A Meaning-Relative Logical Consequence Relation?

Tibo Rushbrooke

November 2021

Introduction

- When we ask whether some English sentence is a logical consequence of some other sentences, does the answer depend on the meanings of those sentences?

Introduction

- When we ask whether some English sentence is a logical consequence of some other sentences, does the answer depend on the meanings of those sentences?
- To some extent, it must. ‘It’s raining or it’s snowing’ would not be a logical consequence of ‘it’s snowing’ if ‘or’ meant what ‘and’ usually means
- So, the standard answer is that the logical consequence relation of a natural language, like English, depends only on the meanings of the ‘logical constants’

The Standard View

This standard view is widely held to have originated with Tarski:

- "The consequence relation cannot be affected by replacing the designations of the objects referred to in these sentences by the designations of any other objects" (1936, p.415)

It has since been endorsed by many authors. Here is Quine:

- Quine: "A logical truth is a statement which is true and remains true under all reinterpretations of its components other than the logical particles" (1951, p.23)

The Standard View

And here is Lewis:

- "Our syntactic surrogate had better not yield narrowly logical consistency...- that is, consistency under some reinterpretation or other of all but the logical vocabulary... That would falsify the facts of modality by yielding allegedly consistent ersatz worlds according to which there are unmarried bachelors. . ." (1986,p.152-3)

So 'There are unmarried bachelors' is logically consistent according to Lewis.

The Standard View

- By '(re)interpretation', Quine and Lewis apparently mean something which behaves, for a natural language, much like a first-order model for a first-order language- assigning objects to terms, extensions to predicates, etc.
- I will call such an entity a *formal interpretation*
- Summary of Tarskian view: logical consequence, $\models$, in a natural language is preservation of truth under all formal interpretations for that language
- Such a relation cannot be meaning-relative as terms have different 'meanings' under different formal interpretations

Formal interpretation $\neq$ meaning

- A formal interpretation for a class of formal sentences fixes the *truth value* of those sentences
- Intuitively, a semantics for a class of natural language sentences seems not to fix the truth value of those sentences
  - The truth of 'Brutus killed Caesar' seems to depend not just on the meanings of the terms, but also on the world (Quine, 1986)
- A formal interpretation seems to perform two quite different functions: *assigning a referent* to each term, and *fixing the truth value* of each sentence

Formal interpretation $\neq$ meaning

- A formal interpretation for a class of formal sentences fixes the *truth value* of those sentences
- Intuitively, a semantics for a class of natural language sentences seems not to fix the truth value of those sentences
  - The truth of 'Brutus killed Caesar' seems to depend not just on the meanings of the terms, but also on the world (Quine, 1986)
- A formal interpretation seems to perform two quite different functions: *assigning a referent* to each term, and *fixing the truth value* of each sentence
- What if we were to separate these functions/ conceive of them as being performed by different kinds of thing?

Interpretations

To perform the first function, we introduce *interpretations*.

Interpretations

To perform the first function, we introduce *interpretations*.

- An *interpretation* of English is a function which maps each English term to an object, a relation, or a logical operation. The notion generalises to other natural languages in the obvious way.
- For example, the interpretation corresponding to the usual semantics of English might map 'Big Ben' to a clock tower in London, 'chair' to a 1-place relation (property), and 'and' to the operation $\wedge$ of conjunction.
- The notion is intended to model natural language meaning. But it is a MASSIVE simplification!! Ignores adverbs, context-sensitivity, ambiguity, vagueness, ...

Propositions

Propositions are given by familiar-looking recursive clauses...

- If R is an n-place relation, $a_1, \ldots, a_n$ are objects, then there is a proposition $Ra_1 \ldots a_n$ which 'says that' the relation R holds between $a_1, \ldots, a_n$.
- If ϕ, ψ are propositions then there are propositions $\phi \wedge \psi, \phi \vee \psi, \neg\phi$
- (Tentative) If $\phi(x)$ is a unary propositional function, then $\forall x \phi(x)$ and $\exists x \phi(x)$ are propositions

Propositions

Propositions are given by familiar-looking recursive clauses...

