cis32-ai — lecture # 19 — wed-5-apr-2006

today's topics:

- finishing propositional logic: proof theory

- predicate logic

- what were those Greek letters from last time?!?!?

$$\phi \quad \text{phi}$$
$$\Phi \quad \text{capital Phi}$$
$$\psi \quad \text{psi}$$
$$\pi \quad \text{pi}$$
$$\tau \quad \text{tau}$$
$$\chi \quad \text{chi}$$

Proof Theory

- What is logic used for? A number of things, but most importantly, it is a *language for representing the properties of things*.

- But also, we hope it will give us a *method for establishing the properties of things*.

- To see how logic may be used to establish the properties of things, it helps to look at its history.

- Logic was originally developed to make the notion of an *argument* precise.
  (We do *not* mean argument as in fighting here!)

- Here is a classic argument:

$$\frac{\text{All men are mortal}}{\text{Socrates is a man}}$$
$$\text{Socrates is mortal}$$

- This example serves to illustrate a number of features of arguments:

  - The argument has a number of *premises* — these are the statements that appear *before* the horizontal line;

  - The argument has a *conclusion* — this is the statement that appears *after* the horizontal line;

  - The argument has the form

    If
    you accept that
    the premises are true
    then
    you must accept that
    the conclusion is true.

- In mathematics, we are concerned with when arguments are *sound*.

- To formalise the notion of a sound argument, we need some extra terminology...

- **Definition:** If $\phi \in$ *WFF*, then:

  1. if there is *some* interpretation π such that

  $$\pi \models \phi$$

  then ϕ is said to be *satisfiable*, otherwise ϕ is *unsatisfiable*.

  2. if

  $$\pi \models \phi$$

  for *all* interpretations π, then ϕ is said to be *valid*.

- Valid formulae of propositional logic are called *tautologies*.

- Theorem:

  1. If ϕ is a valid formula, then $\neg\phi$ is unsatisfiable;
  2. If $\neg\phi$ is unsatisfiable, then ϕ is valid.

- We indicate that a formula ϕ is valid by writing

$$\models \phi.$$

- We can now define the *logical consequence*.

- Definition: If
$$\{\phi_1, \ldots, \phi_n, \phi\} \subseteq \text{WFF}$$
then ϕ is said to be a *logical consequence* of $\{\phi_1, \ldots, \phi_n\}$ iff ϕ is satisfied by all interpretations that satisfy
$$\phi_1 \wedge \cdots \wedge \phi_n.$$

- We indicate that ϕ is a logical consequence of $\phi_1, \ldots, \phi_n$ by writing

$$\{\phi_1, \ldots, \phi_n\} \models \phi.$$

- An expression like this is called a *semantic sequent*.

- Theorem:
$$\{\phi_1, \ldots, \phi_n\} \models \phi.$$

  iff

$$\models (\phi_1 \wedge \cdots \wedge \phi_n) \Rightarrow \phi.$$

- So we have a method for determining whether ϕ is a logical consequence of $\phi_1, \ldots \phi_n$: we use a truth table to see whether $\phi_1 \wedge \cdots \wedge \phi_n \Rightarrow \phi$ is a tautology. If it is, then ϕ is a logical consequence of $\phi_1, \ldots, \phi_n$.

- *Our main concern in proof theory is thus to have a technique for determining whether a given formula is valid, as this will then give us a technique for determining whether some formula is a logical consequence of some others.*

- EXAMPLE. Show that
$$p \wedge q \models p \vee q.$$
To do this, we construct a truth-table for
$$(p \wedge q) \Rightarrow (p \vee q).$$
Here it is:

p	q	(1) $p \wedge q$	(2) $p \vee q$	$(1) \Rightarrow (2)$
F	F	F	F	T
F	T	F	T	T
T	F	F	T	T
T	T	T	T	T

Since
$$(p \wedge q) \Rightarrow (p \vee q).$$
is true under every interpretation, we have that $p \vee q$ is a logical consequence of $p \wedge q$.

- The notion of logical consequence we have defined above is acceptable for a *definition* of a sound argument, but is not very helpful for checking whether a particular argument is sound or not.

