Key Skills for Computer Science

Lecture 3: Reasoning

Michael Wooldridge
(mjw@liv.ac.uk)

UNIVERSITY OF LIVERPOOL

What is Logic?

- Many views of logic:
  - logic is concerned with *consistency*
  - logic is concerned with *deriving new knowledge from existing knowledge*
  - logic is concerned with *valid arguments*
- In fact, as we will see, these are all different ways of looking at the same thing

Set text for this
bit of the course:
"Logic", by W Hodges, Penguin.

Logic as the Study of Consistency

- Here, we do not mean consistency as in "always doing the same thing"
- We mean *compatibility of beliefs*
- A set of beliefs is *consistent* if the beliefs are compatible with one another, and *inconsistent* otherwise

Consistency

- A more precise definition:

  A set of beliefs is *consistent* if there is *some possible way that they could be true*

- Note that we don't require the beliefs to be plausible, or likely, or even to really make sense... Logic is neutral about this.
  - We just require that there is *some conceivable state of affairs in which they are true*

Consistency: An Example

“It would be wrong to censor violent programmes on television, because people’s behaviour isn’t really affected by what they see on the screen. All the same, it would be a good idea to have more programmes showing the good sides of life, because it would straighten out some of the people who are always knocking our country.”

Is this consistent?

Consistency: An Example

"During the last five years, I have been involved in three major accidents and several minor ones while driving my car. After two of the major accidents, the court held me responsible. But basically, I'm a safe driver; I've simply had a run of bad luck."

Is this consistent?

Consistency: An Example

“I’ve invented an amazing new sedative
which makes people faster
and more excited.”

Is this consistent?

Consistency: An Example

"There is no housing shortage in Lincoln today. This is just a rumour that is put about by people who have nowhere to live."

Is this consistent?

Consistency: An Example

"I knew I would never get pregnant.
But somehow it just happened."

Is this consistent?

Consistency: An Example

"I've never drawn anything in my life. But if I sat down to draw now, it would take me two minutes to produce a drawing worth as much as anything Picasso ever produced."

Is this consistent?

Consistency is the Most Basic Test of a Theory or Argument

- The first point of analysis of a theory or argument is for consistency
- If the theory or argument fails this test, we can immediately dismiss it:
  - *there is no conceivable way that the theory/argument could be true*

The Content of Beliefs

- Beliefs are *declarative sentences* or *propositions*
- In English, a proposition is *a sentence that is either true or false*
- We call “true” and “false” the *truth values,* or the *Booleans*
- A test for whether something is a proposition: See if it makes sense when prefixed with the text

  “It is true that...”

Propositions?

- Twice two is four.
- The square root of two is not a rational number.
- 5 > 6
- Pigs can fly.
- The first *Beatles* album is a classic.
- The second *Pirates of the Carribean* film is rubbish
- Please write a specimen of your signature in the space provided.
- Would you believe you're standing in a place where Madonna once stood?
- That's true.
- I promise not to peep.
- Adultery is wicked.

Propositions Have Structure

- Most propositions have *structure*: they are *complex*
- The most common and simplest form of structure is the use of *Boolean connectives*
- Examples:
  - it is raining <u>and</u> I am bored
  - either Liverpool win this weekend <u>or</u> I leave the country
  - <u>if</u> you go out of that door <u>then</u> we are finished
  - the economy is <u>not</u> doing well

Clarifying the Structure of Complex Propositions

- If we have a complex proposition like
  - either Liverpool win this weekend <u>or</u> I leave the country
- then we can emphasise the logical structure by systematically rewriting atomic propositions with *proposition letters*
  - P = “Liverpool win this weekend”
  - Q = “I leave the country”
  - $R = P \vee Q$

The Boolean Connectives

AND	conjunction	$\wedge$
OR	disjunction	$\vee$
NOT	negation	$\neg$
IMPLIES	implication	$\Rightarrow$
IF, AND ONLY IF	equivalence	$\Leftrightarrow$

