Chapter Seven

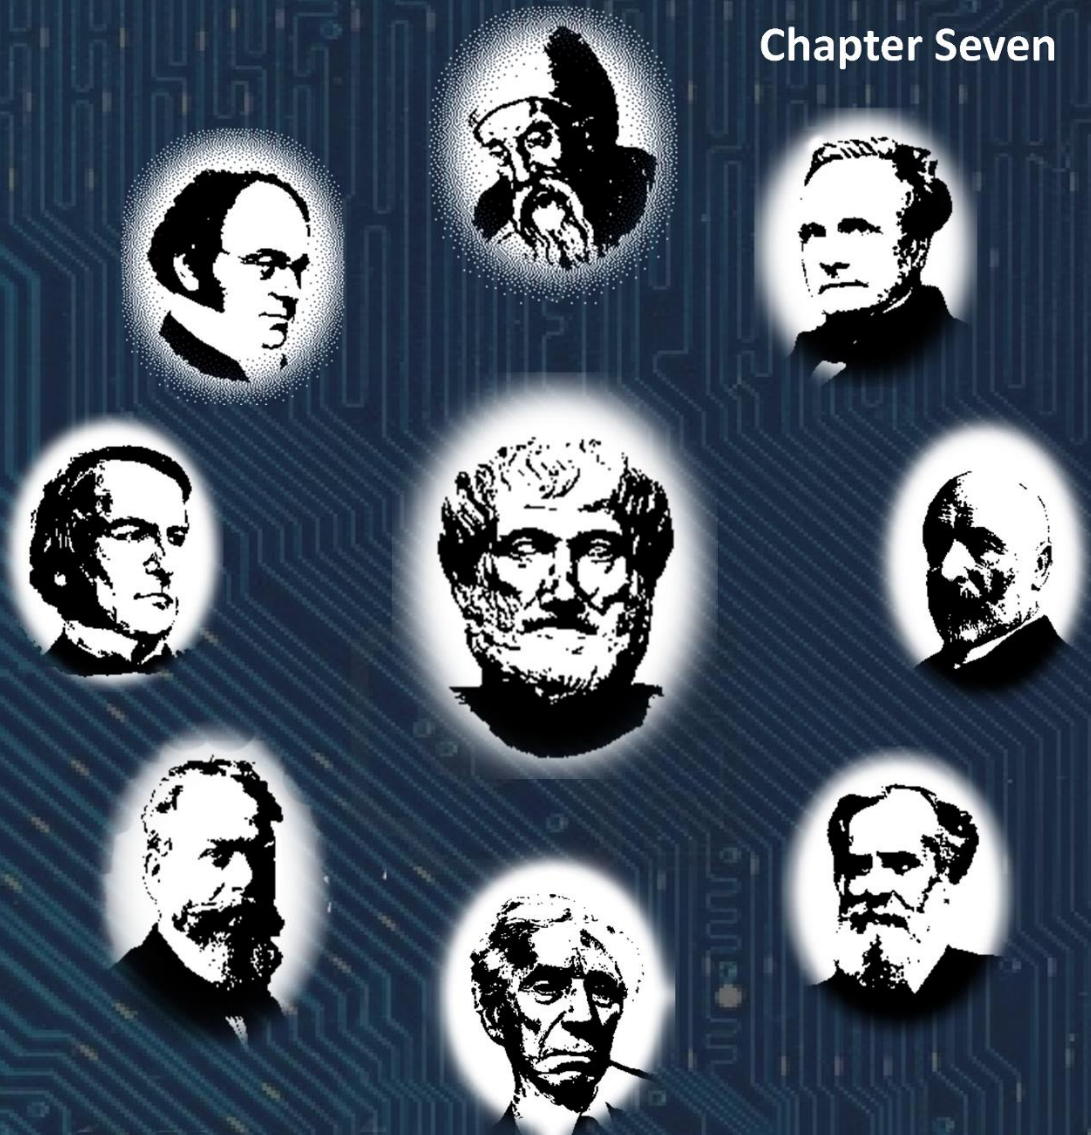

REASONING

elaine rich

alan kaylor cline

The Logicians on our cover are:

Euclid (? - ?)

Augustus De Morgan (1806 – 1871)

Charles Babbage (1791 – 1871)

George Boole (1815 – 1864)

Aristotle (384 BCE – 322 BCE)

George Cantor (1845 – 1918)

Gottlob Frege (1848 – 1925)

John Venn (1834 – 1923)

Bertand Russell (1872 – 1970)

REASONING

AN INTRODUCTION TO LOGIC, SETS, AND FUNCTIONS

CHAPTER 7
ENGLISH INTO LOGIC: ISSUES AND SOLUTIONS

Elaine Rich
Alan Kaylor Cline

The University of Texas at Austin

Image credits:

Girl with cat: © Lynda Trader, 2007.

Bear with rifle: © Lynda Trader, 2007

Potatoes, Rice, Beans:

Balloons: http://www.balloondealer.com/app/images/Balloons21.jpg

Male bird song speeds: http://sites.sinauer.com/animalcommunication2e/chapter08.04.html

White peacock: © David Rich, 2014.

http://www.cs.utexas.edu/learnlogic
Library of Congress Cataloging-in-Publication Data

Rich, Elaine, 1950 -
Reasoning—An Introduction to Logic Sets and Functions / Elaine Rich.— 1st ed. p. cm.
ISBN x-xxx-xxxxx-x 1

Table of Contents

English Into Logic: Issues and Solutions

Getting Off the Ground

Introduction

When we try to translate English sentences into logical statements, we may get stuck as we try to figure out exactly what the English sentences actually mean.

We could write volumes on the subject of English ***semantics***: the process by which meaning is assigned to utterances. We won't do that here. What we will do is just to sketch a few issues so that we're aware, as we use English to describe the reasoning that we'll do, of possible misunderstandings or confusion.

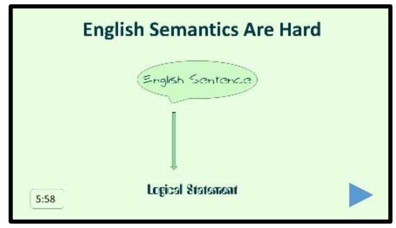

https://www.youtube.com/watch?v=jIGOpFOyFcc

And, by the way, English isn't special in this regard. All natural human languages (Spanish, Urdu, Khmer, or whatever) have the problems that we are describing here, although the details differ.

> ***Big Idea***
>
> English sentences can be ambiguous, vague, unclear, and sometimes even downright misleading.
>
> The language of logic helps us avoid those problems.
>
> But, to use it to express our ideas, we need to understand how it relates to our natural language, English.

Problems

1. We might think that it should be very straightforward to translate into logic claims about particular named individuals. For example, assuming that it is clear which person named Logan we are talking about, we can write:

Student(*Logan*)

But even this is not always so simple. Suppose that we want to encode into logic the sentence:

There is a tooth fairy.

Which one or more of the following expressions is a syntactically legal predicate logic statement (or wff), that may reasonably encode our claim if we have defined all the predicates appropriately:

I. ∃*x* (*ToothFairy*(*x*))
II. ∃ (*ToothFairy*)
III. ∃ *Tooth Fairy*
IV. ∃*x* (*BringsMoneyForTeeth*(*x*))

2. Suppose that we want to encode into logic the sentence:

Santa Claus does not exist.

Which one or more of the following expressions is a syntactically legal predicate logic statement (or wff), that may reasonably encode our claim if we have defined all the predicates appropriately

I. ¬∃ (*SantaClaus*)
II. ¬∃ *SantaClaus*
III. ¬∃*x* (*SantaClaus*(*x*))
IV. ¬∃*x* (*WearsRedSuit*(*x*) ∧ *LivesAtNorthPole*(*x*) ∧ *ClimbsDownChimneys*(*x*) ∧ *BringsXmasPresents*(*x*))

Big Idea

In many cases, if we try to map from English sentences to logical ones without clearly analyzing what the English sentence is saying, we will have trouble. It may be tempting to jump to the conclusion that we have found a weakness in our logical system itself. While there are real limits to the logical system that we are discussing, we should, before assuming that we've found one, see if a more careful translation of English into logic can't make the problem disappear.

We Must Overcome the Perils of English - Ambiguity

Introduction

English has evolved to be an efficient way for people to communicate with each other. The price we pay for that is that many English sentences are ambiguous: they have two or more meanings. Often, we don't even notice most of those meanings since they're nonsense. But logical statements are unambiguous. So, before we can write one, we have to know exactly what we are trying to say.

Structural Ambiguity

A very common source of ambiguity is structural: There may be more than one syntactic parse (structure) for a sentence. Each parse corresponds to a different meaning.

We've already looked at a lot of examples in which ambiguity arises from the lack of parentheses in English.

Recall that: pancakes or bacon and eggs

could mean:

(pancakes or bacon) and eggs

OR

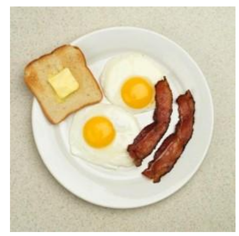

or it could mean:

pancakes or (bacon and eggs)

OR

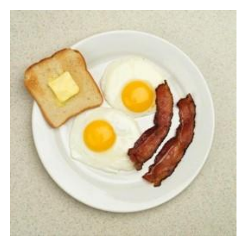

A second source of structural ambiguity is postmodifiers (modifiers that come after whatever they modify). The problem is that those modifiers have to get attached to the constituents they modify. There may be more than one way to do that. Prepositional phrases often cause this kind of ambiguity.

Consider: Chris likes the girl with a cat.

Anyone who reads this sentence is going to assume that the prepositional phrase, "with a cat" modifies, "the girl". There's a girl with a cat and Chris likes her. So we get:

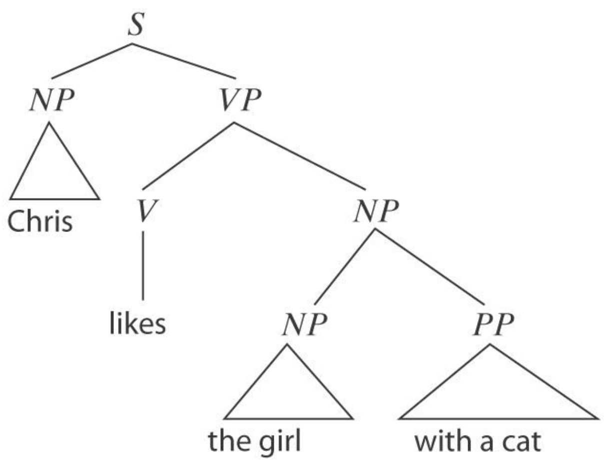

But now, suppose that we replace "like" with "shot", "girl" with "bear", and "cat" with "rifle".

We get: Chris shot the bear with a rifle.

Should we get the same structure?

