditional logic. And until late in the 20th century, attempts have been made to further develop the Syllogistic into a full-fledged calculus of Monotonicity reasoning, witness Sommers 1982 (going back to versions from the 1960s), Englebretsen 1981. The claim of these authors was that this enterprise provided a viable alternative to first-order logic for bringing out key structures in actual human reasoning in a more congenial way. Still they did not propose turning back the clock altogether. E.g., Sommers' book is up to modern standards in its style of development, providing a systematic account of syntactic forms, an arithmetical calculus for computing positive and negative syntactic occurrence, as well as further inferential schemata generalizing traditional inference patterns like Conversion and Contraposition. While this has not led to a Counter-Revolution in logic, these points have resonated in other areas, such as linguistics, computer science, and recently also, cognitive science.

3 Monotonicity in natural language and generalized quantifier theory

Very similar issues came up in the 1970s and 1980s when linguists and logicians started looking at natural language together (there are several famous joint papers) with fresh eyes, following Richard Montague's pioneering work. Suddenly, natural language was no longer 'misleading' as to the correct logical forms, but rather a gold-mine of intriguing insights with remarkable staying power. After all, no one, not even pure mathematicians, has ever seriously switched to predicate logic as a tool of reasoning (but see below). In particular, Montague provided a categorial/type-theoretic analysis of quantifier expressions Q as taking linguistic noun phrases A to noun phrases QA that denote properties of properties B . Deconstructing this forbidding phrasing somewhat, quantifiers may then be viewed semantically as denoting binary relations between predicates, on the pattern



E.g., in this Venn-style diagram format, “All A are B ” says that the area $A-B$ is empty, “Some A are B ” that the intersection $A \cap B$ has at least one object in it, while the more complex “Most A are B ” says that the number of objects in $A \cap B$ exceeds that in $A-B$.

Generalized quantifiers and monotonicity A more general background was given by *Generalized Quantifier Theory*: a research program which charted the variety of the quantifier repertoire of natural human languages (cf. Keenan & Westerståhl 1997). Again, Monotonicity turned out to play a crucial role across this repertoire. One influential example here was the observation by Bill Ladusaw that so-called ‘negative polarity items’ like “at all”, or “ever” flag ‘negative contexts’ in linguistic expressions which allow for downward monotonic entailments from predicates to sub-predicates:

“If you ever *feel an ache*, I will cure it” implies
 “If you ever *feel a headache*, I will cure it”.

Negative polarity items do not generally occur in upward entailing positive contexts. Here are some general facts about monotonicity for basic quantifiers: they can be either upward or downward, in both arguments. E.g., the quantifier “All” is downward monotonic in its left-hand argument, and upward in its right-hand argument, exemplifying the patterns

$$\begin{array}{ll} \downarrow MON & Q AB, A' \subseteq A \ / \ Q A'B \\ MON \uparrow & Q AB, B \subseteq B' \ / \ Q AB' \end{array}$$

It is easy to exemplify the other three possible combinations: e.g., “Some” is $\uparrow MON \uparrow$. By contrast, a quantifier like “Most” is only $MON \uparrow$, being neither ‘down’ nor ‘up’ on the left.

Conservativity But Monotonicity is not the only key property found with natural language quantifiers (and *NP*-forming determiner expressions in general, such as “Mary’s”). Here is a further principle which shows how the first argument sets the scene for the second:

$$\text{Conservativity} \quad Q AB \text{ iff } Q A(B \cap A)$$

Conservativity seems to hold in all human languages. One can think of this as a sort of domain or role restriction imposed by the initial predicate A on the predicate B . More generally, the nouns in sentences give us relevant domains of objects in the total universe of discourse, and quantifiers impose a sort of coherence on the total predication expressed.

Taken together, the preceding semantic properties explain what makes particular (logical) notions so special. The following result from van Benthem 1986 shows that the traditional quantifiers are the simplest level of *conservative, inference-rich* linguistic expressions:

Theorem The quantifiers “All”, “Some”, “No”, “Not All” in the Square of Opposition are the only ones satisfying Conservativity, Double Monotonicity, and Variety.