- The problem is that we must look at all the possible interpretations of the primitive propositions. While this is acceptable for, say, 4 primitive propositions, it will clearly be unacceptable for 100 propositions, as it would mean checking 2^{100} interpretations.

  (Moreover, for first-order logic, there will be an *infinite* number of such interpretations.)

- What we require instead is an alternative version of logical consequence, that does not involve this kind of checking. This leads us to the idea of *syntactic* proof.

'Syntactic' Proof

- The idea of syntactic proof is to replace the semantic checking to determine whether a formula is valid by a procedure that involves purely *syntactic* manipulation.

- The kinds of techniques that we shall use are similar to those that we use when solving problems in algebra.

- The basic idea is that to show that ϕ is a logical consequence of $\phi_1, \ldots, \phi_n$, we use a set of *rules* to manipulate formulae.

  If we can derive ϕ from $\phi_1, \ldots, \phi_n$ by using these rules, then ϕ is said to be *proved* from $\phi_1, \ldots, \phi_n$, which we indicate by writing

  $$\phi_1, \ldots, \phi_n \vdash \phi.$$

- The symbol $\vdash$ is called the *syntactic turnstile*.

- An expression of the form
  $$\phi_1, \ldots, \phi_n \vdash \phi.$$
  is called a *syntactic sequent*.

- A rule has the general form:
  $$\frac{\vdash \phi_1; \cdots; \vdash \phi_n}{\vdash \phi} \quad \text{rule name}$$

  Such a rule is read:

  If
     $\phi_1, \ldots, \phi_n$ are proved
  then
     ϕ is proved.

- EXAMPLE. Here is an example of such a rule:

$$\frac{\vdash \phi; \vdash \psi}{\vdash \phi \wedge \psi} \quad \wedge\text{-I}$$

This rule is called *and introduction*. It says that if we have proved ϕ, and we have also proved ψ, then we can prove $\phi \wedge \psi$.

- EXAMPLE. Here is another rule:

$$\frac{\vdash \phi \wedge \psi}{\vdash \phi; \vdash \psi} \quad \wedge\text{-E}$$

This rule is called *and elimination*. It says that if we have proved $\phi \wedge \psi$, then we can prove both ϕ and ψ; it allows us to eliminate the $\wedge$ symbol from between them.

- Let us now try to define precisely what we mean by *proof*.
- **Definition:** (Proof) If

$$\{\phi_1, \ldots, \phi_m, \phi\} \subseteq \textit{WFF}$$

then there is a proof of ϕ from $\phi_1, \ldots, \phi_m$ iff there exists some sequence of formulae

$$\psi_1, \ldots, \psi_n$$

such that $\psi_n = \phi$, and each formula ψ_k, for $1 \leq k < n$ is either one of the formula $\phi_1, \ldots, \phi_m$, or else is the conclusion of a rule whose antecedents appeared earlier in the sequence.

- If there is a proof of ϕ from $\phi_1, \ldots, \phi_m$, then we indicate this by writing:

$$\phi_1, \ldots, \phi_m \vdash \phi.$$

- It should be clear that the symbols $\vdash$ and $\models$ are related. We now have to state exactly *how* they are related.
- There are two properties of $\vdash$ to consider:
  - *soundness*;
  - *completeness*.
  - Intuitively, $\vdash$ is said to be *sound* if it is correct, in that it does not let us derive something that is not true.
  - Intuitively, *completeness* means that $\vdash$ will let us prove anything that is true.

- **Definition:** (Soundness) A proof system $\vdash$ is said to be *sound* with respect to semantics $\models$ iff

$$\phi_1, \ldots, \phi_n \vdash \phi$$

implies

$$\phi_1, \ldots, \phi_n \models \phi.$$

- **Definition:** (Completeness) A proof system $\vdash$ is said to be *complete* with respect to semantics $\models$ iff

$$\phi_1, \ldots, \phi_n \models \phi$$

implies

$$\phi_1, \ldots, \phi_n \vdash \phi.$$

A Proof System

- There are many proof systems for propositional logic; we shall look at a simple one.