If we parse it the same way, we get:

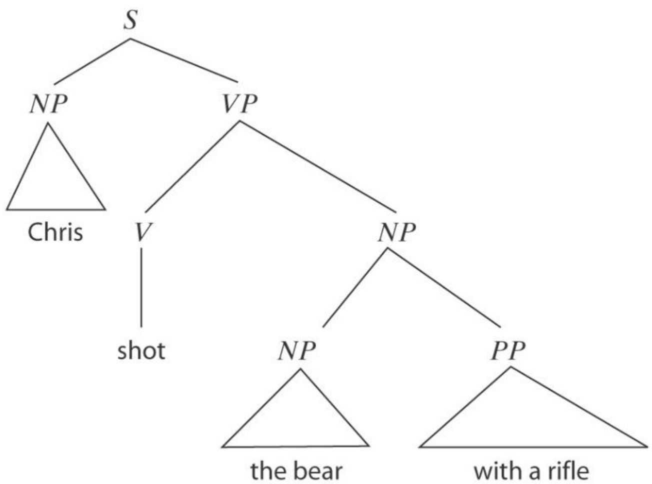

Cute. But almost surely not what was meant. Instead, in this case, we want to attach the prepositional phrase, "with a rifle," not to "the bear," but to the verb phrase, "shot the bear". The parse tree we need is this one:

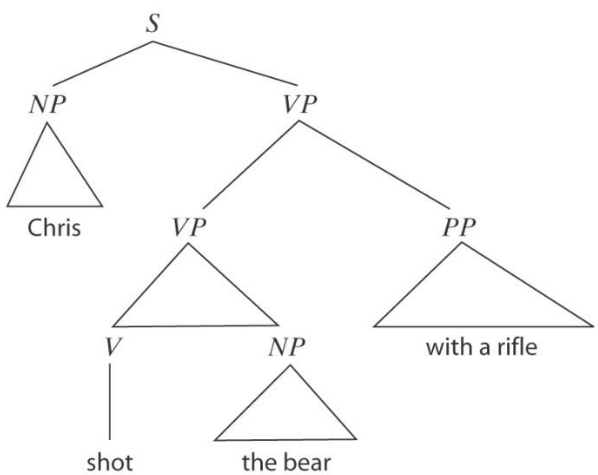

Now the prepositional phrase (as is common when such phrases modify verb phrases) says something about the way in which the action (shooting) happened.

English syntax would have allowed us to produce a parse tree like this for the first sentence as well. In other words, we could have attached, "with a cat," to the verb, "likes", but that wouldn't have made sense – you can't like using a cat.

So the problem is that the syntax (grammar) of English allows prepositional phrases to be attached in many different ways. In the case of most sentences, only one of the possible attachments makes sense. We (people) jump directly to it. But we must be careful, when we assert that the meaning of a given sentence is some particular logical claim, that we are clear on which parse we are assigning that meaning to.

Problems

1. Consider: You can have potatoes or beans and rice.

Which of the following possible parenthesizations is almost certainly the one that is intended:

a) potatoes OR (beans AND rice)

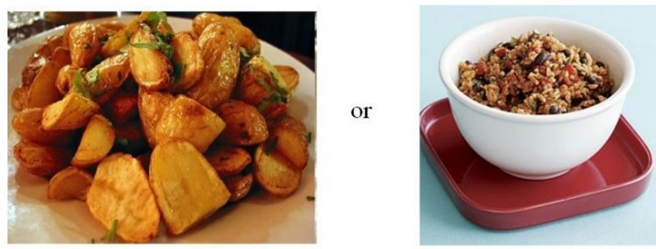

b) (potatoes OR beans) AND rice

2. Consider: Kim ate the fried chicken with a fork.

The prepositional phrase, “with a fork,” modifies:

a) The direct object (“the fried chicken”)
b) The verb phrase (“ate the fried chicken”)

3. Consider: Tracy likes books with surprise endings.

The prepositional phrase, “with surprise endings,” modifies:

a) The direct object (“books”)
b) The verb phrase (“likes books”)

More Structural Ambiguity

Other syntactic structures can also lead to ambiguity.

Suppose we have:

[1] Chris is crazy about Blake.
[2] Chris sort of likes Tracy.
[3] Tracy is far and away Blake's biggest fan.

Now consider:

[4] Chris likes Blake better than Tracy.

Is [4] true (given what we already know)? We can't answer this simple question because [4] is ambiguous. Which of these does it mean:

[4′] Chris likes Blake better than Chris likes Tracy.
[4″] Chris likes Blake better than Tracy likes Blake.

Based on what we know from [1] – [3], [4′] is true. But [4″] is false.

Problems

1. Consider: Jia speaks Chinese more fluently than French.

Possible meanings for this sentence have the form:

Jia speaks Chinese more fluently than (A)____________ speaks (B)_______________.

Write down all the meanings that aren't nonsense (in our everyday world). In other words, write down all the combinations of ways to fill in the two blanks.

2. Consider: Jia speaks Chinese more fluently than Jean.

Possible meanings for this sentence have the form:

Jia speaks Chinese more fluently than (A)____________ speaks (B)_______________.

Write down all the meanings that aren't nonsense (in our everyday world). In other words, write down all the combinations of ways to fill in the two blanks.

3. Consider: Jia knows Harper better than Morgan.

Possible meanings for this sentence have the form:

Jia knows Harper better than (A)____________ knows (B)_______________.

Write down all the meanings that aren't nonsense (in our everyday world). In other words, write down all the combinations of ways to fill in the two blanks.

Logical Ambiguity

Sometimes ambiguity arises more directly from the logic of the claim that is being made. This kind of ambiguity often happens when the English sentence contains some sort of negative. Often, in such statements, the issue is the scope of the *not*.

Consider: I didn't take both of my pills this morning.

This could mean either of these things:

[1] I may have taken one of my pills this morning, but not both of them.
[2] For each of my pills, it's the case that I didn't take it this morning.

Or, writing the claims in our logical language:

[1] $\neg(\forall x\ (pill(x) \rightarrow take(x)))$
[2] $\forall x\ (\neg(pill(x) \wedge take(x)))$

Consider: Mo doesn't want Jo and Bo to come to the party.

This could mean either of these things:

[1] Mo doesn't want either Jo or Bo to come to the party.
[2] Mo doesn't want it to happen that both of them come (maybe that would make things too exciting.)

We can write (very simple versions of) these meanings in Boolean logic if we define:

J: Mo wants Jo to come to the party.
B: Mo wants Bo to come to the party.

[1] $\neg J \wedge \neg B$
[2] $\neg(J \wedge B)$

Problems

1. Consider: Smoking is not permitted on all Korean Air flights.

We want to encode the meaning of this sentence. Define:

KAf(x) : True if x is a Korean Air flight.
SP(x): True if smoking is permitted in or on x.

Shown here five logical expressions. Mark the two that correspond to the two reasonable interpretations of this sentence. (Hint: first write out the two interpretations in unambiguous English. Then translate each of those into logic.)

a) $\exists x\ (\neg(KAf(x) \lor SP(x)))$
b) $\forall x\ (KAf(x) \rightarrow \neg SP(x))$
c) $\forall x\ (\neg KAf(x) \land \neg SP(x))$
d) $\exists x\ (KAf(x) \land \neg SP(x))$
e) $\neg\exists x\ (KAf(x) \rightarrow \neg SP(x))$

2. Consider: Cruz forgot to invite the whole class.

We want to encode the meaning of this sentence. So that we can focus on the issue of its ambiguity, we'll simplify the problem. So we'll ignore the issue of why Cruz didn't invite people. We'll just describe who was invited. Define:

Inclass(x): True if x is in the class.
Invited(x, y): True if x invited y.

Shown here are five logical expressions. Mark the two that correspond to the two reasonable interpretations of this sentence. (Hint: First write out the two interpretations in unambiguous English. Next, translate each of those into logic. Then compare your results to the ones listed below. Be careful: It is possible that one of these is not identical to what you have written but is logically equivalent to something you have written. In that case, mark it as one of the reasonable ones.)

a) $\forall x\ (Inclass(x) \rightarrow \neg Invited(\text{Cruz}, x))$
b) $\neg\forall x\ (Inclass(x) \land Invited(\text{Cruz}, x))$
c) $\exists x\ (Inclass(x) \lor \neg Invited(\text{Cruz}, x))$
d) $\exists x\ (Inclass(x) \land \neg Invited(\text{Cruz}, x))$
e) $\forall x\ (\neg Inclass(x) \rightarrow \neg Invited(\text{Cruz}, x))$

3. Consider: We couldn't do all of these shows without the staff.

We want to encode the meaning of this sentence. So that we can focus on the issue of its ambiguity, we'll simplify a bit and define these predicates:

Staff: True if the staff exists.
Show(x): True if x is a show.
Possible(x): True if we can do x.

Shown here are five logical expressions. Mark the two that correspond to the two reasonable interpretations of this sentence. (Hint: first write out the two interpretations in unambiguous English. Then translate each of those into logic.)

a) $\neg Staff \rightarrow (\forall x\ (Show(x) \vee \neg Possible(x)))$

b) $\neg Staff \rightarrow (\forall x\ (Show(x) \rightarrow \neg Possible(x)))$

c) $\neg Staff \rightarrow (\exists x\ (Show(x) \wedge \neg Possible(x)))$

d) $(\forall x\ (Show(x) \vee \neg Possible(x))) \rightarrow Staff$

e) $(\forall x\ (Show(x) \rightarrow \neg Possible(x))) \rightarrow \neg Staff$

Referential Ambiguity - Pronouns

The job of noun phrases is to refer to things. Ambiguity arises when we can't tell which things.

Pronouns are often a problem. To decide on the meaning of a sentence that contains a pronoun, we have to determine the pronoun's referent (the thing to which it refers). Sometimes that's straightforward. Sometimes it isn't.

Here's an easy case: Jen saw Bill at the movies. She went over to talk to him.

We'll assume that there isn't any larger, complicating context and that names are being used conventionally. Then "she" must refer to Jen and "him" must refer to Bill.

But there are harder cases: The dog spotted the cat on the lawn. It ran away.

Now it's possible that either the dog or the cat decided to run.

Problems

1. Consider:

[1] Crystal told Sherry that her sister wanted to meet her for lunch.

Now consider:

I. Crystal's sister wanted to meet Crystal for lunch.
II. Crystal's sister wanted to meet Sherry for lunch.
III. Sherry's sister wanted to meet Sherry for lunch.
IV. Sherry's sister wanted to meet Crystal for lunch.

(Part 1) How many of these is/are possible paraphrase(s) for what Crystal told Sherry?

(Part 2) Why is [1] ambiguous?

2. A ***Winograd schema*** (named after the linguist/philosopher Terry Winograd) is a sentence that mentions two objects and contains a pronoun that could refer back to either of them. We then ask a person (who can use world knowledge) to decide on the referent of the pronoun. In a properly designed Winograd schema, it is possible, by changing a single word in the sentence, to change the answer that a person will give.

Consider: The toaster won't fit in the box because it is too big.

(Part 1) "it" refers to:

a) the toaster b) the box

(Part 2) Which of the following words would, if changed, change your answer to Part 1:

a) toaster b) box c) big

3. A ***Winograd schema*** is a sentence that mentions two objects and contains a pronoun that could refer back to either of them. We then ask a person (who can use world knowledge) to decide on the referent of the pronoun. In a properly designed Winograd schema, it is possible, by changing a single word in the sentence, to change the answer that a person will give.

Consider: Maria made cookies for the students. They were delicious.

(Part 1) "they" refers to:

a) the cookies b) the students

(Part 2) Which of the following words, if inserted in place of "delicious", would change your answer to Part 1:

a) hungry b) oatmeal c) expensive

Referential Ambiguity – The Definite Article

Pronouns, however, are not the only source of referential ambiguity.

The definite article "the" introduces a description of an object that is assumed (in the current context) to be unique. If it's not, the resulting sentence may have multiple meanings. And, if there is no possible referent, the sentence will typically be meaningless. All this matters if we want to translate the sentence into a logical expression and then evaluate the truth value of that expression.

The purple balloon popped.	No problem.
The green balloon popped.	Which one?
The black balloon popped.	Huh?

(As an aside, we'll point out that we can also use "the" when the referent, while not unique, is unique enough for the current purpose. For example, we can say, "Cam went to the store." There may be many stores around, but we say this when it doesn't matter which one Cam went to.)

Problems

1. Before we can assign a truth value to an assertion, it must be clear what objects are being discussed. In the case of sentences that contain the definite article “the”, that will only be the case if there is exactly one reasonable referent for the noun phrase that “the” introduces.

In each of these examples, assume an arbitrary, reasonable context. Indicate the number of possible referents for the underlined noun phrase.

(Part 1) <u>The state flower of Texas</u> is the bluebonnet.

a) None.
b) Exactly one.
c) There could easily be more than one.

(Part 2) <u>The square root of Paris</u> is 2.

a) None.
b) Exactly one.
c) There could easily be more than one.