The third more technical condition says that the quantifier makes maximal distinctions:

Variety If $A \neq \emptyset$, then $Q AB$ for some B , and $\neg Q AC$ for some C .

For further information on Generalized Quantifier Theory, cf. Peters & Westerståhl 2006. In particular, quantifiers can also be classified by more algebraic types of inferential property for specific lexical items, such as the ‘Conversion’ of traditional logic:

Symmetry $Q AB$ iff $Q BA$.

This holds typically for expressions like “Some”, “At least n ”, “No”, “All but at most n ”. Characterizations exist of all quantifiers satisfying Symmetry and other basic properties.

4 The ‘natural logic’ program

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.

But of course, one can think of many further natural subsystems inside natural language, such as reasoning about collective predication, prepositions, anaphora, tense and temporal perspective, etcetera. Of course, the challenge is then to see how much inference can be done directly on natural language surface form, and we will look at some details below.

Notice however, how this cuts the cake of reasoning differently from the syntax of first-order logic – making the border-line between traditional and modern logic much more intriguing than what either De Morgan or Sommers would have said. E.g., monotonicity inference is both richer and weaker than first-order predicate logic. It is weaker in that it only describes a part of all possible quantifier-based inferences, but it is richer in that it is

not tied to any particular logical system, as we observed above (it works for second-order just as well as first-order logic). One intriguing aspect then becomes *where the surplus of first-order logic is really needed*. We will give two possible answers to this below.

5 Compositional structure, parsing, and marking inferences

To get to the true power of the earlier observations, one needs to see how they play in complex sentences (this Section is taken largely is from van Benthem 1986, 1991).

Spreading conservativity A first example concerns the broader impact of Conservativity. Consider an iterated quantifier sentence of the sort “Every man loves a woman”:

$$Q_1 A R Q_2 B$$

Clearly, both predicates A and B should have restriction effects. How do they do this? It is not hard to see, and in fact it can be computed in many parsing formalisms, that we have

$$Q_1 A R Q_2 B \text{ iff } Q_1 A R \cap (AxB) Q_2 B$$

That is, the first predicate restricts the first argument of the binary relation R , while the second restricts the second argument. Thus in general, nouns constrain predicate roles.

Monotonicity in complex sentences Likewise, in order to get the correct monotonicity inferences in complex sentences, we need grammatical theory for analysis of hierarchical structure. Flat strings will not do. E.g., I cannot tell whether the occurrence of “ladies” in

“Johan harmed no men and ladies”

is upward monotonic unless I resolve the scope of the negation “no”. (Without grammar, no morality.) To achieve this, we can use a mixture of a logic-friendly grammar formalism, viz. *Categorial Grammar* and *Monotonicity Calculus*. For a start, consider the sentence

“No mortal man can slay every dragon.”

We would like to compute all predicate markings in this sentence, but they cannot be taken at face value. E.g., whether “dragon” is negative depends on the scope of other expressions inside which it occurs. Maybe the reader wants to check that, on the narrow scope reading for “every dragon”, readings should come out (intuitively) as follows:

+ - - - - - +
 “No mortal man can slay every dragon.”

E.g., it follows that no mortal *Dutchman* can slay every dragon, or that no mortal man can slay every *animal*. For the wide scope reading of ‘every dragon’ (artificial, but still a nice illustration), these markings should come out as follows:

+ - - - - + -
 “No mortal man can slay every dragon.”

Readers may want to check this with first-order transcriptions if their untutored intuition fails them. Perhaps surprisingly, other expressions than predicates can also be marked here: e.g., we can validly replace the quantifier “no” by the weaker quantifier “no or almost no”.