- First, we have an unusual rule that allows us to introduce any tautology.

$$\frac{}{\vdash \phi} \ \text{TAUT} \quad \text{if } \phi \text{ is a tautology}$$

- Because a tautology is true there is no problem bringing it into the proof.

- Next, rules for *eliminating* connectives.

$$\frac{\vdash \phi \wedge \psi}{\vdash \phi; \vdash \psi} \ \wedge\text{-E}$$

$$\frac{\begin{array}{c} \vdash \phi_1 \vee \cdots \vee \phi_n; \\ \phi_1 \vdash \phi; \\ \cdots; \\ \phi_n \vdash \phi \end{array}}{\vdash \phi} \ \vee\text{-E}$$

- An alternative $\vee$ elimination rule is:

$$\frac{\begin{array}{c} \vdash \phi \vee \psi; \\ \vdash \phi \Rightarrow \chi; \\ \vdash \psi \Rightarrow \chi \end{array}}{\vdash \chi} \ \vee\text{-E}$$

- Next, a rule called *modus ponens*, which lets us eliminate $\Rightarrow$.

$$\frac{\vdash \phi \Rightarrow \psi; \vdash \phi}{\vdash \psi} \ \Rightarrow\text{-E}$$

- Next, rules for *introducing* connectives.

$$\frac{\vdash \phi_1; \cdots; \vdash \phi_n}{\vdash \phi_1 \wedge \cdots \wedge \phi_n} \ \wedge\text{-I}$$

$$\frac{\vdash \phi_1; \cdots; \phi_n}{\vdash \phi_1 \vee \cdots \vee \phi_n} \ \vee\text{-I}$$

- We have a rule called the *deduction theorem*. This rule says that if we can prove ψ from ϕ, then we can prove that $\phi \Rightarrow \psi$.

$$\frac{\phi \vdash \psi}{\vdash \phi \Rightarrow \psi} \ \Rightarrow\text{-I}$$

- There are a whole range of other rules, which we shall not list here.

Proof Examples

- In this section, we give some examples of proofs in the propositional calculus.
- Example 1:

$p \wedge q \vdash q \wedge p$

1. $p \wedge q$ Given
2. p From 1 using $\wedge$-E
3. q 1,$\wedge$-E
4. $q \wedge p$ 2, 3, $\wedge$-I

- Example 2:

$p \wedge q \vdash p \vee q$

1. $p \wedge q$ Given
2. p 1, $\wedge$-E
3. $p \vee q$ 2, $\vee$-I

- Example 3:

$p \wedge q, p \Rightarrow r \vdash r$

1. $p \wedge q$ Given
2. p 1, $\wedge$-E
3. $p \Rightarrow r$ Given
4. r 2, 3, $\Rightarrow$-E

- Example 4:

$p \Rightarrow q, q \Rightarrow r \vdash p \Rightarrow r$

1. $p \Rightarrow q$ Given
2. $q \Rightarrow r$ Given
3. p Assumption |
4. q 1, 3, $\Rightarrow$-E |
5. r 2, 4, $\Rightarrow$-E |
6. $p \Rightarrow r$ 3, 5, $\Rightarrow$-I

- Example 5:

  $(p \wedge q) \Rightarrow r \vdash p \Rightarrow (q \Rightarrow r)$

1.	$(p \wedge q) \Rightarrow r$	Given	
2.	p	Assumption	$\vert$
3.	q	Assumption	$\vert\vert$
4.	$p \wedge q$	2, 3, $\wedge$-I	$\vert\vert$
5.	r	1, 4, $\Rightarrow$-I	$\vert\vert$
6.	$q \Rightarrow r$	3–5, $\Rightarrow$-I	$\vert$
7.	$p \Rightarrow (q \Rightarrow r)$	2–6, $\Rightarrow$-I	

- Example 6:

  $p \Rightarrow (q \Rightarrow r) \vdash (p \wedge q) \Rightarrow r$

1.	$p \Rightarrow (q \Rightarrow r)$	Given	
2.	$p \wedge q$	Assumption	$\vert$
3.	p	2, $\wedge$-E	$\vert$
4.	q	2, $\wedge$-E	$\vert$
5.	$q \Rightarrow r$	1, 3, $\Rightarrow$-E	$\vert$
6.	r	4, 5, $\Rightarrow$-E	$\vert$
7.	$(p \wedge q) \Rightarrow r$	2–6, $\Rightarrow$-I	