(Part 3) <u>The magazine on the table</u> has a great article about breakfast tacos.

a) None.
b) Exactly one.
c) There could easily be more than one.

2. It’s not possible to assign a truth value to an English sentence until it has been disambiguated. If two people agree to accept as a premise an English sentence that is ambiguous, they may find out later that the arrive at different conclusions (because they have, without realizing it, accepted different logical premises)

For each of the following sentences, indicate whether this could be a problem because there is not a unique referent of the underlined noun phrase:

a) <u>The square of every even number</u> is even.

b) <u>The best actor on the planet</u> is an American.

c) <u>The capital of the happiest country in the world</u> is Canberra.

Referential Ambiguity – Rhetorical Devices

The problem of assigning referents is further complicated by the use of various kinds of rhetorical devices (figures of speech), such as metonymy.

Is this statement true, false, or nonsense:

Michigan is in Nebraska now.

Taken literally, it is false. The state of Michigan is not contained in the state of Nebraska.

But this statement does have a sensible meaning that could, in some circumstances be true: Something (very likely a sports team) associated with the state of Michigan is now in Nebraska. In fact, a newscaster said it when the University of Michigan's solar car entered the state of Nebraska on its way from Austin to Minnesota, where it won the 2014 solar car challenge.

Problems

1. For each of the following sentences mark the statement that is true of it:

(Part 1) The White House just announced a new policy on greenhouse gases.

a) The literal meaning could be true.
b) The literal meaning is false or nonsense but there is a meaning that exploits a figure of speech and that could be true.
c) It's hard to find any meaning that isn't nonsense.

(Part 2) Wall Street overreacted to the latest unemployment numbers.

a) The literal meaning could be true.
b) The literal meaning is false or nonsense but there is a meaning that exploits a figure of speech and that could be true.
c) It's hard to find any meaning that isn't nonsense.

Situated Truth

Whenever we say or write something, we do so in some particular situation. A claim can be true in one situation, but not in another.

[1] I live in Texas. (Depending on who is speaking, and thus what the referent of "I" is, this sentence is either true or false.)

[2] The principal is from Kansas. (When we use the word, "the", we presuppose that there exists some unique object to which we're referring. Clearly there are many, many principals in the world. But said in the context of a particular school, this sentence refers to one person and is either true or false.)

[3] There's an ice cream store down the street. (True in some places and false in others.)

[4] Tomorrow is a holiday. (True on some days and in some places, and false in others.)

We will need to assume, in our discussions, that we begin with claims for which any situational ambiguity (including whether we're talking about the real world or some imaginary one) has already been resolved.

Problems

1. Consider the claim, “A Presidential election is held every four years.”

Consider the following statements about this claim:

I. To assign it a truth value requires additional information about the time at which it was asserted.
II. To assign it a truth value requires additional information about the place in which it was asserted.
III. To assign it a truth value requires additional information about the entities (such as people or things) in the situation where it was asserted.

Which of them is/are true (don’t stretch super far to find some exceptional case):

2. Consider the claim, “Kerry likes ice cream.”

Consider the following statements about this claim:

I. To assign it a truth value requires additional information about the time at which it was asserted.
II. To assign it a truth value requires additional information about the place in which it was asserted.
III. To assign it a truth value requires additional information about the entities (such as people or things) in the situation where it was asserted.

Which of them is/are true (don’t stretch super far to find some exceptional case):

3. Consider the claim, “5 is a prime number”.

Consider the following statements about this claim:

I. To assign it a truth value requires additional information about the time at which it was asserted.
II. To assign it a truth value requires additional information about the place in which it was asserted.
III. To assign it a truth value requires additional information about the entities (such as people or things) in the situation where it was asserted.

Which of them is/are true (don’t stretch super far to find some exceptional case):

Ambiguity of Some Logical Operators – NEGATION

In logic, for any predicate *P*, we can write both:

[1] $P(x)$
[2] $\neg P(x)$

Given some appropriate definition of *P*, The Law of the Excluded Middle tells us that, for any *x* in *P*'s domain, one of these claims must be true. And, if we write many simple claims in English, it will also be the case that either the claim or its apparent negation must be true.

So, for example, we could write:

[1]	*French*(*Jean*)	Jean is French.
[2]	¬*French*(*Jean*)	Jean isn't French.

(with respect to whatever definition of the predicate *French* we have provided).

But this doesn't always happen, particularly when we allow for the variety of ways of expressing negation in English.

Consider:

[1] Travis likes Jody.
[2] Travis doesn't like Jody. (an apparently simple negation of [1])

Must one of those sentences be true? It depends on what [2] means. It could mean either:

[3] It is not the case that Travis likes Jody.
[4] Travis dislikes Jody.

If it means [3], then yes, one of [1] and [3] must be true. We can express this as a logical claim:

Likes(*Travis*, *Jody*) ∨ ¬*Likes*(*Travis*, *Jody*)

But if it means [4], no, it is not necessarily the case that one of [1] and [4] must be true. It is possible that Travis knows nothing about Jody. Or perhaps Travis is neutral about Jody. Then both [1] and [4] are false. Travis neither likes nor dislikes Jody.

Have we violated the Law of the Excluded Middle in a case like this? No. "Dislikes" is not equivalent to "not likes".

The problem is that many of the more natural tools that English gives us for expressing negations (including "doesn't", and prefixes such as "dis-", "un-", "il"-, and so forth) are almost like logical negation, but not quite. Often this happens because the property that is being described can best be thought of as describing, not a dichotomy, but rather points on some spectrum. In this case, we often define a logical predicate that corresponds to some part (typically one end) of that spectrum. Then its negation must cover all of the rest. Yet related negative words in English often apply only at the other end.

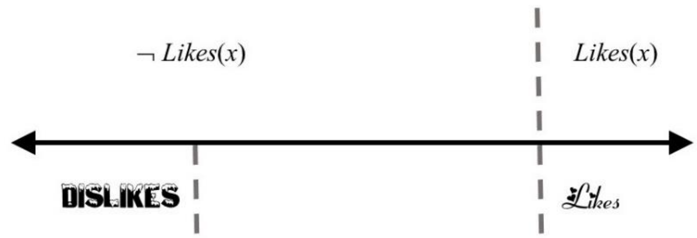

Some adjectives correspond to dichotomies, but many don't.

"Legal" does correspond to a dichotomy:

Illegal	is equivalent to	¬*Legal*

Anything in the domain of *Legal* must either be legal or not legal (illegal).

"Likely" doesn't correspond to a dichotomy:

Unlikely	is not equivalent to	¬*Likely*

Something with probability close to 50/50 is neither likely nor unlikely.

Later, we'll have more to say about the difficulty of dichotomizing our analog world and the vague words we tend to use when we attempt to do that.

Big Idea

The Law of the Excluded Middle says that, for any statement *P*, *P* or ¬*P* is true. It doesn't say that *P* or some Q that may seem like not *P* must be true. If Q is not exactly ¬*P*, then it's possible that there are more than two alternatives. Failure to recognize that can lead to the logical error called a ***False Dichotomy***.

A classic example of this is the expression, "If you're not with us, you're against us." The phrase "against us" is not exactly equivalent to "not with us". It's possible to be neutral. Unfortunately, people who are trying to make a point may say things like this and hope that their listeners will be fooled into believing their argument.

Problems

1. Let's say that a pair of claims C_1 and C_2 sets up a "true dichotomy" if, for any x in the domain of C_1 and C_2, $C_2(x)$ is equivalent to $\neg C_1(x)$. In this case, the Law of the Excluded Middle says that $C_1(x)$ or $C_2(x)$ must be true. For example, *Legal*(x) and *Illegal*(x) set up a true dichotomy. *Likely*(x) and *Unlikely*(x) don't.

In each of the following examples, the second of the two claims is an *English* negation of the first. Indicate whether or not the two claims set up a true dichotomy. In other words, is the second claim the *logical* negation of the first?

a) Let the domain be the set of tasks to be accomplished.

C_1 : x is complete.
C_2 : x is incomplete.

b) Let the domain be the set of people. Assume that we're using "sensitive" in the sense of a personality trait.

C_1 : x is sensitive.
C_2 : x is insensitive.

c) Let the domain be the set of people.

C_1 : x is smart.
C_2 : x isn't smart.

d) Let the domain be the set of decisions.

C_1 : x is fair.
C_2 : x isn't fair.

Ambiguity of Some Logical Operators - OR

And, finally, we must deal with the fact that some English words that appear to refer to logical operators are themselves ambiguous.

We've already seen that the English word "or" can be used to indicate both:

- the logical operator $\vee$ (inclusive or) – Recall that $p \vee q$ is true whenever p or q or both is/are true.
- the logical operator XOR (exclusive or) – Recall that p XOR q is true whenever p or q but not both is true.

Sometimes, context makes it clear what meaning is intended. Sometimes, however, there is ambiguity.

In these examples, context tells us that "or" means inclusive or:

- If it rained or the sprinklers went off, the sidewalk will be wet. (The sidewalks will be wet even if both rain and sprinklers.)
- You'll love the concert if you're into percussion or you love everything Japanese. (You'll like the concert even if both you're into percussion and you love everything Japanese.)
- If you or your partner works at ZZZ Corp, you are eligible for their insurance. (You can get the insurance even if both of you work for ZZZ.)

In these examples, context tells us that "or" means exclusive or:

- Lightening or Black Thunder will win the race tomorrow. (They can't both win.)
- If you have the winning ticket, you get the car or the trip. (You can't have both.)
- I will buy my Calculus textbook from the University Co-op or Amazon. (Students don't buy multiple Calculus textbooks.)

But in these examples, there is ambiguity that could lead different people to different interpretations (and thus ways of encoding the sentence in logic):

- You can have cake or pie for dessert. (Maybe you have to choose, but maybe you can have both.)
- Every house in the neighborhood has a play room or a home theater. (Maybe there was space for only one of these, but maybe it's an upscale neighborhood and some houses have both.)

Problems

1. In each of the following examples, there is one meaning of “or” that most people will find obvious. Indicate which ($\vee$ or XOR) it is.

(Part 1) Whenever I go to Seattle it is cold or rainy.

(Part 2) Whenever I try to call FlyByNightCo, I get a busy signal or the call doesn’t even go through.

(Part 3) Whenever I go to Cupcakes Forever, I run into Maria or Jose.

(Part 4) Whenever I work on Sunday, the head security guard is Tim or Kristin.

Ambiguity of Some Logical Operators - IMPLIES

The logical operator $\rightarrow$ (implies) has a single clear meaning. If you're not sure you remember our discussion of that, you may want to rewatch our Implies video.

https://www.youtube.com/watch?v=3RFamYOCEHA

It's much less clear, however, what we mean when we say, in English, "implies" or "if" or "if/then".

There are at least three common meanings of those terms (and the many other ways there are of saying the same thing):

- ***Material implication***. This is the meaning of $\rightarrow$ that we have been using. So, in particular, the meaning of $p \rightarrow q$ is given by this truth table:

p	q	$p \rightarrow q$
T	T	T
T	F	F
F	T	T
F	F	T

- ***If*** Drew comes, there will be ice cream.
- ***If*** Drew comes ***then*** there will be ice cream.
- Drew coming definitely ***implies*** there will be ice cream.

In all three of these cases, we're saying that Drew's coming means there will be ice cream. We haven't said anything about what will happen if Drew doesn't come. Possibly there are other things that might cause ice cream to appear.

- ***Equivalence*** (if and only if). Sometimes, "if" actually means "if and only if", whose meaning is given by this truth table:

p	q	$p \leftrightarrow q$
T	T	T
T	F	F
F	T	F
F	F	T

In other words, you might say $p \rightarrow q$ but mean $(p \rightarrow q) \wedge (\neg p \rightarrow \neg q)$. Notice that $(\neg p \rightarrow \neg q)$ is the converse of your actual claim. The converse of a statement P is not necessarily true whenever P is. So I (the listener) can't *logically* conclude $\neg p \rightarrow \neg q$. But it may nevertheless part of the meaning that you intended. How can I know that? I will generally assume that you are trying to communicate effectively. So you'll make the strongest claim you reasonably can. If q is true regardless of p, you would just have asserted q. We'll say more about this idea, called ***conversational implicature***, later.