Categorial monotonicity calculus Speaking generally, we need a linguistic mechanism marking positive/negative occurrences in tandem with syntactic analysis of an expression. This can be done elegantly in a *categorial grammar*, of the Ajdukiewicz type or of the more sophisticated Lambek type which is more like a simple system of function application plus limited (‘single-bind’) lambda abstraction. For details, we refer to van Benthem 1991, Sanchez Valencia 1991. Here we only state the major rules of the procedure:

Rules come in two kinds:

- (a) General rules of composition:
 occurrences in a function head A in applications $A(B)$ retain their polarity,
 occurrences in a body A of a lambda abstract $\lambda x.A$ retain their polarity,

where function heads can block the monotonicity marking in their arguments. Notice, e.g., how “Most AB ” made its left-hand argument ‘opaque’: it is neither upward or downward. But “Most AB ” does pass on monotonicity information in its right-hand argument B , and this demonstrates a second crucial source of information for our calculus:

- (b) Specific information about lexical items:
 e.g. “All” has functional type $e^- \rightarrow (e^+ \rightarrow t)$.

Here is how the two kinds of information combine. First, in general, a function application $A(B)$ may block the polarity marking of positions in the argument B . E.g., “best (friend)” has no marking left for “friends”, as there is no inference to either “best girlfriend” or “best acquaintance”. The adjective “best” is highly context-dependent and hence steals the show.

But sometimes monotonicity marking does percolate upwards, when the meaning of the function head A ‘helps’. E.g., “blonde friend” does imply “blonde acquaintance”, because the adjective “blonde” has a simple ‘intersective’ meaning forming a Boolean conjunction “blonde $\cap B$ ”. More generally, if a function head A has type $a \rightarrow b$ where the argument type a is marked, the argument position B in applications $A(B)$ will assume that same polarity.

This explains how negations switch polarity, how conjunctions just pass them up, and so on. Such markings will normally come from lexical information, but there is one nice twist. They can also be introduced via lambda abstractions $\lambda x_a \bullet M_b$ of type $a \rightarrow b$, where the type a gets positive marking. Readers may want to check the semantics for an explanation.

The final ingredient to make this work is the following self-evident rule of calculation. Markings can be computed as long as there is an unbroken string in the parse tree:

$$(c) \quad \begin{array}{cccc} + & + & = & + \\ - & - & = & + \\ + & - & = & - \\ - & + & = & - \end{array}$$

This is just one mechanism for making natural logic precise. Its basic categorial structure has been rediscovered by several people, including David Dowty: it just *is* rather natural! The general insight is this. Monotonicity marking can work in tandem with one's preferred style of syntactic analysis for natural language – providing fast inferences ‘on the fly’.

A bit of deconstruction Here is a summary of all that has been said in a *Boolean typed lambda calculus*, the system behind much of the above (van Benthem 1991). Expressions now have ‘marked types’, and we define inductively what it means for an occurrence of a sub-expression to be positive or negative in an expression:

The occurrence x_a is positive in the term x_a ,

The head M occurs positively in applications $M_{a \rightarrow b}(N_a)$,

If M has type $a^+ \rightarrow b$, then N occurs positively in $M_{a \rightarrow b}(N_a)$,

If M has type $a^- \rightarrow b$, then N occurs negatively in $M_{a \rightarrow b}(N_a)$,

The body M occurs positively in $\lambda x_a \bullet M_b$, and the resulting type is $a^+ \rightarrow b$.

The rest is the earlier computation rule (c), or alternatively, we could have built this feature into the inductive definition. Clearly, this definition can be extended to deal with, not just functional types, but also *product types* $a \bullet b$ allowing for pair formation of objects.

Metatheory: Lyndon theorems This logical perspective raises further issues of its own. One is the ‘completeness’ of the above syntactic marking procedure. Can we be sure that every semantically monotone inferential position will be found in this way?

Is there a Lyndon Theorem stating that every semantically monotone occurrence must be positive in the above categorial sense?

This would extend the first-order result. The answer is ‘Yes’ in the categorial single-bind Lambek Calculus (this is a model-theoretic result proved by brute force in van Benthem 1991) – but the problem is still open for type theory with Boolean operators in general.

A lambda calculus for natural language has no internal first-/higher-order boundary. The only special thing to first-order quantifiers here is that they have more monotonicity markings than others, in line with our earlier observation about their inferential richness.