(there are others, but these are the common ones,
and all the others can be defined in terms of these...
in fact, just $\neg$ and $\vee$ is enough to define every one)

Negation

- A complex proposition ¬*P* (read as "not P") is true under an interpretation exactly when *P* is false under that interpretation
- We can summarise the behaviour of ¬ in the following *truth table*

P	¬*P*
false	true
true	false

Disjunction

- A complex proposition $P \vee Q$ (read as "P or Q") is <u>true</u> under an interpretation exactly when either P is <u>true</u> or Q is <u>true</u> or both are <u>true</u>

P	Q	$P \vee Q$
<u>false</u>	<u>false</u>	<u>false</u>
<u>false</u>	<u>true</u>	<u>true</u>
<u>true</u>	<u>false</u>	<u>true</u>
<u>true</u>	<u>true</u>	<u>true</u>

Conjunction

- A complex proposition $P \wedge Q$ (read as "P and Q") is <u>true</u> under an interpretation exactly when *both* P is <u>true</u> *and* Q is <u>true</u>

P	Q	$P \wedge Q$
<u>false</u>	<u>false</u>	<u>false</u>
<u>false</u>	<u>true</u>	<u>false</u>
<u>true</u>	<u>false</u>	<u>false</u>
<u>true</u>	<u>true</u>	<u>true</u>

If, and only if (a.k.a. Equivalence)

- A complex proposition $P \Leftrightarrow Q$ (read as "P iff Q") is <u>true</u> under an interpretation exactly when P and Q have the *same* truth value

P	Q	$P \Leftrightarrow Q$
<u>false</u>	<u>false</u>	<u>true</u>
<u>false</u>	<u>true</u>	<u>false</u>
<u>true</u>	<u>false</u>	<u>false</u>
<u>true</u>	<u>true</u>	<u>true</u>

Implication

- By *far* the hardest to understand!
- A complex proposition $P \Rightarrow Q$ (read as "P implies Q") is true if, when P is true then Q is true also
  - If P is false then we say that $P \Rightarrow Q$ is true by definition

P	Q	$P \Rightarrow Q$
false	false	true
false	true	true
true	false	false
true	true	true

Implication

- In a proposition $P \Rightarrow Q$ we call
  - P the ***antecedent***
  - Q the ***consequent***
- Why the weird reading of implication makes sense:
  - Consider: “If 2+2=5 then I’m a Martian”
  - Antecedent is “2+2=5”
  - Consequent is “I’m a Martian”
  - Clearly the antecedent is false, so there is no suggestion I’m a Martian -- the overall proposition is true

Tautologies and Contradictions

- How can we tell whether a proposition is true or false?
- Usually, this depends on the how the propositions are *interpreted* -- we'll say more about this shortly
- Sometimes, though, we don't need an interpretation to know whether a proposition is true or false
- A proposition than *cannot be anything other than true* is called a ***tautology***
- A proposition than *cannot be anything other than false* is called a ***contradiction***

Some Tautologies

- Either Liverpool will win or they will not win
- If you’re old, then you’re not young
- Every woman who ever walked on Mars is a scouser

Some Contradictions

- I am young and I am old
- There is a number bigger than 5 and less than 4
- Every positive integer is exactly divisible by 2

Interpretations

- Propositions can be true *in one state of affairs and* false *in another*
- We capture the notion of a "state of affairs" in an *interpretation*
- An interpretation defines, for every *atomic* proposition, whether that proposition is true or false
- Interpretations are sometimes called "valuations" or "assignments"

An Example Interpretation

- Consider the interpretation *v1* such that:

  $v1(P) =$ <u>true</u> $v1(Q) =$ <u>true</u> *and* $v1(R) =$ <u>false</u>

- Under this interpretation,
  - the proposition $P \wedge (Q \vee R)$ is <u>true</u>
  - the proposition $P \wedge Q \wedge R$ is ...
  - the proposition $R \Rightarrow (P \vee Q)$ is ...
- We write $\boldsymbol{v} \vDash \boldsymbol{P}$ to mean ***P is true under interpretation v***