- If it rains, we'll move the picnic indoors.

But what if it doesn't rain? On hearing this sentence, most of us would conclude that, unless it rains, we'll have the picnic outside. Why? Partly because we can't imagine why the picnic would be moved unless there is rain. But also because, if the picnic is going to be moved regardless of the weather, why would you not simply have said, "We're going to move the picnic inside"?

By the way, it is common for mathematicians (and others) to say "if" in *definitions*, when they actually mean "if and only if".

- An integer greater than 1 is prime ***if*** it has no divisors other than itself and 1.

- ***Causality***.

- If the sweater gets wet, the colors will run. (The color run will be caused by the water.)
- If you eat too much, you'll get sick. (Eating too much causes one to be sick.)

Each of these sentences is making a claim of material implication *and* an additional claim about causality. That claim must be represented separately in logic if we want to reason with it.

Problems

1. Indicate for each of the following sentences, which meaning of if/implies is *most likely* intended by the speaker. (Do not stress if your answer to some of these differs from ours. These sorts of sentences can be ambiguous and open to misinterpretation.)

(Part 1) If Koko is tired, she'll be grumpy.

a) Material implication.
b) Equivalence.

(Part 2) Morgan will sing if Casey does.

a) Material implication.
b) Equivalence.

(Part 3) We run out of hot dogs whenever the Astros play.

a) Material implication.
b) Equivalence.

(Part 4) You are a senior if you have at least 90 credits.

a) Material implication.
b) Equivalence.

(Part 5) On Fridays, we go to Torchy's. (Hint: start by rewriting this to make the "if" explicit.)

a) Material implication.
b) Equivalence.

“Paradoxes” of Material Implication

Material implication, the definition of $\to$ that we are using, is similar to the English expressions “implies” and “if/then”. But it isn’t identical. The differences lead to various kinds of confusion, sometimes thought of as paradoxes.

Define: *A*: Austin is in Texas
M: The Moon is made of green cheese.
C: Dallas is the coolest city in Texas.

Now we can look at some of these “paradoxes”.

The “Paradox” of Entailment (or the Principle of Explosion)

> This sentence is true:
>
> If Austin is in Texas and Austin isn’t in Texas then the moon is made of green cheese.
>
> $(A \wedge \neg A) \to M$
>
> This is so even though there is no connection whatever between $(A \wedge \neg A)$ and *M*.

More generally, for any claims *p* and *q*:

$$(p \wedge \neg p) \to q$$

$$False \to p$$

This observation, which follows directly from the truth table *definition* of $\to$, is also called the ***Principle of Explosion*** (because, if we accept *even one* contradiction, we must accept *every* claim as true). It’s also called, in the Latin of classical logic, *ex falso quodlibet* or *ex contradictione sequitur quodlibet*.

If *p* is True

> This sentence is true:
>
> If Austin is in Texas then “The moon is made of green cheese implies that Austin is in Texas.”
>
> $A \to (M \to A)$
>
> Again, this is so even though there is no connection between *A* and *M*.

More generally, there is a simple truth table proof that, for any p and q, if p is true, then any q implies it:

$$p \rightarrow (q \rightarrow p)$$

One Implication Must Be True

This sentence (under the interpretation given by the parentheses shown here) is true:

> (Austin being in Texas implies that Dallas is the coolest city in Texas) or (Dallas being the coolest city in Texas implies that the moon is made of green cheese).
>
> $(A \rightarrow C) \vee (C \rightarrow M)$

More generally (and provable by truth table) for any p, q, and r:

$$(p \rightarrow q) \vee (q \rightarrow r)$$

An English argument for the truth of this claim is: q must be either true or false. If it's true, then $(p \rightarrow q)$ is true (regardless of the truth of p). If it's false, then $(q \rightarrow r)$ is true (regardless of the truth of r).

If *p* Doesn't Imply *q*

This sentence is true:

> If it's not true that Austin being in Texas implies that Dallas is the coolest city in Texas, then Austin is in Texas and Dallas isn't the coolest city in Texas.
>
> $\neg(A \rightarrow C) \rightarrow (A \wedge \neg C)$

More generally (and provable by truth table) for any p and q:

$$\neg(p \rightarrow q) \rightarrow (p \wedge \neg q)$$

An English argument for the truth of this claim is: The only way that the claim $(p \rightarrow q)$ can be false is if p is true, yet q is false.

Problems

1. On the same subject as the old adage that, if there's smoke there's fire, let's assert one premise:

[1] It's not true that fire implies smoke.

For each of the following claims, indicate whether (given the premise) it must be true, it must be false, or it could be either true or false (depending on the truth of other claims).

Hint: Write out the premise and each of these claims in clear Boolean logic.

(Part 1) There's fire and not smoke.
(Part 2) If there's smoke, it's sunny.
(Part 3) There's smoke or there isn't fire.
(Part 4) It's sunny and there's a fire.
(Part 5) If there's no fire, then it's Tuesday.
(Part 6) If it's Tuesday, there's no fire.

Summary

Big Idea

In many cases, if we try to map from English sentences to logical ones without clearly analyzing what the English sentence is saying, we will have trouble. It may be tempting to jump to the conclusion that we have found a weakness in our logical system itself. While there are real limits to the logical system that we are discussing, we should, before assuming that we've found one, see if a more careful translation of English into logic can't make the problem disappear.

We Must Overcome the Perils of English - We Leave Out a Lot and Are Sloppy

Introduction

Talking	Approximate number of words a moderately fast-talking English speaker can utter per minute:	160
Wifi	Approximate number of words a not very state-of-the-art wifi connection can transmit per minute:	30,000,000

English evolved to support talking, not wifi, communication. So:

- In the interest of efficiency, English arguments leave out a lot. We assume that the people to whom we're talking (or writing) share much of our understanding of the world around us.

- We're also often sloppy. Again, we assume that the people we're communicating with will be able to figure out, from what we literally say, what we actually mean.

But the language of logic doesn't allow this. In order to use it to represent the intended meanings of English sentences, we must first determine exactly what those meanings are. And then we must be very careful as we map those meanings into logical statements.

The Cooperative Principle and Conversational Implicature

https://youtu.be/Nocwm9-bH7Y

In our everyday use of language, what we explicitly say and what we actually mean (and intend others to understand) are often quite different.

This happens for a variety of reasons, including our desire to be brief and our need to follow social conventions, such as the rules that govern what it means to be polite.

One important idea that goes a long way toward explaining why communication works even in the face of this problem is the ***Cooperative Principle***.

The philosopher/linguist Paul Grice described a set of four ***conversational maxims***: rules that describe cooperative conversation:

- *Quality*: Say only things that you believe to be true and for which you have adequate evidence.
- *Relevance*: Say only things that are relevant.
- *Quantity*: Give as much information (no more and no less) as is appropriate to the situation. Put differently: make the strongest true and relevant statement you can.
- *Manner*: Be clear and brief.

In some contexts, other maxims also apply. For example:

- *Politeness*: Be polite.
- *Legality*: Don't say things that lawyers tell you not to say.

If we assume that others are following these rules, then sentences can generate ***implicatures***: conclusions that follow from what was literally said.

Suppose that we are trying to find a way to get our whole class to the museum. I say:

> My car can hold 4 people.

If you assume that I'm being cooperative, you will conclude:

> My car can hold 4 people *and* *no more than 4 people*.

If my car can hold 6 people, I have violated the maxim of quantity. The claim, "and no more than 4 people" is an implicature.

Cooperative conversation is efficient because it lets speakers leave out claims that they know can be derived by listeners who assume that the maxims are being followed. But the rules of cooperative conversation complicate the mapping between English sentences and logical meaning.

Problems

1. Suppose that I ask you, "Do you know what time it is?" Mark each of the following possible responses as Cooperative or Not Cooperative. If you're on the fence, choose Not Cooperative. Keep in mind that being succinct is cooperative, but that must be balanced against providing just enough information as to be useful in the current situation. Assume that everything that is stated in any of these responses is true and known to be true by the responder.

a) Yes.
b) Um. Let me see. Um. 2:15.
c) No.
d) No. But there's a clock out in the hall.
e) No. But it's Tuesday.

2. Suppose that I ask you, "Shall we go to a movie tonight?" You reply, "I have to study for an exam tomorrow." Which of the following describes your response:

a) You have not been cooperative. You didn't answer the question.
b) You have been cooperative. I should be able to infer that you've said yes.
c) You have been cooperative. I should be able to infer that you've said no.t

Presuppositions

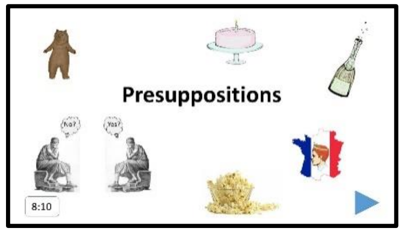

Recall that cooperative conversation is brief. One way to make short claims possible is to leave out shared assumptions. We call such unstated assumptions ***presuppositions***.

The definite article "the" typically carries the presupposition that some object exists (and, often, is unique).

https://www.youtube.com/watch?v=TNFpdNNSays

> Consider: "The President of the United States lives in the White House." This sentence carries the presupposition that such a President exists. That presupposition is true.
>
> Consider: "The King of the United States lives in a castle." This sentence carries the presupposition that such a king exists. But now the presupposition is false.

When presuppositions are true, we generally don't notice them. We have no difficulty interpreting the sentences that carry them. And assigning truth values to those sentences is straightforward.

But when presuppositions are false, we do notice them. We are forced to assume that whoever uttered the sentences that carry them must (unless deliberately talking nonsense) be assuming them to be true. But we are left unsure how to interpret those sentences or to assign truth values to them.

The use of presuppositions isn't limited to our discussions about everyday ideas like kings. We can use them in talking about mathematics.

> Consider: "The identity for multiplication is 1." This sentence carries the (true) presupposition that there is a multiplicative identity.
>
> Consider: "The last decimal digit of the largest prime number is 1." This sentence carries the (false) presupposition that there is a largest prime number.

It is common for presuppositions to assume:

- The existence, and sometimes the uniqueness, of something.

> "The President/King of the United States lives in the White House/a castle."

> "My cat's favorite school subject was math" carries the false presupposition that there exists a school subject that was my cat's favorite.

- That some event occurred or is occurring.

"There are no longer just 13 states in the U.S." carries the true presupposition that there once were just 13 states in the U. S.

"The Flying Pizza Monster was re-elected President of the U.S." carries the false presupposition that the Flying Pizza Monster exists and was already the President.

"While Thomas Dewey was president, Harry Truman sat and licked his wounds" carries the false presupposition that Dewey was ever president.

"After the rain stopped, the worms came out" carries a cascaded pair of presuppositions: It rained. Then the rain stopped.

- That two or more things are comparable.

"The University of Texas is a bigger school than Harvard is" carries the true presupposition that Harvard is a school.

"The University of Texas is a bigger school than Hollywood is" carries the false presupposition that Hollywood is a school.

Notice that, if we negate a sentence that carries a presupposition, its presuppositions don't change.

"There are no longer just 13 states in the U.S." carries the true presupposition that there once were just 13 states in the U. S.

If we change the sentence to "There are still just 13 states in the U.S", the presupposition is still that there were once 13 states. The presupposition is still true, even though the sentence that carries it is now false.

Problems

1. Mark the most accurate truth value claim for each of these English sentences:

(Part 1) The ocean in the middle of Texas is full of fish.

a) True.
b) False.
c) Hard to assign a truth value because it presupposes something false.