- If R is an n-place relation, $a_1, \ldots, a_n$ are objects, then there is a proposition $Ra_1 \ldots a_n$ which 'says that' the relation R holds between $a_1, \ldots, a_n$.
- If ϕ, ψ are propositions then there are propositions $\phi \wedge \psi, \phi \vee \psi, \neg\phi$
- (Tentative) If $\phi(x)$ is a unary propositional function, then $\forall x \phi(x)$ and $\exists x \phi(x)$ are propositions

What's a propositional function??? Intuitively: take a proposition, then remove one or more occurrences of a particular object.

Propositions

- Propositions are intended to be the interpretations of natural language sentences.
- For example, let $\mathcal{I}$ be the 'usual' interpretation of English. Let b be the object Big Ben, C the property of being a chair.
- Then $\mathcal{I}$('Big Ben is a chair'), that is, the image of this sentence under the 'usual' interpretation, might be the proposition Cb.

Valuations

To perform the second function, we introduce *valuations*.

Fix a set D of objects, and a set Rel of relations. A *valuation* is a function, V, which assigns to each n-ary relation a set of n-tuples of D: $V(R) \subseteq D^n$.
A valuation V determines a map $\mathbb{V}$ assigning each proposition a truth value, as follows:

- $\mathbb{V}(Ra_1 \dots a_n) = 1$ if $\langle a_1, \dots, a_n \rangle \in V(R)$, 0 otherwise
- $\mathbb{V}(\phi \wedge \psi) = 1$ iff $\mathbb{V}(\phi) = 1$ and $\mathbb{V}(\psi) = 1, \mathbb{V}(\phi \vee \psi) = 1$ iff $\mathbb{V}(\phi) = 1$ or $\mathbb{V}(\psi) = 1, \mathbb{V}(\neg\phi) = 1$ iff $\mathbb{V}(\phi) = 0$
- $\mathbb{V}(\forall x \phi(x)) = 1$ iff $\mathbb{V}(\phi(a)) = 1$ for every object $a \in D$, $\mathbb{V}(\exists x \phi(x)) = 1$ iff $\mathbb{V}(\phi(a)) = 1$ for some object $a \in D$

Consequence

- We may now define a meaning-relative consequence relation for English.

Consequence

- We may now define a meaning-relative consequence relation for English.
- Let Γ be a set of English sentences, ϕ an English sentence, I an interpretation of English
- Say that Γ entails ϕ relative to I, $\Gamma \models_I \phi$, iff for every valuation V, if $\mathbb{V}(I(\gamma)) = 1$ for every $\gamma \in \Gamma$, then $\mathbb{V}(I(\phi)) = 1$.
- Say that ϕ is *valid* relative to I, $\models_I \phi$, iff $\emptyset \models_I \phi$

New inferences

- Now for the fun part! We shall see that the relation $\models_{\mathcal{I}}$ licenses a wider range of inferences than the classical consequence relation.
- Let the object c be the Roman orator Cicero, assume 'Tully' names this same object. Let W be the property of wisdom.
- 'Cicero is wise' $\not\models$ 'Tully is wise', because a formal interpretation can be found within which 'Cicero' refers to an individual in the extension of 'wise', while 'Tully' refers to a *different* individual, not in the extension of 'wise'
- However, $\mathcal{I}$('Cicero') $=$ $\mathcal{I}$('Tully') and so $\mathcal{I}$('Cicero is wise') and $\mathcal{I}$('Tully is wise') are the very same proposition, Wc. Then $\mathbb{V}(\mathcal{I}$('Cicero is wise'$)) = 1$ implies $\mathbb{V}(\mathcal{I}$('Tully is wise'$)) = 1$, so 'Cicero is wise' $\models_{\mathcal{I}}$ 'Tully is wise'.

New inferences

- Another example- assume a particular context where 'That stuff' refers to some stuff a. Let H be the chemical property of being H_2O.
- 'That stuff is water' $\models$ 'That stuff is H_2O' does not hold, as Glanzberg (2015) points out
- However, assuming $\mathcal{I}$('water') is also the property of being H_2O, $\mathcal{I}$('That stuff is water') $= \mathcal{I}$('That stuff is H_2O) $= Ha$
- So we have 'That stuff is water' $\models_{\mathcal{I}}$ 'That stuff is H_2O'

New inferences

Let's try cracking a tougher nut.