- Example 7:

  $p \Rightarrow q, \neg q \vdash \neg p$

1.	$p \Rightarrow q$	Given	
2.	$\neg q$	Given	
3.	p	Assumption	$\vert$
4.	q	1, 3, $\Rightarrow$-E	$\vert$
5.	$q \wedge \neg q$	2, 4, $\wedge$-I	$\vert$
6.	$\neg p$	3, 5, $\neg$-I	

- Example 8:

  $p \Rightarrow q \vdash \neg(p \wedge \neg q)$

1.	$p \Rightarrow q$	Given	
2.	$p \wedge \neg q$	Assumption	$\vert$
3.	p	2, $\wedge$-E	$\vert$
4.	$\neg q$	2, $\wedge$-E	$\vert$
5.	q	1, 3, $\Rightarrow$-E	$\vert$
6.	$q \wedge \neg q$	4, 5, $\wedge$-I	$\vert$
7.	$\neg(p \wedge \neg q)$	6, $\neg$-I	

- Example 9:

  Jim will party all night and pass AI? That must be wrong. If he works hard he won't have time to party. If he doesn't work hard he's not going to pass AI.

  Let:

  $$\begin{array}{ll} p & \text{Jim will party all night} \\ q & \text{Jim will pass AI} \\ r & \text{Jim works hard} \end{array}$$

  Formalisation of argument:

  $$r \Rightarrow \neg p, \neg r \Rightarrow \neg q \vdash \neg(p \wedge q)$$

1.	$r \Rightarrow \neg p$	Given	
2.	$\neg r \Rightarrow \neg q$	Given	
3.	$p \wedge q$	Assumption	|
4.	r	Assumption	||
5.	$\neg p$	1, 4, $\Rightarrow$-E	||
6.	p	3, $\wedge$-I	||
7.	$p \wedge \neg p$	5, 6, $\wedge$-I	||
8.	$\neg r$	4, 7, $\neg$-I	|
9.	$\neg q$	2, 9, $\Rightarrow$-E	|
10.	q	3, $\wedge$-E	|
11.	$q \wedge \neg q$	9, 10, $\wedge$-I	|
12.	$\neg(p \wedge q)$	3, 11, $\neg$-I	

Predicate Logic

- And now, we introduce *first-order predicate logic*.

- *More expressive* than propositional logic.

- Consider the following argument:

  - *all monitors are ready*;
  - *X12 is a monitor*;
  - therefore *X12 is ready*.

- Sense of this argument *cannot* be captured in propositional logic.

- Propositional logic is too *coarse grained* to allow us to represent and reason about this kind of statement.

Syntax

- We shall now introduce a generalisation of propositional logic called first-order logic (FOL). This new logic affords us much greater expressive power.

- **Definition:** The alphabet of FOPL contains:

  1. a set of *constants*;
  2. a set of *variables*;
  3. a set of *function symbols*;
  4. a set of *predicates symbols*;
  5. the connectives $\vee$, $\neg$;
  6. the *quantifiers* $\forall$, $\exists$, $\exists_1$;
  7. the punctuation symbols), (.

Terms

- The basic components of FOL are called *terms*.
- Essentially, a term is an object that *denotes* some object other than $\top$ or $\bot$.
- The simplest kind of term is a *constant*.
- A value such as 8 is a constant.
- The *denotation* of this term is the number 8.
- Note that a constant and the number it denotes are different!
- Aliens don't write "8" for the number 8, and nor did the Romans.

- The second simplest kind of term is a *variable*.
- A variable can stand for anything in the *domain of discourse*.
- The domain of discourse (usually abbreviated to domain) is the set of all objects under consideration.
- Sometimes, we assume the set contains "everything".
- Sometimes, we explicitly *give* the set, and *state* what variables/constants can stand for.