(Part 2) They don't sell Twinkies any more.

a) True.
b) False.
c) Hard to assign a truth value because it presupposes something false.

(Part 3) They don't sell flying purple pizza monsters any more.

a) True.
b) False.
c) Hard to assign a truth value because it presupposes something false.

2. Consider the following English sentences:

I. George Washington named his son George, Jr.
II. The bacterium that causes colds can be killed with penicillin.
III. The capital of France is London.

Which of them has/have a presupposition that's false in the world in which we live/have lived?

3. Consider the following formal claims:

I. For every integer, there is another integer that is larger than it is.
II. Each input record will be processed on time.
III. Every employee has a social security number.

Which of them has/have any presupposition at all?

4. Consider the claim: Maria always watches her daughter's soccer games on Saturday.

This claim carries the presuppositions:

[1] Maria has a daughter.
[2] Her daughter plays soccer.

Which one or more of the following sentences carry these same two presuppositions:

a) Maria never watches her daughter's soccer games on Saturday.
b) Maria sometimes watches her daughter's soccer games on Monday.
c) Maria rarely watches her son's soccer games on the weekend.

5. We've talked about three common kinds of presuppositions:

[A] The existence of some object.
[B] That some event has occurred or is occurring.
[C] That two or more things are comparable.

Consider the claim: George, the USA's first President, lost his cell phone again.

This claim carries false presupposition(s) of which type(s)?

Interpreting Sentences that Carry Presuppositions

What truth value shall we assign to sentences like "The largest prime number ends in 1." This sentence carries the false presupposition that there is a largest prime number. Here are two possible approaches:

1) If a sentence *S* carries a false presupposition, assign the truth value False to both *S* and $\neg S$.

2) If a sentence *S* carries a false presupposition, then assign no truth value to either *S* or $\neg S$.

But have we now identified a challenge to the Law of the Excluded Middle (if we take approach 2) or the Principle of Noncontradiction (if we take approach 1)?

No.

Both the Law of the Excluded Middle and the Principle of Noncontradiction are claims about *logical statements*. They are not claims about English sentences.

There is, however, one thing that we *can* say about English sentences: The relationship between them and useful logical statements can be complex. The problem of assigning meaning to English sentences is hard. Centuries of philosophers and linguists have worked on it. We can only scratch the surface of that problem here.

Recall that presuppositions are unstated assumptions. So a good first step, in attempting to map English sentences into logical expressions is to make those assumptions explicit.

Consider again: "The last decimal digit of the largest prime number is 1."

We can rewrite it as: "There exists a largest prime number and the last decimal digit of that number is 1."

Or, in our logical language: $\exists x\ (Prime(x) \wedge (\forall y\ (Prime(y) \rightarrow x > y)) \wedge LastDigitOf(x, 1))$

Now there is no confusion. Since there exists no x such that $(\forall y\ (Prime(y) \rightarrow x > y))$, this entire statement is simply false.

Big Idea

It's critical to keep in mind the difference between English sentences and logical ones. The mapping from the former to the latter is not always straightforward.

Problems

1. Consider: "Taylor's brother's son likes ice cream."

Assume that the referent of Taylor is clear, so there is no issue there.

We want to encode this English sentence as a logical claim that will be either true or false. So we need to make its presuppositions explicit. Which one or more of the following logical expressions correctly does that? (Read *BrotherOf*(x, *y*) as x is the brother of *y*, *SonOf*(x, *y*) as x is the son of y, and *Likes*(x, *y*) as x likes *y*.)

I. $\exists x\ (\exists y\ (BrotherOf(x, Taylor) \wedge SonOf(y, x) \wedge Likes(y, ice\ cream)))$
II. $\exists y\ (Likes(y, ice\ cream) \wedge \exists x\ (BrotherOf(x, Taylor) \wedge SonOf(y, x)))$
III. $\exists y\ (Likes(y, ice\ cream) \wedge SonOf(y, x) \wedge BrotherOf(x, Taylor))$

2. Let the domain be the reals. Consider:

[1] $\sqrt{-2} > 0.$

Which of the following is true:

a) [1] carries no presuppositions.
b) [1] carries one or more presuppositions but they are all true.
c) [1] carries one or more false presuppositions.

3. Let the domain be the reals. Consider:

[1] $\sqrt{-2} > 0.$

We want to encode this claim as a logical claim that will be either true or false. So we need to make any presuppositions explicit. Which one or more of the following is a well-formed logical expressions (assuming the domain is the reals) that correctly does that?

I. $\sqrt{-2} > 0$
II. $\exists x\ (x = \sqrt{2} \wedge x > 0)$
III. $\exists x\ (\sqrt{-2} > 0)$

We Omit the "Obvious"

But presuppositions aren't the only things we leave out. We often construct arguments in which every statement has a truth value. But it isn't possible to prove the conclusion using only the claims that have been explicitly mentioned.

[1] Jamie is holding Alex's Ming vase.
[2] If Jamie lets go of the vase, Alex will be furious.

Probably none of us would question that [2] follows from [1]. But notice that, in evaluating the validity of argument, we bring to the table at least the following additional (unstated) premises:

[3] If Jamie lets go of the vase, it will fall to the floor. (Gravity)
[4] If a porcelain vase falls to the floor, it will break.
[5] A broken vase has very little value.
[6] An unbroken Ming vase has substantial value.
[7] People get furious when the value of their possessions is destroyed.

When we argue to each other, we rely on a pre-existing set of shared premises about our world.

If we want to apply our logical tools, we'll have to be careful. We must make sure that our inference engine has all the premises it needs. We can do that either by providing them explicitly ourselves or by appealing to a knowledge base that has already been crafted by someone else.

We’re Often Sloppy

Sometimes the problem is even worse than just omission. We’re often sloppy and we count on our listeners/readers to figure out what we really must have meant.

Look at the sign shown below on the right. To its left is a picture that makes it clear that the sign doesn’t mean what it says. It probably is meant to say something like, “Must be at least 42″ tall to climb.

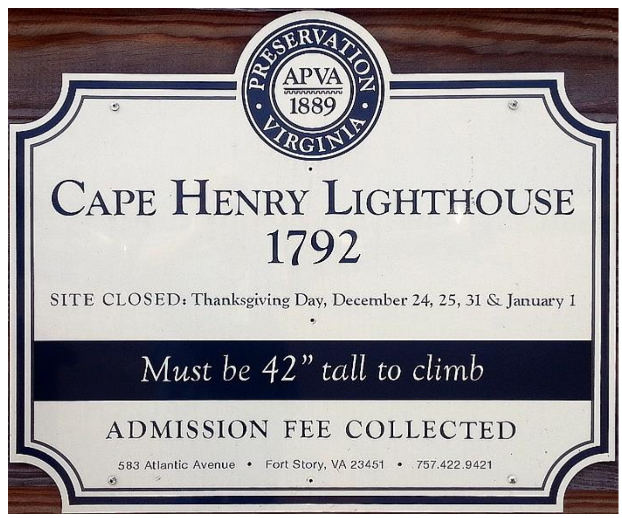

Problems

1. Consider this airport sign. Taken literally, it says that if you're 12 *and* under, the rule applies to you. Let's try to figure out whether that's what it really means.

(Part 1) Reese is 10. Is this claim true or false:

$age(Reese) = 12 \wedge age(Reese) < 12.$

(Part 2) Reese is 12. Is this claim true or false:

$age(Reese) = 12 \wedge age(Reese) < 12.$

(Part 3) Reese is 15. Is this claim true or false:

$age(Reese) = 12 \wedge age(Reese) < 12.$

Back to the Law of the Excluded Middle

Recall that the Law of the Excluded Middle (LEM) says that every logical sentence has a truth value: it must be true or false.

Take another look at the second Law of the Excluded Middle video that we watched back at the time of our discussion of the importance of the LEM is as a theorem-proving tool. Recall that we considered a set of English sentences that appear to challenge the LEM.

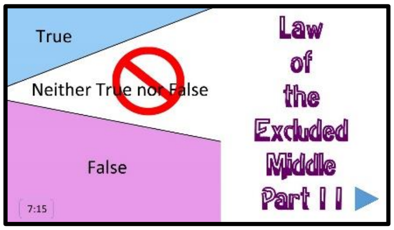

The key to understanding why these examples don't in fact challenge the LEM as a *logical* tool is that their issues are *linguistic*. The way to resolve them is with a careful mapping of the English sentences into the language of formal logic.

https://www.youtube.com/watch?v=r_KG3EZuJmw

Consider:	4 and 7 are friends.
The logical predicate *Friends* isn't defined on numbers.	

Consider:	The king of France has red hair.
When we translate sentences with presuppositions into logic, we must make the (implicit) presuppositions explicit.	

Consider:	The chili is hot.
While English adjectives are often vague, logical predicates cannot be.	

Let's now look at a new example, where we must be careful in how we translate an English adjective into a reasonable logical predicate.

Consider: John was brave.

Suppose that John died peacefully, as an old man, never having faced any situations in which bravery was called for. Could one argue that neither:

Brave(John) nor ¬*Brave*(John)

is true? If so, we have a challenge to the LEM. But let's look more closely at how the English word "brave" can reasonably be translated into logic. What should its definition be?

How about:

Brave(x) iff: *HasShownBravery*(x)

Then, (because he had no opportunity to behave bravely), ¬*Brave*(John) is true. The LEM is still intact, but this may not seem fair to John.

So how about a definition that describes a personality trait, independent of whether that trait has had a chance to manifest itself:

Brave(x) iff: *WouldShowBraveryWheneverNecessary*(x)

Now we can say that one of *Brave*(John) or ¬*Brave*(John) is true. We'll just never know which (because John never got a chance to show us). But, again, the LEM is intact. It says only that there is a truth value. Not that we know what it is.

Nifty Aside

This example is due to Michael Dummett.

Finally, let's consider a class of sentences called ***Liar Paradox*** sentences, which appear to pose a substantive threat to the LEM.

Consider this sentence, which we will call *S*: This sentence is false.

If *S* is true then, taking it at its word, it is false.

If, on the other hand, *S* is false, again taking it at its word, it cannot be false.

It appears, then, that *S* is neither true nor false.

But, again, we must dig deeper. Exactly how shall we represent liar sentences in our logical language? Liar sentences are self-referential: They make a claim about themselves. So we need a way to encode their claims and then assert something about those encoded claims. There is no straightforward way to do that in the first-order predicate logic language that we have defined.

Using other logical mechanisms, however, there are ways to handle Liar sentences. Here are two:

- Define a set of logical languages, arranged in levels. At one level, one can make claims only about sentences at some lower level. Thus no sentence, on any level, can make a claim about itself. (This is the Alfred Tarski solution.)

- Interpret every sentence as implicitly asserting that it, itself, is true (in addition to whatever else it says). So we translate the Liar sentence *S*, above, as:

  This sentence is true $\wedge$ This sentence is false.

  The challenge to the LEM has disappeared. That sentence must be false (since every statement of the form *P* and $\neg P$ is false). (This approach is due to Arthur Prior, among others.)

> ***Big Idea***
>
> The Law of the Excluded Middle applies to *logical* statements, not English ones. And the mapping between English and logic is not always straightforward.

Problems

1. Which of the following sentence pairs constitute Liar Paradoxes? To answer each question, try to find a consistent way of assigning truth values to the two sentences, without appeal to either the Tarski or Prior solutions, described above. If you cannot do so, answer that there is a paradox. If you can, then indicate the truth values that you have found, where [v_1, v_2] means assign the value v_1 to sentence 1 and the value v_2 to sentence 2. If there is more than one consistent assignment, show all of them.

(Part 1) [1] Sentence 2 is true.
[2] Sentence 1 is false.

(Part 2) [1] Sentence 2 is false.
[2] Sentence 1 is false.

(Part 3) [1] Sentence 1 is true.
[2] Sentence 2 is false.

(Part 4) [1] Sentence 2 is true.
[2] Sentence 1 is true.