Checking for Tautologies and Consistency

- Suppose we have a complex proposition *P,* and we want to know whether *P* is consistent. How can we tell?
  - construct a truth table for *P*
  - P is consistent if *P* evaluates to <u>true</u> in any line of the truth table
- If we want to know whether *P* is a tautology:
  - construct a truth table for *P*
  - P is a tautology if *P* evaluates to <u>true</u> in *every* line of the truth table

Checking for Consistency

Consider the formula: $P \Rightarrow (Q \wedge \neg P)$

P	Q	$\neg P$	$Q \wedge \neg P$	$P \Rightarrow (Q \wedge \neg P)$
false	false	true	false	true
false	true	true	true	true
true	false	false	false	false
true	true	false	false	false

So the proposition is consistent.

Checking for Tautologies

Consider the formula: $P \Rightarrow (P \vee Q)$

P	Q	$P \vee Q$	$P \Rightarrow (P \vee Q)$
false	false	false	true
false	true	true	true
true	false	true	true
true	true	true	true

So the proposition is a tautology.

Checking for Inconsistency

Consider the formula: $\neg(P \Rightarrow (Q \Rightarrow P))$

P	Q	$Q \Rightarrow P$	$(P \Rightarrow (Q \Rightarrow P))$	$\neg(P \Rightarrow (Q \Rightarrow P))$
false	false	true	true	false
false	true	false	true	false
true	false	true	true	false
true	true	true	true	false

So the proposition is inconsistent.

Some Facts

- If P is a tautology, then $\neg P$ is inconsistent.
- If P is inconsistent, then $\neg P$ is a tautology.
- All tautologies are consistent.

Pictorially...