Functions

- We can now introduce a more complex class of terms — *functions*.
- The idea of functional terms in logic is similar to the idea of a function in programming: recall that in programming, a function is a procedure that takes some arguments, and *returns a value*.
  In C:

  ```
  T myfunction( T1 a1, ..., Tn an ) {
    ...
  }
  ```

  this function takes n arguments; the first is of type T1, the second is of type T2, and so on. The function returns a value of type T.
- In FOL, we have a set of *function symbols*; each symbol corresponds to a particular function. (It denotes some function.)

- Each function symbol is associated with a number called its *arity*. This is just the number of arguments it takes.
- A *functional term* is built up by *applying* a function symbol to the appropriate number of terms.
- Formally . . .
  Definition: Let f be an arbitrary function symbol of arity n. Also, let $\tau_1, \ldots, \tau_n$ be terms. Then

  $$f(\tau_1, \ldots, \tau_n)$$

  is a functional term.

- All this sounds complicated, but isn't. Consider a function $plus$, which takes just two arguments, each of which is a number, and returns the first number added to the second. Then:
  - $plus(2, 3)$ is an acceptable functional term;
  - $plus(0, 1)$ is acceptable;
  - $plus(plus(1, 2), 4)$ is acceptable;
  - $plus(plus(plus(0, 1), 2), 4)$ is acceptable;

- In maths, we have many functions; the obvious ones are
$$+ \quad - \quad / \quad * \quad \sqrt{\quad} \quad \sin \quad \cos \quad \ldots$$
- The fact that we write
$$2 + 3$$
instead of something like
$$plus(2, 3)$$
is just convention, and is not relevant from the point of view of logic; all these are functions in exactly the way we have defined.

- Using functions, constants, and variables, we can build up *expressions*, e.g.:
$$(x + 3) * \sin 90$$
(which might just as well be written
$$times(plus(x, 3), sin(90))$$
for all it matters.)

Predicates

- In addition to having terms, FOL has *relational operators*, which capture *relationships* between objects.
- The language of FOL contains *predicate symbols*.
- These symbols stand for *relationships between objects*.
- Each predicate symbol has an associated *arity* (number of arguments).
- **Definition:** Let P be a predicate symbol of arity n, and $\tau_1, \ldots, \tau_n$ are terms. Then
$$P(\tau_1, \ldots, \tau_n)$$
is a predicate, which will either be $\top$ or $\bot$ under some interpretation.

- EXAMPLE. Let gt be a predicate symbol with the intended interpretation 'greater than'. It takes two arguments, each of which is a natural number.

  Then:

  – $gt(4, 3)$ is a predicate, which evaluates to $\top$;

  – $gt(3, 4)$ is a predicate, which evaluates to $\bot$.

- The following are standard mathematical predicate symbols:

$$> \: < \: = \: \geq \: \leq \: \neq \: \ldots$$

- The fact that we are normally write $x > y$ instead of $gt(x, y)$ is just convention.

- We can build up more complex predicates using the connectives of propositional logic:

$$(2 > 3) \land (6 = 7) \lor (\sqrt{4} = 2)$$

- So a predicate just expresses a relationship between some values.

- What happens if a predicate contains *variables*: can we tell if it is true or false?

  Not usually; we need to know an *interpretation* for the variables.

- A predicate that contains no variables is a proposition.

- Predicates of arity 1 are called *properties*.

- EXAMPLE. The following are properties:

$$Man(x)$$
$$Mortal(x)$$
$$Malfunctioning(x).$$

- We interpret $P(x)$ as saying x is in the set P.

- Predicate that have arity 0 (i.e., take no arguments) are called *primitive propositions*.

  These are identical to the primitive propositions we saw in propositional logic.

Quantifiers

- We now come to the central part of first order logic: *quantification*.

- Consider trying to represent the following statements:

  – *all men have a mother*;

  – *every positive integer has a prime factor*.

- We can't represent these using the apparatus we've got so far; we need *quantifiers*.

- We use three quantifers:

  $\forall$ — *the universal quantifier*;
  is read 'for all...'

  $\exists$ — *the existential quantifier*;
  is read 'there exists...'

  $\exists_1$ — *the unique quantifier*;
  is read 'there exists a unique...'