Bottom-up modern style after all? A final issue is that maybe, our ‘natural logic’ has turned out to be a rather modern Fregean system, requiring a fully-fledged parse of a sentence – just as logical systems require construction of fully explicit formulas from the atomic level up. Of course the above analysis can work in a top-down manner (van Eijck 2005). The only thing we need to know for the monotonicity marking of a constituent is the hierarchical structure of the sentence *above it*, leaving all other ‘side parts’ unanalyzed. Even so, I must confess that I am not entirely happy. Even so, the categorical monotonicity calculus is definitely not carefree surface analysis, and like much of current linguistics and logic, it analyzes more than the bare minimum which seems involved in natural reasoning. How can we ‘hit and run’ as reasoners, or is that idea just a chimera?

6 Further issues in natural logic

There are many further questions at this stage about the range of our inferences so far. For instance, there is much more to quantifier patterns in natural language than the above single and iterated cases. Through the 1980s, further quantifier combinations have come to light which do not reduce to simple iterations, such as *cumulative forms* “Ten firms own 100 executive jets” or *branching patterns* “Most boys and most girls knew each other”. These require new forms of monotonicity marking, depending on how one takes their meanings. Also, quantifiers also lead to *collective predication* (“The boys lifted the piano”), as well as *mass quantification* (“the teachers drank most of the wine”), whose inferential behaviour is far from being generally understood – either in linguistic semantics or in modern logic.

Other fast subsystems Maybe more interesting is the issue whether there are *other fast surface inference systems* in natural language. I already mentioned the general functioning of Conservativity as a mechanism of general *Role Restriction* for predicates in sentences. And I can think of several other examples, such as ‘individual positions’ X in expressions allowing for arbitrary *distribution over disjunctions*, as in $\varphi(X_1 \vee X_2) \leftrightarrow \varphi(X_1) \vee \varphi(X_2)$.

Another potential candidate was brought up just recently at the Stanford RTE seminar: *disjointness of predicates*. Which expressions preserve this? I would cast this issue as follows. Monotonicity was about abstract inferences of the form:

$$P \leq Q \text{ implies } \varphi(P) \leq \varphi(Q)$$

But now, we are now given, not an inclusion premise, but an *exclusion premise* (note that $P \cap Q = \emptyset$ iff $P \subseteq \neg Q$), and we want to know what follows then in the right contexts:

$$P \leq \neg Q \text{ implies } \varphi(P) \leq \neg \varphi(Q)$$

In first-order logic, this amounts to stating a monotonicity-like inference between the formula and its *dual* obtained by working the prefix negation inward switching operators:

$$P \leq Q \text{ implies } \varphi(P) \leq \varphi^{\text{dual}}(Q)$$

I think one can find first-order syntax which guarantees this. In any case, more generally, I think that many classical model-theoretic preservation results in first-order logic can be re-interpreted as giving simple special-purpose syntax for specialized inferences.

Interactions So much for separate inferential systems. Another issue is of course how these interact with other major features of natural language. E.g., *anaphora* with pronouns can wreak havoc with monotonicity inferences, as in the following classical example:

“Everyone with a child owns a garden.
Every owner of a garden waters it. *So*:
Everyone who has a child sprinkles it?”

Here, the pronoun “it” has picked up the wrong antecedent. Again, information about the total sentence composition is crucial to block these inferences, and keep the correct ones.

Inference without scope resolution? Finally, here is a more daunting challenge for surface reasoning. If we are to do inference as close to the linguistic surface string as possible, it would be nice to not have to resolve all quantifier scope ambiguities – and infer as much as possible from ambiguous expressions. (Note this was the philosophy of LFG in the 1980s.) This is the way language often comes to us. E.g., in the above two dragon examples, note that 5 out of the 7 positions in the string retain the same marking in either scope reading. Can we find a still more surfacy natural logic for inferences unaffected by ambiguity?

7 Back to the scholastics after all

Here is another direction in current research on natural logic. We saw that the medieval logicians started classifying multi-quantifier inferences. Thus, they were well-aware that

“Some $P R$ all Q ” implies “All Q are R -ed by some P ”

and that the converse fails. Now the iteration depth in ordinary discourse seems limited (at best 3 levels, as in “You can fool some people all of the time” seem to occur in practice). Thus it makes sense, pace Dummett, to define small special-purpose notations for such

combinations, and try to axiomatize them. Moss 2005 is a first attempt in this direction of rehabilitating Scholastic efforts, with very appealing sets of valid principles. (It has to be admitted though that higher-order quantifiers like “Most” have proved recalcitrant so far.)