- Does the inference from 'Vasily is a bachelor' to 'Vasily is not married' come out logically valid, according to $\models_{\mathcal{I}}$?
- Intuitively it should. $\mathcal{I}$('bachelor') should be a property somehow logically connected to $\mathcal{I}$('married')...

New inferences

Let's try cracking a tougher nut.

- Does the inference from 'Vasily is a bachelor' to 'Vasily is not married' come out logically valid, according to $\models_{\mathcal{I}}$?
- Intuitively it should. $\mathcal{I}$('bachelor') should be a property somehow logically connected to $\mathcal{I}$('married')...
- We need some more formal machinery!

Complex properties

- We allow new properties to be constructed from old
- If θ, ρ are properties then $[\theta \wedge \rho], [\theta \vee \rho], [\neg\theta]$ are properties
- We extend the definition of a valuation $\mathbb{V}$:
  $\mathbb{V}([\theta \wedge \rho]a) = \mathbb{V}(\theta a \wedge \rho a)$, ..., $\mathbb{V}([\neg\theta]a) = \mathbb{V}(\neg\theta a)$
- Intuition: given properties Red, Car, we can build the property: Red car (Red-and-Car)

Back to bachelor-ness

- Let M be the property of being married, Y the property of being a man. Then reasonable to say $\mathcal{I}$ maps 'bachelor' to the complex property $[\neg M \wedge Y]$
- This gives $\mathcal{I}$('Vasily is a bachelor') $= [\neg M \wedge Y]v$, $\mathcal{I}$('Vasily is not married') $= \neg Mv$
- Suppose $\mathbb{V}([\neg M \wedge Y]v) = 1$. Then $\mathbb{V}(\neg Mv \wedge Yv) = 1$, implying $\mathbb{V}(\neg Mv) = 1$
- We can conclude $\mathbb{V}(\mathcal{I}($'Vasily is a bachelor'$)) = 1$ implies $\mathbb{V}(\mathcal{I}($'Vasily is not married'$)) = 1$
- So 'Vasily is a bachelor' $\models_{\mathcal{I}}$ 'Vasily is not married'.

A final example

- Assume ‘R’ refers to some relation, R.
- Certainly ‘R is Euclidean’ $\not\models$ ‘R is directed’
- However, I claim that ‘R is Euclidean’ $\models_{\mathcal{I}}$ ‘R is directed’.

A final example

- Assume ‘R’ refers to some relation, R.
- Certainly ‘R is Euclidean’ $\not\models$ ‘R is directed’
- However, I claim that ‘R is Euclidean’ $\models_{\mathcal{I}}$ ‘R is directed’.

Proof: Exercise!

Conclusion

- The relations $\models$ and $\models_{\mathcal{I}}$ capture two different notions of logical consequence, offering different perspectives on logic
- From the new perspective, the primary objects of logical study are *non-linguistic*- viz. interpreted sentences, or propositions
- $\models_{\mathcal{I}}$ allows a richer array of inferences to be treated with logical methods. Perhaps this is exciting for us logicians.

Thanks for Listening!

References

- Tarski, A. (1936), 'O pojciu wynikania logicznego', *Przeglad Filozoficzny* 39:58-68, trans. J.H. Woodger in A. Tarski, *Logic, Semantics and Metamathematics*, ed. J. Corcoran (Indianapolis: Hackett)
- Quine, W.V. (1951), 'Two Dogmas of Empiricism', *The Philosophical Review*, vol. 60 (1), pp.20-43
- Lewis, D.K. (1986), *On the Plurality of Worlds* (Oxford: Blackwell)
- Quine, W.V. (1986), *Philosophy of Logic*, 2nd ed. (Harvard University Press)
- Glanzberg, M. (2015), 'Logical Consequence and Natural Language', in C. Caret and O. Hjortland, eds., *Foundations of Logical Consequence* (Oxford University Press), p.71-120