- The simplest form of quantified formula is as follows:

$$quantifier\ variable\ \cdot predicate$$

where

  – $quantifier$ is one of $\forall$, $\exists$, $\exists_1$;
  – $variable$ is a variable;
  – and $predicate$ is a predicate.

Examples

- $\forall x \cdot Man(x) \Rightarrow Mortal(x)$
  'For all x, if x is a man, then x is mortal.'
  (i.e. all men are mortal)
- $\forall x \cdot Man(x) \Rightarrow \exists_1 y \cdot Woman(y) \wedge MotherOf(x, y)$
  'For all x, if x is a man, then there exists exactly one y such that y is a woman and the mother of x is y.'
  (i.e., every man has exactly one mother).

- $\exists m \cdot Monitor(m) \wedge MonitorState(m, ready)$
  'There exists a monitor that is in a ready state.'
- $\forall r \cdot Reactor(r) \Rightarrow \exists_1 t \cdot (100 \leq t \leq 1000) \wedge temp(r) = t$
  'Every reactor will have a temperature in the range 100 to 1000.'

- $\exists n \cdot posInt(n) \land n = (n * n)$

  "Some positive integer is equal to its own square."
- $\exists c \cdot EUcountry(c) \land Borders(c, Albania)$

  "Some EU country borders Albania."
- $\forall m, n \cdot Person(m) \land Person(n) \Rightarrow \neg Superior(m, n)$

  "No person is superior to another."
- $\forall m \cdot Person(m) \Rightarrow \neg \exists n \cdot Person(n) \land Superior(m, n)$

  (same as previous)

Domains & Interpretations

- Suppose we have a formula $\forall x \cdot P(x)$.

  What does x *range over*?

  Physical objects, numbers, people, times, . . . ?
- Depends on the *domain* that we intend.
- Often, we *name* a domain to make our intended interpretation clear.

- Suppose our intended interpretation is the positive integers. Suppose $>, +, *, \ldots$ have the usual mathematical interpretation.
- Is this formula *satisfiable* under this interpretation?

$$\exists n \cdot n = (n * n)$$

- Now suppose that our domain is all living people, and that $*$ means "is the child of".
- Is the formula satisfiable under this interpretation?

Comments

- Note that universal quantification is similar to *conjunction*.

  Suppose the domain is the numbers $\{2, 4, 6\}$. Then

$$\forall n \cdot Even(n)$$

  is the same as

$$Even(2) \land Even(4) \land Even(6).$$

- Existential quantification is similar to *disjunction*. Thus with the same domain,

$$\exists n \cdot Even(n)$$

  is the same as

$$Even(2) \lor Even(4) \lor Even(6).$$

- The universal and existential quantifiers are in fact *duals* of each other:

$$\forall x \cdot P(x) \;\Leftrightarrow\; \neg \exists x \cdot \neg P(x)$$

Saying that everything has some property is the same as saying that there is nothing that does not have the property.

$$\exists x \cdot P(x) \;\Leftrightarrow\; \neg \forall x \cdot \neg P(x)$$

Saying that there is something that has the property is the same as saying that its not the case that everything doesn't have the property.

Decidability

- In propositional logic, we saw that some formulae were tautologies — they had the property of being true under all interpretations.

- We also saw that there was a procedure which could be used to tell whether any formula was a tautology — this procedure was the truth-table method.

- A formula of FOL that is true under all interpretations is said to be *valid*.

- So in theory we could check for validity by writing down all the possible interpretations and looking to see whether the formula is true or not.

- Unfortuately in general we can't use this method.

- Consider the formula:

$$\forall n \cdot Even(n) \Rightarrow \neg Odd(n)$$

- There are an infinite number of interpretations.

- Is there any other procedure that we can use, that will be guaranteed to tell us, in a finite amount of time, whether a FOL formula is, or is not, valid?

- The answer is *no*.

- FOL is for this reason said to be *undecidable*.

Proof in FOL

- Proof in FOL is similar to propositional logic (PL); we just need an extra set of rules, to deal with the quantifiers.

- FOL *inherits* all the rules of PL.

- To understand FOL proof rules, need to understand *substitution*.

- The most obvious rule, for ∀-E.