In addition to this focus on completeness, Pratt 2006 has performed computational complexity analysis on several small decidable fragments of natural language. Here outcomes are a bit more gloomy, in that high complexity can arise – but as so often with bad news of this sort, it remains to be seen what this means in practice.

8 Computational issues

The story of natural logic has also crossed to computer science.

Efficient computation For instance, it has often been observed that simple reasoning with relational databases only uses small parts of predicate logic, and that monotonicity accounts for most of what is needed. Accordingly Purdy 1991 and follow-up publications have come close to the material covered in the above. Likewise, modern research on natural language processing, and in particular, intelligent text analysis appears to be arriving at similar aims and results. Results from the earlier linguistic phase have been promoted in van Eijck 2005, using various programming techniques for optimizing the monotonicity calculus and related forms of inference. For a more extensive empirical investigation of actual data, see Manning & MacCartney 2007. Likewise, polarity calculi for hierarchical marking to extend the scope of natural lexical inferences with factive verbs have been proposed in Nairn, Condoravdi & Karttunen 2006. Whether all this really becomes less complex than first-order alternatives is partly a matter of detailed complexity-theoretic analysis, as indicated earlier – and the jury is still out.

A technical aside The computational setting also suggests natural constructions outside of first-order logic, such as transitive closure (Kleene iteration) and recursive definitions. These involve extensions of first-order logic with *fixed-point operators*. But monotonicity still makes sense here, and indeed, it is crucial. A recursive definition $Px \leftrightarrow \varphi(P)(x)$ of a new predicate P does not make sense in general – but it does when $\varphi(P)$ is semantically monotone with respect to P . Is this just a technical coincidence, or does this mean something from the perspective of natural logic? Maybe it supports circular definitions?

From classical to default logic But there are further, less standard aspects to the computer science connection. In particular, ‘common sense reasoning’ has been analyzed extensively in Artificial Intelligence by John McCarthy and his school. Now this reasoning involves not just monotonic inferences, but also *non-monotonic* ones, where inclusion of predicates need not keep the conclusion true. This brings us to the area of default logics, which

involve both classical reasoning with rock-solid conclusions, and *defeasible inferences* which can be retracted when new information comes in. E.g., by default “birds fly”, but there are exceptions such as penguins, who tend to march... A systematic extension of the above monotonicity calculus to deal also with default implications based on predicate inclusions would be a highly interesting project! Existing logical systems in this area, which already combine material and default conditionals might provide a guide.

Combination, architecture and complexity My final computational theme returns to the natural logic program as such. Analyzing natural inference as a large family of simple fast subsystems is not enough! In reality, we are not a bare set of isolated processors. All this information must be *combined*, and one module must be able to feed quickly into another. So, what sort of ‘natural’ reasoning system are we? Here, an insidious challenge to the whole enterprise emerges. Much experience with logical systems over the past decade has shown that the analytical strategy of ‘divide and conquer’ does not always hold. The complexity of the total logic is not just a maximum of the complexities of the components. It can explode dramatically, because the additional source of complexity is *the nature of the combination*, i.e., the communication between the different inferential modules. Indeed, several natural examples are known where apparently innocuous combinations of decidable logics create *undecidable* over-all inference systems. Note, this does not have to happen, and Dov Gabbay’s technique of ‘fibered logics’ provides a way out in many cases. But the danger is there. Unless we have an additional idea about the *global architecture* of natural logic, claims about its performance may be premature and ill-founded.

9 Cognitive science

A final arena where natural logic is coming up these days is experimental cognitive science. We definitely know that inference in the human brain is not one unified phenomenon, but a joint venture between many modules, some related to our language abilities, some more to immediate visual processing or to schematic representation, and yet others to brain areas dedicated to planning and executive function (Knauff 2007). Current neuroscience experiments, guided by hypotheses about linguistic data (Geurts and van der Slik, 2005) are now beginning to map out how, e.g., monotonicity inference is located in different parts of the brain than heavy-duty first-order logic (if that is available at all).