Predicate Logic Doesn't Solve All Our Representation and Reasoning Problems

Sketching Some of the Problems

Unfortunately, predicate logic all by itself, without substantial additional theory and, in some cases, significant structural changes, doesn't do a good job of capturing the full range of statements that we often make about the world around us.

Consider the following story:

[1] It doesn't rain very often in Austin.
[2] Kelly only likes dancing in the rain.
[3] Kelly won't go anywhere unless he can plan a long time in advance and be pretty sure he'll be able to dance.
[4] Judy won't go anywhere without Kelly unless something really unusual happens.
[5] Fran does exactly what Judy does.

---

[6] Judy probably isn't coming to Austin any time soon. (Because Kelly isn't.)
∴ [7] Fran's mom isn't counting on seeing her in Austin any time soon. (Because Judy isn't coming.)

The derivation of [5] from the premises [1] – [4] seems right. But we can't do it in our system.

Let's look at the lines one at a time to see what's going on:

[1] We need to be able to represent statistical truth. What does "very often" mean? Let's say we could agree that it means that the chances of rain on a given day are less than 5%. How should we represent and reason with even that more concrete fact?
[2] This one we can do if we stretch. But we'll need some way to represent time and place since this sentence is saying that Kelly likes dancing at a particular time, in a particular place, if and only if it's raining at that time in that place.
[3] Again we need statistical reasoning. What does pretty sure mean? And we need to be able to reason from the fact that it doesn't rain very often to the fact that, a long time in advance, it won't be possible to know whether it's going to rain.
[4] How can we represent the "unless" clause here? It's saying that, in the absence of information about some unusual event, we should assume that Judy won't go if Kelly doesn't. In other words, we are to take our lack of knowledge as telling us something. But we must be prepared, if suddenly we are told about an unusual event, to undo our reasoning and give up on the conclusion that Judy won't go.
[5] We can represent actions with predicates, such as *VisitsAustin*(Judy) or *Dances*(Kelly). Then what we want to say here is something like: $\forall P\ (P(Fran) \equiv P(\text{Judy}))$. Read this as, for all predicates *P*, *P* is true of both Fran and Judy or neither of them. But our logical system doesn't allow us to quantify over predicates.
[6] Again we need statistical reasoning.
[7] First, we need to reason that, since Judy isn't coming to Austin, neither is Fran. In addition, how should we represent not just basic facts (such as it's not likely that Judy will come to Austin), but also people's beliefs about those facts?

Summarizing, we've recognized the following issues:

1. We need a way to talk about statistical truth.
2. We need a way to reason about time and place.
3. We need a way to reason with default information (assume unless told otherwise).
4. We need a way to reason not just about *objects*, but also about *predicates*.
5. We need a way to reason about what we know or believe.

There exist reasoning systems that solve these (and other) problems. We can't go into them here. But in the next few slides we'll say a little more about the issues that we've just raised.

Problems

1. Consider the following argument:

[1] Most peppers are spicy.
[2] Chris doesn't like spicy food.
[3] Chris saw Skip put a lot of peppers in the chili for the party.
[4] Chris won't go to a party unless he loves the food or his favorite movie star will be there.

[5] Taylor has given up hoping to see Chris at the party.

Only one of the first four claims is fairly straightforwardly representable in the logical framework we've been using. The others aren't. Which one is representable?

Dichotomizing the Analog World

Introduction

A logical statement (in the systems we have defined) is either true or false. There's no such thing as, "slightly true", or "mostly true" or "really true." Unfortunately we live in an analog world.

"The chili is hot."

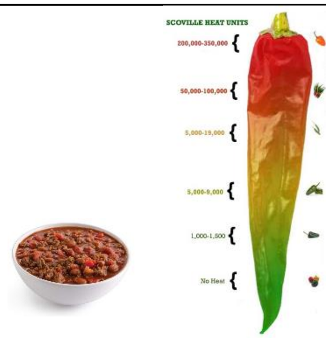

How hot does the chili have to be to count as "hot"?

"It's likely to rain."

What counts as "likely"? 95% certainly does. What about 65%

"That school is hard to get into."

What counts as "hard to get into"? Only 1 out of 25 applicants get in? Or 1 out of 10? Or 1 out of 5?

The Sorites Paradox

Suppose that we accept the following two premises:

[1] One grain of sand is not a heap.

[2] One grain of sand is too small to make a difference in determining whether something is or is not a heap.

Then, nothing, including this can be a heap:

But it clearly is a heap.

Thus we appear to have a logical paradox, classically called the ***Sorites Paradox***.

The English word "heap" is vague: Are 75,000 grains enough? 120,000 grains? So we have a problem in deciding how to map English sentences about heaps into logic.

https://www.youtube.com/watch?v=brz8tIYV1U8

But we do not have a logical problem. The reason that we appeared, above, to get into trouble, is that we chose, as premises, to claims that don't, in fact, do a good job of characterizing heapness.

Problems

1. Sometimes, a Sorites-like problem formulation may look reasonable (as in the case of the heap of sand). But sometimes it's obvious that such a pair of premises doesn't correspond to the situation. In these latter cases, it's easy to see how to write premises that do clearly define the boundary between having some property *P* and not having it.

Mark each of the following pairs of premises as Sorites-like if, as in the sand example, it's hard to decide how to get out of the apparent paradox. Mark it not Sorites-like if it is straightforward to write a more accurate set of premises.

(Part 1) A one word story isn't a novel.
Adding one word to a short story won't make it a novel.

(Part 2) A one-year old is not old enough to drive.
Getting one year older doesn't change whether or not you're old enough to drive.

(Part 3) A one gram package is easy to carry.
Adding one gram to a package doesn't change how hard it is to carry.

(Part 4) One credit isn't enough to get you a degree.
Getting one additional credit can't make you degree-eligible.

(Part 5) One molecule of detergent won't get the clothes clean.
Adding one additional molecule of detergent to the wash won't change how clean the clothes get.

Explanation: There doesn't seem to be a hard and fast boundary that separates "enough detergent" from "not enough" detergent.

Taming Vagueness in Describing the Everyday World

> Consider the question of height:
>
> [1] Professional basketball players are tall.
> [2] Elephants are very tall.
> [3] Giraffes are super tall.

Sometimes we make comparative statements:

> Indian food is spicier than Irish food.

How shall we represent claims such as these in the logical language that we've defined? The answer is that it depends on how we want to reason with them.

One standard approach is to define a numeric scale (say 1 to 10). Then we can make numeric claims. We can specify cutoffs that define adjectives like tall.

If we want to be more sophisticated, we can make such cutoffs relative to a reference set.

> For example, it's possible to be a tall person at a height that would be in midget territory for a giraffe.

There are other digitizing issues that don't so naturally correspond to numeric scales.

> Suppose that I have an apple in the refrigerator. I take it out and eat one bit and put it back, I probably still have an apple in the fridge. What if I eat half the apple and put the rest back. Do I still have an apple in the fridge? What if I eat all but one bite?

A lot has been written about issues such as these. We'll have to skip most of them for now. The key for us will be to stay focused on the reasoning that we want to do. Then we must make representation decisions that let us do that reasoning. We should never imagine that we've captured everything anyone might want to say.

Taming Vagueness in Formal Applications

We can't typically tolerate vagueness in formal applications, including:

- Mathematics
- Software and hardware specifications
- Critical databases

So, while philosophers and linguists have, over the centuries, devoted substantial attention to the problem of vagueness, there are many practical situations in which we effectively define it away.

Unacceptable specification for a factory control system:

> The line must shut down if the machine gets ***hot***.

Acceptable specification for a factory control system:

> The line must shut down if the internal temperature of the compressor exceeds 150° F.

Unacceptable specification for a Human Resources database application:

> An employee who has worked here a ***long time*** can retire.

Acceptable specification for a Human Resources database application:

> An employee *E* can retire if and only if:
>
> $AgeInYears(E) + NumYearsEmployeed(E) \geq 80$

Problems

1. Consider this vague rule about who may invest in our new, risky hedge fund:

> Only rich people may invest in the Gotcha Fund.

Assume that these predicates mean what they appear to mean. Indicate, for each of the following clear statements, whether or not it could be what was intended:

a) $\forall x\ ((Networth(x) > \$10M) \equiv MayInvest(x, GF))$

b) $\forall x\ ((AnnualIncome(x) > \$1M) \equiv MayInvest(x, GF))$

c) $\forall x\ ((Networth(x) > \$100M) \equiv MayInvest(x, GF))$

d) $\forall x\ ((NumberOfHousesOwned(x) > 3) \rightarrow MayInvest(x, GF))$

Statistical (Likelihood) Reasoning

Consider these two statements:

[1] People have more neurons than earthworms do.

[2] Healthy male birds sing faster than sick ones do.

These may sound like similar claims. But they're different in an important way. The first is a categorical statement. All people have more neurons than all earthworms. The second, though, is a statistical statement. We might rephrase it as:

[3] On average, healthy male birds sing faster than sick ones do.

Or we could actually show the distributions for song speed for two groups of birds:

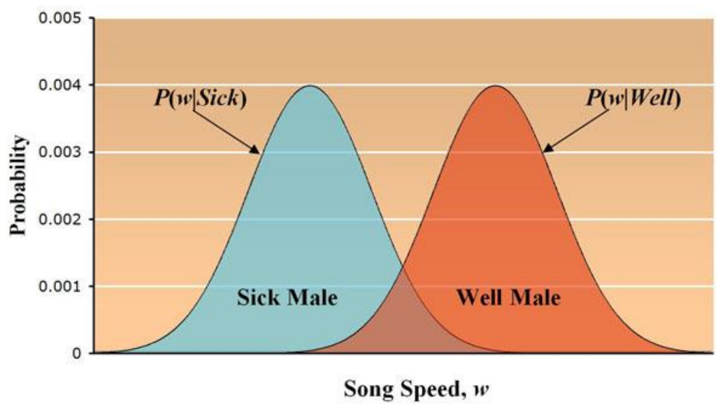

As the distribution makes clear, it isn't true that all healthy birds sing faster than all unhealthy ones. But the statistical claim that healthy birds are faster singers than sick ones is correct.

Many everyday claims are statistical ones.

[4] Older children are taller than younger ones. (As in the bird song example, this is a claim about distributions. There certainly are 8-year olds who are taller than some 10-year olds.)
[5] In Florida, it rains in the afternoon. (This sentence means that the probability of rain in any given afternoon is very high, although not certain.)
[6] Children like ice cream better than broccoli. (Again, I won't accuse you of lying just because I find a single child who likes broccoli better. This is a statistical claim.)
[7] It's likely to rain tomorrow. (Sometimes we even make explicit the statistical nature of our claim.)

Fortunately, the science of statistics gives us precise tools for reasoning with these kinds of statements when we need to. And that science rests on the core logical structures that we will study. Logic is not irrelevant to this kind of reasoning. It's just not enough.

Most

A common kind of everyday claim involves the notion of “most”. Let’s analyze “most” reasoning more carefully since it illustrates one fundamental way in which statistical reasoning differs from the kind of absolute reasoning that we’ve been studying.

We’ve shown that:

$$\forall x\,(P(x) \to Q(x))$$
$$\forall x\,(Q(x) \to R(x))$$

$$\therefore\ \forall x\,(P(x) \to R(x))$$

So, for example, if all cats are mammals and all mammals have lungs, then all cats have lungs.

But now suppose that we tried to introduce a new quantifier (sort of an upside down M) corresponding to the idea of “most”:

Ɯ $x\,(P(x) \to Q(x))$
Ɯ $x\,(Q(x) \to R(x))$

$\therefore$ Ɯ $x\,(P(x) \to R(x))$

If we allowed this reasoning, we’d get some useful things.

Most soda has a lot of sugar.
Most sugary stuff is unhealthy.

∴ Most soda is unhealthy ☑

But we’d also get some junk.