  Tells us that if everything in the domain has some property, then we can infer that any *particular* individual has the property.

$$\frac{\vdash \forall x \cdot \phi(x);}{\vdash \phi(a)} \quad \text{∀-E} \quad \text{for any } a \text{ in the domain}$$

  Going from *general* to *specific*.

- Example 1.

  Let's use $\forall$-E to get the Socrates example out of the way.

$$Man(s); \forall x \cdot Man(x) \Rightarrow Mortal(x)$$
$$\vdash Mortal(s)$$

1.	$Man(s)$	Given
2.	$\forall x \cdot Man(x) \Rightarrow Mortal(x)$	Given
3.	$Man(s) \Rightarrow Mortal(s)$	2, $\forall$-E
4.	$Mortal(s)$	1, 3, $\Rightarrow$-E

- Existential Introduction Rule 1 ($\exists$-I(1)).

- We can also go from the general to the slightly less specific!

$$\frac{\vdash \forall x \cdot \phi(x);}{\vdash \exists x \cdot \phi(x)} \;\; \exists\text{-I}(1) \;\; \text{if domain not empty}$$

Note the *side condition*.

The $\exists$ quantifier *asserts the existence* of at least one object.

The $\forall$ quantifier does not.

- Existential Introduction Rule 2 ($\exists$-I(2)).

- We can also go from the very specific to less specific.

$$\frac{\vdash \phi(a);}{\vdash \exists x \cdot \phi(x)} \;\; \exists\text{-I}(2)$$

- In other words once we have a concrete example, we can infer there exists something with the property of that example.

- We often informally make use of arguments along the lines...

  1. We know somebody is the murderer.
  2. Call this person a.
  3. ...

  (Here, a is called a *Skolem constant*.)

- We have a rule which allows this, but we have to be careful how we use it!

$$\frac{\vdash \exists x \cdot \phi(x);}{\vdash \phi(a)} \;\; \exists\text{-E} \;\; a \text{ doesn't occur elsewhere}$$

- Here is an *invalid* use of this rule:

$$
\begin{array}{lll}
1. & \exists x \cdot Boring(x) & \text{Given} \\
2. & Lecture(AI) & \text{Given} \\
3. & Boring(AI) & 1, \exists\text{-E}
\end{array}
$$

- (The conclusion may be true, the argument isn't sound.)

- Example 2:

  1. Everybody is either happy or rich.
  2. Simon is not rich.
  3. Therefore, Simon is happy.

  Predicates:

  – $H(x)$ means x is happy;
  – $R(x)$ means x is rich.

- Formalisation:

$$\forall x.H(x) \vee R(x); \neg R(Simon) \vdash H(Simon)$$

$$
\begin{array}{lll}
1. & \forall x.H(x) \vee R(x) & \text{Given} \\
2. & \neg R(Simon) & \text{Given} \\
3. & H(Simon) \vee R(Simon) & 1, \forall\text{-E} \\
4. & \neg H(Simon) \Rightarrow R(Simon) & 3, \text{defn} \Rightarrow \\
5. & \neg H(Simon) & \text{Assumption} \mid \\
6. & R(Simon) & 4, 5, \Rightarrow\text{-E} \quad \mid \\
7. & R(Simon) \wedge \neg R(Simon) & 2, 6, \wedge\text{-I} \quad \mid \\
8. & \neg\neg H(Simon) & 5, 7, \neg\text{-I} \\
9. & H(Simon) \Leftrightarrow \neg\neg H(Simon) & \text{PL axiom} \\
10. & (H(Simon) \Rightarrow \neg\neg H(Simon)) & \\
 & \quad \wedge(\neg\neg H(Simon) \Rightarrow H(Simon)) & 9, \text{defn} \Leftrightarrow \\
11. & \neg\neg H(Simon) \Rightarrow H(Simon) & 10, \wedge\text{-E} \\
12. & H(Simon) & 8, 11, \Rightarrow\text{-E}
\end{array}
$$

Summary

- This lecture looked at predicate (or first order) logic.
- Predicate logic is a generalisation of propositional logic.
- The generalisation requires the use of quantifiers, and these need special rules for handling them when doing inference.
- We looked at how the proof rules for propositional logic need to be extended to handle quantifiers.