10 History once more

Logic started simultaneously in at least three geographical areas and cultures: Greece, India, and China. Without much commentary, I just cite a few telling examples from Liu &

Zhang 2007 about Mohist logic (5th century B.C.), a school of thought clearly manned by logicians. Here are some examples from the Mohist Canon:

“A white horse is a horse. To ride a white horse is to ride a horse.”

This is clearly the pattern of upward monotonicity.

But now, here are two further examples which seem to contradict this:

“Robbers are people, but to abound in robbers is not to abound in people.”

“A cart is a wooden object. To ride a cart is not to ride a wooden object.”

But actually, the latter examples are subtle, and both highlight a further phenomenon. The second seems a failure of upward monotonicity due to *context dependence* of a quantifier. If “Many” just means ‘more than a fixed threshold value N ’, it is upward monotonic in both arguments. But if we assume that the norm is dependent on the predicate, as seems much more likely, than “Many” is not upward monotonic in either argument. Another form of context dependence played in an example brought up in a colloquium at PARC Palo Alto:

Does “They verbally attacked the president”
imply “They attacked the president”?

The conclusion suggests (incorrectly) physical attacks. To manage correct and incorrect inferences here, one would seem to need a dynamic mechanism of context management. In other words, the precise parameters of the total enterprise need to be determined.

The second example seems one of *intensionality*. ‘To ride a cart’ may be read extensionally as just ‘being transported’, and then the conclusion should indeed follow by upward monotonicity. But surely, our Mohist colleagues meant something more subtle. Intensionally, one rides a cart *as a vehicle*, and read in that way, the stated inference is invalid, since one does not ride a wooden object qua wooden object.

These refinements of the above monotonicity setting are as valid now as they were then.

Coda Mohist logic had many more subtle features, including versions of the Paradox of the Cretans, and very nicely, the following pragmatic paradox of conversation:

Telling someone that “You can never learn anything” can never be successful.

11 Conclusion: modern and traditional logic once more

Natural logic is a proposal for taking a new look at natural inferential systems in human reasoning, not through the lenses of modern logic which see ‘formal systems’ everywhere.

I think this is well-worth doing for all the reasons that I mentioned. But let me also be clear that I do not see this as a resumption of warfare between traditional logic and modern logic. If only by this ‘cheap shot’: the only respectable systems of natural logic that I know use thoroughly modern logical techniques and standards of exposition. There is no way back.

Architecture and transitions To me, by now, the more interesting question is one of *architecture* and *transitions*. Clearly, modern logic provides more subtle tools for analyzing reasoning than traditional logic, and it is of great interest to see where these must come into play. The analysis of mathematical reasoning is of course one clear case in point where traditional logic failed eventually in its accounts of density, continuity, limits, etcetera (cf. Friedman 1985 on Kant’s views of mathematics). But the transitions do not lie where De Morgan and many modern teachers claim they lie. Traditional logic was much richer than many people think, and it still deals in attractive ways with large areas of natural language and common sense reasoning. First-order logic probably takes over when explicit variable management and complex object constellations become inevitable. Thus, it seems to me, there is no inevitable conflict between ‘natural logic’ versus modern logic.

Redefining the issue: mixtures and merges But again, one cannot just see this peaceful co-existence in sweeping terms like ‘mathematics is modern logic’, ‘natural language is traditional logic’. The more interesting perspective for research may rather be

mixtures of natural and formal language and reasoning!

For instance, it is a telling fact that mathematicians have never abandoned natural language in favour of logical formalisms. Just read any mathematics paper, or go to any mathematics colloquium. The real situation is that mathematicians use mixtures of both, with the logical notation coming in *dynamically* when natural language needs to be made *more precise*. This mixture suggests that ‘natural logic’ and ‘modern logic’ can coexist harmoniously, because both have their place. And modern logic might even learn something from natural logic in combating its ‘system imprisonment’, trying to look for more general systems-free formulations of its basic insights, the same way, say, Monotonicity seems a general insight about human reasoning, which does not really seem to depend on any specific formal language and semantics. But I take this also as a shift in the whole issue. What we really need to understand is how our ‘natural logic’ can be ‘naturally extended’ with formal notations, and other technical inventions which somehow seem to fit our cognitive abilities.