Most American adults can read.
Most adults who can read are not Americans.

∴ Most American adults are not Americans. ☒

This sort of attempt to reason with *most* can fail even if one of the quantifiers is universal.

All babies are children.
Most children go to school.

∴ Most babies go to school. ☒

We need a statistically based reasoning system to work with problems like these.

Problems

1. Define:

$P(x)$: True if x is an American soldier fighting in Afghanistan.
$Q(x)$: True if x was born in the United States.
$R(x)$: True if x has never been to Afghanistan.

Assume that most American soldiers fighting in Afghanistan were born in the United States and that most people born in the United States have never been to Afghanistan. Consider this argument:

$$\begin{array}{l} \text{W}\ x\ (P(x) \rightarrow Q(x)) \\ \text{W}\ x\ (Q(x) \rightarrow R(x)) \\ \hline \therefore\ \text{W}\ x\ (P(x) \rightarrow R(x)) \end{array}$$

In this case, is the conclusion true or false?

2. Define:

$P(x)$: True if x is a knife.
$Q(x)$: True if x is sharp.
$R(x)$: True if x is dangerous.

Assume that most knives are sharp and that most sharp things are dangerous. Consider this argument:

$$\begin{array}{l} \text{W}\ x\ (P(x) \rightarrow Q(x)) \\ \text{W}\ x\ (Q(x) \rightarrow R(x)) \\ \hline \therefore\ \text{W}\ x\ (P(x) \rightarrow R(x)) \end{array}$$

In this case, is the conclusion true or false?

3. Define:

$P(x)$: True if x is a kitten.
$Q(x)$: True if x is fluffy.
$R(x)$: True if x is cute.

Assume that most kittens are fluffy and that most fluffy things are cute. Consider this argument:

$$\mathsf{W}\, x\, (P(x) \rightarrow Q(x))$$
$$\mathsf{W}\, x\, (Q(x) \rightarrow R(x))$$
$$\therefore \quad \mathsf{W}\, x\, (P(x) \rightarrow R(x))$$

In this case, is the conclusion true or false?

4. Define:

$P(x)$: True if x is a bird.
$Q(x)$: True if x can fly.
$R(x)$: True if x is an insect.

Assume that most birds can fly and that most things that can fly are insects. Consider this argument:

$$\mathsf{W}\, x\, (P(x) \rightarrow Q(x))$$
$$\mathsf{W}\, x\, (Q(x) \rightarrow R(x))$$
$$\therefore \quad \mathsf{W}\, x\, (P(x) \rightarrow R(x))$$

In this case, is the conclusion true or false?

5. Define:

$PB(x)$: True if x is a professional basketball player.
$RH(x)$: True if x is right-handed.
$PA(x)$: True if x is a professional athlete.

Assume that most professional basketball players are right-handed and that most right-handed people are not professional athletes. Consider this argument:

$$\mathsf{W}\, x\, (PB(x) \rightarrow RH(x))$$
$$\mathsf{W}\, x\, (RH(x) \rightarrow \neg PA(x))$$
$$\therefore \quad \mathsf{W}\, x\, (PB(x) \rightarrow \neg PA(x))$$

In this case, is the conclusion true or false?

Extending the Logical Framework

A variety of ways of extending classical logic to solve some of these problems have been proposed. We'll look briefly at a few of them.

Nonmonotonic Reasoning

So far, the reasoning system that we've been studying has the following key property:

- As we add premises, it is possible (in fact, likely) that new statements will become theorems.
- As we add premises, no old theorems get eliminated. In other words, if it was possible to prove P before adding the new premise(s), it is still possible to prove it.

Thus we'll say that our reasoning system is ***monotonic***. The set of provable claims can change in only one (thus the prefix "mono-") direction as we add premises (in this case, it can only grow).

But not all everyday reasoning has that property. In particular, whenever we make do with incomplete information by exploiting some kind of default fact (for example, assume it's not raining in Austin unless you know otherwise), we must be prepared to undo some conclusions if new information (for example, that it is actually raining in Austin) comes in.

In ***default reasoning***, we conclude some fact *in the absence of* some indication that we should do otherwise.

In the system we already have, we can write:

[3] $\forall x\, ((P(x) \land \neg Q(x)) \rightarrow C(x))$

But, to use [3] to conclude $C(x)$, we must actually be able to prove $\neg Q(x)$. What if we simply don't know anything about $\neg Q(x)$? We're stuck.

But suppose we could write:

[4] $\forall x\, (P(x) \rightarrow C(x)$ UNLESS $Q(x))$

We w ant to interpret this to mean that we can conclude $C(x)$ *unless* we have explicit information that $Q(x)$ is true. Of course, if new information about $Q(x)$ shows up, we'll have to add it to our system. We'll also have to eliminate any conclusions that we derived based on its absence.

Consider the following claim that might be useful in trying to solve the problem of getting to the airport:

CanDriveToAirport UNLESS (¬CarWillStart ∨
FlatTire ∨
FireTrucksBlockingDirveway ∨
StreetsCoveredInIce ∨
CityLockedDownInEmergency)

Notice that, in everyday planning, we are rarely even conscious of all the terms in the UNLESS clause. If we wanted to build a problem-solving robot (perhaps even a self-driving car), we'd need it also to plan an action without having to stop first and verify that all of the UNLESS terms are in fact false.

Our inference system can't do this. But there exist formal systems that can. And, of course, people do it all the time.

Big Idea

We use default reasoning to enable us to reason about ***typical situations*** without getting completely bogged down worrying about all the unlikely things that *might* occur. People do this all the time. Intelligent agents and robots will also have to be able to do it.

Problems

1. Consider the statement:

[1] $\forall x\ ((Room(x) \wedge \exists y\ (Lamp(y) \wedge In(y, x) \wedge TurnedOn(y))) \rightarrow Lit(x))$

If there's a lamp in the room and the lamp is turned on, then the room will be lit. This fact could be useful, for example, to a household robot. It suggests that, upon entering a dark room, a reasonable thing to do would be to turn on a lamp.

But what we really should have written here is:

[2] $\forall x\ ((Room(x) \wedge \exists y\ (Lamp(y) \wedge In(y, x) \wedge TurnedOn(y))) \rightarrow Lit(x)$ UNLESS (…))

In the real world, many things could go wrong and prevent the room from lighting up when the lamp is turned on.

List at least three such things.

2. Consider the statement:

[1] $\forall x\ ((Warehouse(x) \wedge Unlocked(x)) \rightarrow CanHideIn(x))$

If you need to hide someplace and there's a warehouse nearby, you can hide there. This fact could be useful, for example, to an agent in a first person shooter game.

But what we really should have written here is:

[2] $\forall x\ ((Warehouse(x) \wedge Unlocked(x)) \rightarrow CanHideIn(x)$ UNLESS (…))

In the real world, many (generally very low probability) things could prevent it being possible to hide in a warehouse.

List at least three such things.

Inheritance

Recall that we've already considered the issue of the flying capabilities of birds. We might be tempted to say:

[1] $\forall x\ (Bird(x) \rightarrow CanFly(x))$

And, if we did that, we'd be right most of the time. But emus and penguins can't fly. Neither can birds that have just been born or ones with crude oil on their wings or ones with broken wings. What we really want to say here is something like:

[2] $\forall x\ (Bird(x) \rightarrow CanFly(x)$ UNLESS $(Emu(x) \vee Penguin(x) \vee Baby(x) \vee WingBroken(x) \vee \ldots))$

But we want [2] to function like [1] most of the time. If I know that x is a bird and I know nothing else, I want to assume that it can fly.

The flying birds situation is an example of a very common sort of default reasoning in which individuals are assumed to ***inherit*** (take on) the properties of a typical member of some class to which they belong. Birds typically fly. Dogs typically have tails. Houses typically have kitchens. People typically can talk.

As with any sort of nonmonotonic reasoning, we must be prepared to undo conclusions if new information (such as the fact that our bird has a broken wing) comes in.

Problems

1. We might be tempted to say that anyone who is a friend inherits the property of being trustworthy:

[1] $\forall x\ (Friend(x) \rightarrow CanBeTrusted(x))$

If we did that, we'd be right most of the time. But there are exceptions. Suppose that we're trying to program a game agent. We might want to say:

[2] $\forall x\ (Friend(x) \rightarrow CanBeTrusted(x)$ UNLESS (…))

Think of at least three claims that could go inside the UNLESS clause.

The Closed World Assumption

There's a specific kind of reasoning in the absence of facts that often makes sense when we're working with the contents of a database.

> Suppose that we want a rule that says that we'll drop from our approved supplier list any company from whom we haven't bought anything in the last 24 months. We might write (for simplicity, assume that the date of a transaction is stated in months):
>
> [1] $\forall x\ ((\neg\exists y\ (Order(y) \wedge (Supplier(x, y) \wedge (Month(y) > Now - 24))) \rightarrow \neg ApprovedSupplier(x))$
>
> Read this as, "For any x, if there doesn't exist an order y that was supplied by x more recently than 24 months ago, then x isn't an approved supplier.
>
> Fine. But how shall we prove that there *doesn't exist* an order?

Let's generalize a bit. Suppose that, in some arbitrary application, we have:

[2] $\neg\exists x\ (P(x)) \rightarrow Q$

In other words, if there's no x that satisfies some property P, then we can assert some claim Q. If we want to reason with [2], how shall we go about showing that there is no such x? If we can actually prove that, great. Absent such a proof, we could look around for a while to see if we can see an example x. But failing to find one is not a proof that none exists. For example, we could look around for quite a while and fail to find an albino peacock. That is not, however, a proof that none exists. In fact, they do. So we can't, in general admit "I tried to find one but I couldn't," as a proof technique.

In a database application, on the other hand, we often want to do exactly that. Returning to our outdated supplier list problem: If we systematically check the database but fail to find a recent order from some supplier, we can conclude that there hasn't been such an order. (Or at least we can if we assume that our order entry process works as it should.) We can thus conclude that the supplier should be dropped.

In the special case in which we believe that we are operating in a ***closed world*** in which all relevant facts are explicitly asserted to be true, we can take the absence of an assertion as equivalent to the claim that it is false. When we do this, we'll say that we're exploiting the ***closed world assumption***.

Implementing this requires that we add a technique that can't be described in the predicate logic framework that we've been using.

Problems

1. Assume that you have access to the registrar's database at University U. You also have the following rule:

$\forall x\ ((Student(x) \land (\neg\exists c\ (Course(c) \land Taken(x, c) \land DeptOf(c, English)))) \rightarrow \neg CanGraduate(x))$

In other words, no one can graduate without taking an English course.

Suppose that we want to try to prove that some student *x* cannot graduate. Which of the following is true:

a) The closed world assumption is valid here. Failure to find an English course taken by *x* should enable us to conclude that *x* cannot graduate.
b) The closed world assumption is not valid here. Failure to find an English course taken by *x* tells us nothing for sure.

2. Assume that you can look at your friend Chris's Facebook account and see her list of Facebook friends. You also take the following as a premise:

$\forall x\ (\neg Knows(Chris, x) \rightarrow \neg WorthKnowing(x))$

In other words, anyone Chris doesn't know isn't worth knowing.

Suppose that we want to try to prove that some person *x* isn't worth knowing. Which of the following is true:

a) The closed world assumption is valid here. Failure to find *x* on Chris's list of Facebook friends should enable us to conclude that Chris doesn't know *x*.
b) The closed world assumption is not valid here. Failure to find *x* on Chris's list of Facebook friends tells us nothing for sure.

Higher Order Logic

The logical framework that we have been describing is called ***first-order*** because it allows quantification over variables (and thus the objects to which they refer) but not over predicates or functions. So, for example:

- We can say: $\forall x\ (Bear(x) \rightarrow Mammal(x))$

  We've quantified over the objects of which *Bear* or *Mammal* might be true.