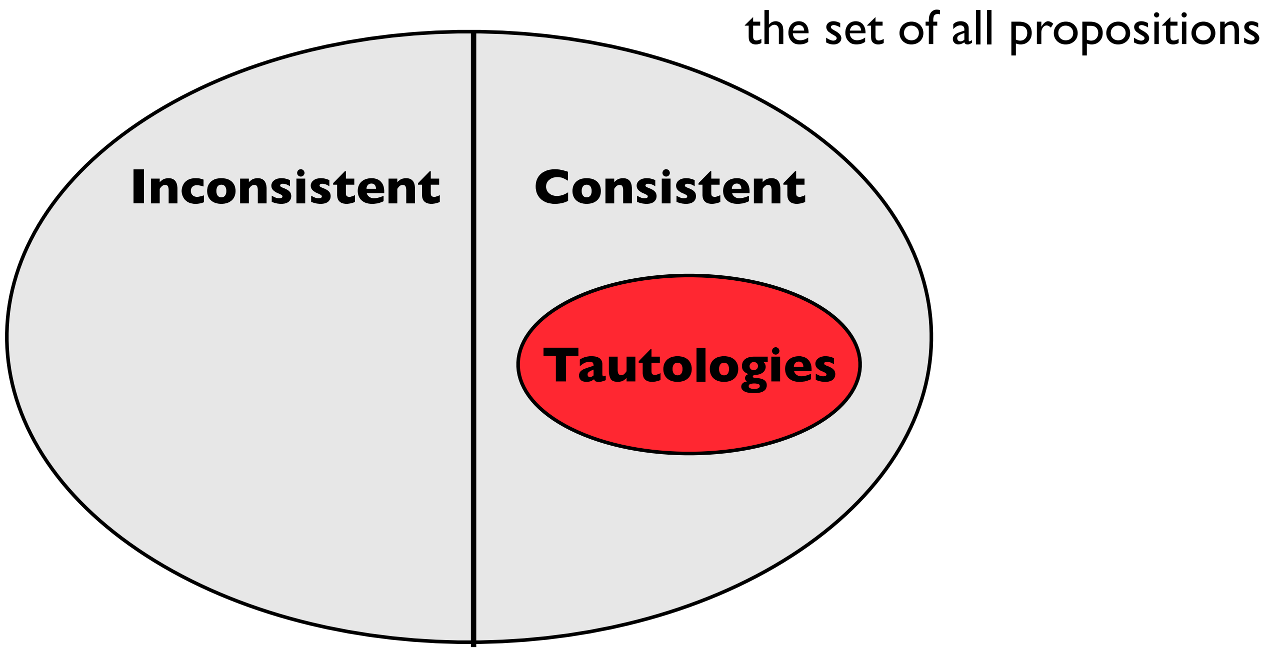

What is a Theory?

- We have been talking about the development of theories, and finally we're in a position to say what a theory *is*...
- A ***theory***, *T*, is simply a ***set of propositions***
- The elements of *T* are called the ***axioms*** of the theory
- A very simple example theory:

  $T = \{P, P \Rightarrow Q, R, R \Rightarrow S\}$

Logical Consequence

- Given a theory *T*, it is natural to ask what *predictions* the theory makes: what the ***logical consequences*** of the theory are
- We say a proposition *P* is a logical consequence of theory *T* if ***T could not be <u>true</u> without P also being <u>true</u>***
  - i.e., if *P* is <u>true</u> under all the interpretations that make *T* <u>true</u>
  - i.e., if $T \Rightarrow P$ is a tautology.
- We write $\boldsymbol{T} \vDash \boldsymbol{P}$ to mean that *P is a logical consequence of T*

Inconsistent Theories

- Suppose you have a theory *T* that is *inconsistent*
  - e.g., $T = \{P \wedge \neg P\}$
- What logical consequences does such a theory have?
- The set of interpretations that make *T* true is *empty*
  - *and so **every*** proposition is a logical consequence of T!

Properties of Logical Consequence: Monotonicity

- Suppose T is a theory and U is an extension of T
- Then if

  $T \models P$

- then

  $U \models P$

- For example let $T = \{P, P \Rightarrow Q\}$ and $U = \{P, P \Rightarrow Q, R\}$
  - clearly $T \models P$

Properties of Logical Consequence: Reflexivity

- Suppose T is a theory and P is an element of T
- Then

  $T \models P$

- For example if $T = \{P, P \Rightarrow Q\}$ then
  - $T \models P$
  - $T \models P \Rightarrow Q$

Properties of Logical Consequence: Cut

- Suppose T is a theory such that

$$T \models P \quad \textit{and} \quad T, P \models Q$$

  then $T \models Q$

- For example if $T = \{P, P \Rightarrow Q, Q \Rightarrow R\}$ then

  $T \models Q$

  $T, Q \models R$

  and so $T \models R$

Properties of Logical Consequence: Transitivity

- Suppose T is a theory
- Then if

  $T \models P$ *and*

  $P \models Q$

  then

  $T \models Q$

Proof

- We have a method to check whether a proposition P is a logical consequence of theory T:
  - use a truth table to check whether $T \Rightarrow P$ is a tautology
- But for complex theories, this isn't practical:
  - if a theory contains n propositions, then the truth table has 2^n rows
- More usually, we use a ***proof procedure,*** denoted $\vdash$

What is a Proof?

- Suppose we have a theory T, and we want to derive the consequences of that theory.
- A ***proof*** of proposition P from theory T is a sequence of propositions $Q_1, \ldots, Q_n$ such that
  - $Q_n = P$ (the last thing in the sequence is the thing to be proved)
  - For all 1, ..., n, Q_i is either:
    - an axiom (i.e., an element of T)
    - the result of applying an ***inference rule*** to elements that occurred earlier in the sequence.