The price of inferential holism Here is a final speculation. There might be another role then for modern first-order logic. Maybe modern logic is the only system which really *integrates all separate natural reasoning modules*. And then, as in my story of architecture

and module combination, there may be a price to pay. This might be the true reason for the *undecidability* of first-order logic: not because its subsystems are so hard by themselves, but *because their combination is*. This may be seen by linking undecidability to Tiling Problems and interaction axioms (van Benthem 1996), but I will leave the matter here.

12 References

- J. van Benthem, 1986, *Essays in Logical Semantics*, Reidel, Dordrecht.
- J. van Benthem 1987, ‘Meaning: Interpretation and Inference’, *Synthese* 73:3, 451-470.
- J. van Benthem, 1991, *Language in Action. Categories, Lambdas and Dynamic Logic*, North-Holland Amsterdam & MIT Press, Cambridge (Mass.).
- J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford.
- M. Dummett, 1973, *Frege. The Philosophy of Language*, Duckworth, London.
- G. Englebretsen, 1981, *Three Logicians: Aristotle, Leibniz, Sommers and the Syllogistic*, Van Gorcum, Assen.
- J. van Eijck, 1985, *Aspects of Quantification in Natural Language*, Dissertation, Philosophical Institute, University of Groningen.
- J. van Eijck, 2005, ‘A Natural Logic for Natural Language’, CWI Amsterdam & UiL-OTS Utrecht. Updated version 2007.
- M. Friedman, 1985, ‘Kant’s Theory of Geometry’, *The Philosophical Review* XCIV:4, 455 – 506.
- B. Geurts and F. van der Slik, 2005, ‘Monotonicity and Processing Load’, *Journal of Semantics* 22:1, 97 – 117
- W. Hodges, 1998, ‘The Laws of Distribution for Syllogisms’, *Notre Dame Journal of Formal Logic* 39, 221 – 230.
- E. Keenan & D. Westerståhl, 1997, ‘Quantifiers’, in J. van Benthem & A. ter Meulen, eds., *Handbook of Logic and Language*, Elsevier, Amsterdam.
- M. Knauff, 2007, ‘How our Brains Reason Logically’, Center for Cognitive Science University of Freiburg. To appear in H. Hodges, W. Hodges & J. van Benthem, eds., Special Issue of *Topoi* on “Logic and Psychology”.
- F. Liu & J. Zhang, 2007, ‘Some Thoughts on Mohist Logic’, Chinese Academy of Social Science, Beijing, and ILLC, University of Amsterdam.
- Ch. Manning & B. MacCartney, 2007, ‘Natural Logic for Textual Inference’, Departments of Linguistics/Computer Science, Stanford University. To appear at the Workshop on Textual Entailment and Paraphrasing, *ACL 2007*.
- L. Moss, 2005, ‘Completeness in Natural Logic: What and How’, LSA Linguistics Summer Institute, Harvard University.

- R. Nairn, C. Condoravdi & L. Karttunen, 2006, 'Computing Relative Polarity for Textual Inference', Palo Alto Research Center, in *Proceedings Inference in Computational Semantics*.
- S. Peters & D. Westerståhl, 2006, *Quantifiers in Language and Logic*, Oxford University Press, Oxford.
- I. Pratt, 2004, 'Fragments of Language', *Journal of Logic, Language and Information* 13:2, 207 – 223.
- W. Purdy, 1991, 'A Logic for Natural Language', *Notre Dame Journal of Formal Logic* 32:3, 409–425.
- V. Sanchez Valencia, 1991, *Studies on Natural Logic and Categorical Grammar*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.
- V. Sanchez Valencia, 2007, 'The Algebra of Logic', to appear in J. Woods, ed. *The Handbook of the History of Logic*, Elsevier, Amsterdam.
- F. Sommers, 1982, *The Logic of Natural Language*, Clarendon Press, Oxford.