- We cannot say: $\forall P\ (\forall x\ ((LivingProp(P(x)) \wedge Dead(x)) \rightarrow \neg P(x)))$

  What we're trying to do here is to quantify over predicates to say that if *P* is any predicate that describes a property of living things, then it ceases to be true if the thing in question becomes dead. In a first-order system, we can't do this.

Here's an example of how we could exploit quantification over predicates if only it were allowed:

We'll let *LivingProp* be true of any predicate that is true only of living things. Then we can assert:

[1]	*LivingProp(Breathing(x))*	Breathing is a property of living things.
[2]	*LivingProp(NeedsWater(x))*	Needing water is a property of living things.
[3]	*LivingProp(CellsDivide(x))*	Cells dividing is a property of living things.
[4]	*LivingProp(CanSing(x))*	Can sing is a property of living things.
[5]	$\forall P\ (\forall x\ ((LivingProp(P(x)) \wedge Dead(x)) \rightarrow \neg P(x)))$	Properties that hold of living things become false when something dies.
[6]	*Dead(Elvis)*	

Now we'd like to be able to conclude:

[7]	$\neg$*Breathing(Elvis)*	
[8]	$\neg$*NeedsWater(Elvis)*	
[9]	$\neg$*CellsDivide(Elvis)*	
[10]	$\neg$*CanSing(Elvis)*	This is the one that makes many people sad.

A system that allows quantification over predicates is called ***second-order***. Even higher order logics exist. In them, it's possible to quantify over properties of properties, and so forth.

The reason that we don't allow these sorts of quantification in the formal system that we're defining is that, if we do, both the computational and the formal logical properties of the resulting system would make using it very difficult.

This means that, to get the same effect, we must say what we have to say separately for each predicate.

Continuing with our example, we'd need to replace [5] with this set of premises:

[5a] $\forall x\ (Dead(x)) \rightarrow \neg Breathing(x))$
[5b] $\forall x\ (Dead(x)) \rightarrow \neg NeedsWater(x))$
[5c] $\forall x\ (Dead(x)) \rightarrow \neg CellsDivide(x))$
[5d] $\forall x\ (Dead(x)) \rightarrow \neg CanSing(x))$

Of course, when we reason informally and don't care about how computationally difficult it would be to create a formal proof automatically, we do use this sort of higher order reasoning quite often.

Problems

1. Which of the following statements is/are legal in first-order logic?

a) $\forall x\ (P(x) \rightarrow \mathrm{Q}(x))$
b) $\forall P\ (\forall x\ (P(x) \rightarrow Q(x)))$
c) $\forall x\ (P(x) \rightarrow \neg P(x))$
d) $\forall x\ (\forall P\ (\mathrm{P}(x) \rightarrow Q(x)))$
e) $\forall P\ (P(Chirpy) \rightarrow P(Elvis))$

2. Recall this claim about what happens to all predicates of which *LivingProp* is true:

[5] $\forall P\ (\forall x\ ((LivingProp(P(x)) \wedge Dead(x)) \rightarrow \neg P(x)))$ Properties that hold of living things become false when something dies.

Suppose that we allowed reasoning about predicates and that we already had the axiom given above about what happens to living properties if an object dies. If you were axiomatizing the world in which we live, which (one or more) of the following axioms/premises could you assert (assume that the predicate names are mnemonic and correctly describe the property in question):

I. *LivingProp*(*WillDieLater*(*x*))
II. *LivingProp*(*BlowsInTheWind*(*x*))
III. *LivingProp*(*HasMass*(*x*))

Equality

But, getting back to the formal system that we're setting up here: There is one restricted use of higher order reasoning that is so important that even most computational logic systems support it. We'll allow it too.

The thing we can't live without is ***equality***. We want to be able to say (typically because we've managed to prove it) that two things are equal.

But what does it actually *mean* for two things to be equal? What we want it to mean is that the two things share all the same properties. Oops. That's a second-order logic idea. What we want to be able to say is:

$$\forall P\ (\forall x, y\ (((x = y) \wedge P(x)) \rightarrow P(y)))$$

What this says is that, for any property *P*, if there are two objects *x* and *y* and they are known to be equal, then if *P* is true of *x* it must also be true of *y*.

Let's look at an example where equality is exactly what we need. Suppose that we are given the following premises:

[1] Roommates share an address. In other words, if x and y are roommates and if z is the address of one of them, then z is the address of the other.
[2] Kelly's address is prestigious.
[3] Kelly and Chris are roommates.

We can write these formally as follows, assuming the existence of a function, *addressOf*, that returns the address of its argument:

[1] $\forall x, y\ (Roomates(x, y) \rightarrow (addressOf(x) = addressOf(y)))$
[2] $Prestigious(addressOf(Kelly))$
[3] $Roomates(Kelly, Chris)$

We'd like to be able to show that Chris also lives at a prestigious address. Let's assume that we have the premise that tells us that if two objects are equal then they share all properties. That's (as stated above):

[4] $\forall P\ (\forall x, y\ (((x = y) \wedge P(x)) \rightarrow P(y)))$

Then we can reason as follows:

- Since Kelly and Chris are roommates, *addressOf*(*Kelly*) = *addressOf*(*Chris*). (from [1])
- *Prestigious*(*addressOf*(*Kelly*)) is a premise.
- But *addressOf*(*Chris*) has all the same properties as *addressOf*(*Kelly*).
- So we have *Prestigious*(*addressOf*(*Chris*)).

We need this interpretation of what equality actually means if we want to do arithmetic and algebra in the way we are used to doing them.

Suppose that we have:

[1] $x = y$
[2] $x + b > 10$

We should be able to reason that, since x and y are equal, we can substitute one for the other anywhere. So we should be able to derive:

[3] $y + b > 10$

A first-order system that branches out just enough to allow equality, in this meaningful sense, is called (of all things) a ***first-order system with equality***.

Big Idea

Equality is such an important concept that almost all practical first-order systems are extended to support it.

Explicit Reasoning about Knowledge and Belief

So far, we've seen that we can write logical expressions that describe what we believe to be true. We call them our premises. Then we can reason with them and derive additional claims that we will then also *believe* to be true because we've proved them so.

So we've been, in some sense, reasoning about belief. Or perhaps you want to make the stronger claim that we're reasoning about facts that we *know* to be true.

But what we have not done, nor are we able to do in our framework, is to reason about claims that explicitly mention believing or knowing.

For example, consider:

[1] I know that Kelly is a student.
[2] I believe that Chris will come to the party.
[3] I think that Chris won't come to the party, so I won't either.

It's not that we can't write such claims at all. We can, in a clunky way:

Give names to the following statements:

K: Kelly is a student.
P: Chris will come to the party.
C: Chris won't come to the party.

And define:

Know(x, *y*):	True if x knows *y*.
Believe(x, *y*):	True if x believes *y*.
Think(x, *y*):	True if x thinks *y*.
WillAttendParty(x):	True if x will attend the party.

Then we can write:

[1′] *Know*(*I*, *K*)
[2′] *Believe*(*I*, *P*)
[3′] *Think*(*I*, *C*) $\wedge$ $\neg$*WillAttendParty*(*I*)

In some sort of way, we've encoded the facts of [1] – [3]. The problem is that we haven't encoded them in a useful way.

Take another look at [3′], for example. Notice that the thing that I think (namely that Chris won't come to the party) is an atomic claim, *C*. It doesn't explicitly represent party attendance in any way that would make it straightforward to reason about my party attendance versus Chris's.

Perhaps you're wondering why we can't be clearer.

For example, why can't we define:

Student(x):	True if x is a student.
WillAttendParty(x):	True if x will attend the party.

And then we could write:

[1″] *Know*(*I*, *Student*(*K*))
[2″] *Believe*(*I*, *WillAttendParty*(*C*))
[3″] *Think*(*I*, $\neg$*WillAttendParty*(*C*)) $\wedge$ $\neg$*WillAttendParty*(*I*)

The problem is that, in a first-order logic system, we can't make claims about predicates such as *Student* or *WillAttendParty*. Predicates have truth value and they can be combined using the logical operators. But they can't themselves be the arguments of other predicates. Nor can complete logical statements. The only things that we can assert claims about are specific objects.

Recall that, in our system, objects can be indicated either by a specific name or value or as the value of a function.

- *Chris*, *London*, and 5 are values. They can be arguments to predicates.
- *age*(x) is a function. Given a specific value for x, it returns another specific value (maybe 25). That value can then be an argument to a predicate.
- *WillAttendParty*(x) is already a predicate. It cannot be an argument to another predicate.

So in particular, in our system, we cannot make explicit assertions about belief or knowing. In fact, we can't make explicit assertions about anything that we're already representing as a predicate.

There exist belief logics that can solve this problem, but they are beyond the scope of this class.

Problems

1. We're going to consider a set of claims. Try to represent them in our logical system. Think carefully about the predicates you use. Even if you do that, however, you'll have trouble encoding one of these claims. Which one?

a) Shelby likes peanut butter.
b) Everyone who likes peanut butter likes jelly.
c) Jody likes everyone who likes jelly.
d) Casey likes it when Jody likes someone..
e) Casey hates peanut butter and jelly.

So Where Does That Leave Us?

By now, you've seen that, in the predicate logic framework that we've described, we can:

- Naturally represent and reason with mathematical facts. (By the way, there are many more examples of this in our follow-on course, Sets, Relations and Functions.)

$$\forall x\ (Prime(x) \equiv (\neg(x = 1) \wedge \forall y\ (Div(x, y) \rightarrow ((x = y) \vee (y = 1)))))$$

- Naturally represent and reason with formal specifications for programs.

$$\forall x\ (Disk(x) \rightarrow On(x, Pole_1, time_1))$$

- Naturally represent and reason with constraints and policies that are stated in terms of objects in databases.

$$\forall x\ ((Loan(x) \wedge Approved(x)) \rightarrow (\exists y\ (\exists z\ (Super(y) \wedge Signed(y, x) \wedge Signed(z, x) \wedge \neg(y = z)))))$$

- Sometimes represent and reason with our knowledge about the everyday world.

$$\forall x\ (\forall y\ (Friends(x, y) \rightarrow Friends(y, x)))$$

Of course, we wish that reasoning about the everyday world were easier. But even if that's what you care primarily about, the logical foundation that we've built will serve you well. There exist formal extensions to predicate logic that can solve many of the problems that we've encountered. And sound everyday reasoning, while not sticking to our formalism, rests on its foundation.

[1] Everyone who's invited to the party works at MegaLogicLand.
[2] Every couple who's coming to the party has to bring something.
[3] ☑ Jamie is planning to attend the party, but she knows that her partner is bringing cookies, so she's not worried about getting something.
[4] ☒ Since Jordan works at MegaLogicLand, he assumes he's invited to the party.

Jamie's reasoning is sound. In its essence, [2] says:

[2] Party → Bring something

Jamie knows that her partner's bringing cookies makes that claim true for them. So she need take no further action.

But Jordon's reasoning is flawed. In its essence, [1] says:

[1] Invited → MegaLogicLand

Jordan has used Modus Ponens backwards to attempt to reason:

☒ MegaLogicLand → Invited

And that, we know, does not preserve truth.

Big Idea

The logical tools that we have developed form the basis for sound reasoning in mathematics, in many areas of computer science, and in our everyday world.

Problems

1. Consider the following claims:

[1] Everyone who thinks that it's going to rain tomorrow will bring an umbrella.
[2] It's going to rain tomorrow.
[3] The weather forecast is very accurate if you ask for just one day in the future.
[4] Bailey just checked the weather forecast.
[5] Dana never checks the weather forecast.
[6] Bailey will bring an umbrella tomorrow.
[7] Dana won't bring an umbrella tomorrow.

Take [1] – [5] as premises. Which of the following statements is correct:

a) [6] follows from the premises but [7] does not.
b) [7] follows from the premises but [6] does not.
c) Both [6] and [7] follow from the premises.
d) Neither [6] nor [7] follows from the premises.