Notation

- We write

  $T \vdash P$

  to mean that we can prove P from theory T

Inference Rules

- An inference rule is a rule that allows us to “safely” derive new propositions from previously proved propositions
- An inference rule is ***sound*** if the derivations it allows us to make really are logical consequences
- A collection of inference rules is ***complete*** if the derivations possible are exactly the logical consequences
- We write

  $\vdash P$

  to mean that P has been proved

Inference Rules: Introducing Tautologies

- The first inference rule we consider tells us that at any stage, we can introduce a tautology into a proof
- The soundness of this is obvious: a tautology is a logical consequence of *any* theory T

Inference Rules: Modus Ponens

- The most basic proof rule is ***modus ponens***
- This rule says:
  - if you've proved $P \Rightarrow Q$ and you've proved P
  - then you can conclude Q
- Notation:
  - from $\vdash P \Rightarrow Q$ and $\vdash P$ **infer** $\vdash Q$

A Simple Proof

- Let $T = \{P, P \Rightarrow Q, Q \Rightarrow R\}$.
- Let's do the following proof:

  $T \vdash R$

A Simple Proof

1. P	(axiom of T)
2. $P \Rightarrow Q$	(axiom of T)
3. Q	(1, 2, using modus ponens)
4. $Q \Rightarrow R$	(axiom of T)
5. R	(3, 4, modus ponens)

Inference Rules: And Elimination

- This rule says:
  - if you've proved $P \wedge Q$
  - then you can infer P
  - (and you can infer Q)
- Notation:
  - from $\vdash P \wedge Q$ infer $\vdash Q$

Inference Rules: And Introduction

- This rule says:
  - if you've proved P
  - and you've proved Q
  - then you can infer $P \wedge Q$ (as well as $Q \wedge P$)
- Notation:
  - from $\vdash P$ and $\vdash Q$ infer $\vdash P \wedge Q$

Another Simple Proof

Let $T = \{P \wedge R, P \Rightarrow Q, Q \wedge R \Rightarrow S\}$. We will prove $T \vdash S$

1. $P \wedge R$	(axiom)
2. P	(1, and elimination)
3. $P \Rightarrow Q$	(axiom)
4. Q	(2, 3, modus ponens)
5. R	(1, and elimination)
6. $Q \wedge R$	(2, 5, and introduction)
7. $Q \wedge R \Rightarrow S$	(axiom)
8. S	(7, 6, modus ponens)

Inference Rules: Or Elimination

- This rule says:
  - if you've proved $P \vee Q$
  - and you've proved $\neg P$
  - then you can infer Q
- Notation:
  - from $\vdash P \vee Q$ and $\vdash \neg P$ infer $\vdash Q$

Inference Rules: Or Introduction

- This rule says:
  - if you've proved P
  - then you can infer $P \vee Q$ for ***any*** proposition Q
- Notation:
  - from $\vdash P$ infer $\vdash P \vee Q$ for any Q

Inference Rules: Reductio ad Absurdum

- This one may seem a bit weird...
- The idea is to start out by *assuming* some proposition P is true, and then see if that leads us to a contradiction
  - if it does, then we can conclude that P must be false: $\neg P$
- Formally:
  - if $P \vdash C$, where C is a contradiction, we can infer $\vdash \neg P$

A Simple Proof with Reductio ad Absurdum

- We will prove $\vdash \neg(P \wedge \neg P)$

  1. $P \wedge \neg P$ (hypothesis)
  2. $\neg(P \wedge \neg P)$ (1, reductio ad absurdum)

- The first line was an *assumption* -- *"suppose* $(P \wedge \neg P)$ were true" -- but since it is a contradiction, immediately allows us to obtain that it cannot be true!

Inference Rules: Introducing Implications

- This is another rule that relies on proving something from an assumption
- We start by assuming P is true, and try to derive Q
  - if we can, then we can conclude that $P \Rightarrow Q$
- Formally:
  - if $P \vdash Q$, then infer that $\vdash P \Rightarrow Q$