Logic: Propositional Logic (Part I)

Alessandro Artale

Free University of Bozen-Bolzano
Faculty of Computer Science
http://www.inf.unibz.it/˜artale

Descrete Mathematics and Logic — BSc course

Thanks to Prof. Enrico Franconi for provoding the slides

Knowledge bases

Inference engine	$\longleftarrow$ domain-independent algorithms
Knowledge base	$\longleftarrow$ domain-specific content

- Knowledge base = set of *sentences* in a *formal* language = logical *theory*
- Declarative approach to build an intelligent agent: TELL him what he needs to know
- Then he can ASK himself what to do—answers should follow from the KB
- Agents can be viewed at:
  - the *knowledge level*—what they know, regardless of how it is implemented;
  - or at the *implementation level*—data structures in KB and algorithms that manipulate them

Logic in general

Logic is a formal language for representing information such that conclusions can be drawn.

- Syntax defines the sentences in the language
- Semantics define the “meaning” of sentences; i.e., define *truth* of a sentence in a world
- E.g., the language of arithmetic

  $x+2 \geq y$ is a sentence; $x2+y>$ is not a sentence
  $x+2 \geq y$ is true iff the number $x+2$ is no less than the number y
  $x+2 \geq y$ is true in a world where $x=7,\ y=1$
  $x+2 \geq y$ is false in a world where $x=0,\ y=6$
  $x+2 \geq x+1$ is true in every world.

The *one and only* Logic?

- Logics of higher order
- Modal logics
  - epistemic
  - temporal and spatial
  - ...
- Description logic
- Non-monotonic logic
- Intuitionistic logic
- ...

But: There are “standard approaches”

$\rightsquigarrow$ propositional and predicate logic

Reasoning: Entailment – Logical Implication

$$\boxed{KB \models \alpha}$$

- Knowledge base KB entails (or, logically implies) sentence α

  if and only if

  α is true in all worlds where KB is true
- E.g., the KB containing "Manchester United won" and "Manchester City won"
  entails "Either Manchester United won or Manchester City won"

Models

Semantics in Logic is in terms of Models: structured worlds with respect to which truth of sentences can be evaluated.

- We say M is a model of a sentence α if α is true in M.
- $\mathcal{M}(\alpha)$ is the set of all models of α
- **Semantics of Entailment**: $KB \models \alpha$ if and only if $\mathcal{M}(KB) \subseteq \mathcal{M}(\alpha)$
- E.g. $KB = \{\text{United won, City won}\}$

  $$\alpha = \text{City won}$$

  or

  $$\alpha = \text{Manchester won}$$

  or

  $$\alpha = \text{either City or Manchester won}$$

Reasoning: Inference – Deduction – Derivation

$$\boxed{KB \vdash_i \alpha}$$

- $KB \vdash_i \alpha$ = sentence α can be inferred (or, derived or deduced) from KB by **procedure** i
- It refers to an **algorithmic procedure** that manipulate sentences in the input KB to produce α as an output
- Soundness: i is sound if

  whenever $KB \vdash_i \alpha$, it is also true that $KB \models \alpha$
- Completeness: i is complete if

  whenever $KB \models \alpha$, it is also true that $KB \vdash_i \alpha$
- Ideally a logic must be expressive enough to say almost anything of interest, and equipped with a sound and complete inference procedure.

Reasoning: Entailment Vs. Inference

We are interested in the questions:

- when a statement is **entailed** by a set of statements, in symbols: $\Theta \models \phi$,
- can we define **inference**, in symbols: $\Theta \vdash_i \phi$, in such a way that inference and entailment coincide?
- Formally, we are looking for an inference procedures, $\vdash_i$, such that:

$$\Theta \vdash_i \phi \quad \text{iff} \quad \Theta \models \phi$$

- If this is the case, then the inference procedure is said to be Sound and Complete.

Propositional Logics: Basic Ideas

The elementary building blocks of propositional logic are *atomic statements* that cannot be decomposed any further: propositions. E.g.,

- "The block is red"
- "One plus one equals two"
- "It is raining"

Using logical connectives "and", "or", "not", we can build **propositional formulas**.

Syntax of Propositional Logic

Countable alphabet Σ of **atomic propositions**: $a, b, c, \ldots$.

Propositional formulas:

$$
\begin{array}{lcll}
\phi, \psi & \longrightarrow & a & \textit{atomic formula} \\
 & | & \bot & \textit{false} \\
 & | & \top & \textit{true} \\
 & | & \neg\phi & \textit{negation} \\
 & | & \phi \wedge \psi & \textit{conjunction} \\
 & | & \phi \vee \psi & \textit{disjunction} \\
 & | & \phi \rightarrow \psi & \textit{implication} \\
 & | & \phi \leftrightarrow \psi & \textit{equivalence}
\end{array}
$$

- **Atom:** atomic formula
- **Literal:** (negated) atomic formula
- **Clause:** disjunction of literals

Syntax of Propositional Logic

Operator Precedence: from high to low is: $\neg$, $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$

Examples of Formulas

- $\neg p \vee q \vee \neg r$
- $p \wedge q \rightarrow a$
- $(p \wedge \neg q) \vee (a \wedge b) \leftrightarrow c \vee d$
- $\ldots$

Semantics: Intuition

- Atomic statements can be *true* `T` or *false* `F`.
- The truth value of formulas is determined by the truth values of the atoms (truth value assignment or interpretation).

Example: $(a \vee b) \wedge c$

- If a and b are false and c is true, then the formula is not true.
- Then *logical entailment* could be defined as follows:
- ϕ is logically implied by Θ, $\Theta \models \phi$, if ϕ is true in all "states of the world" in which Θ is true.

Semantics: Formally

A **truth value assignment** (or **interpretation**) of the atoms in Σ is a function $\mathcal{I}$:

$$\mathcal{I} : \Sigma \to \{\mathtt{T}, \mathtt{F}\}.$$

Note: Instead of $\mathcal{I}(a)$ we also write $a^{\mathcal{I}}$.

Definition: A formula ϕ is *satisfied* by an interpretation $\mathcal{I}$ ($\mathcal{I} \models \phi$), or is *true under* $\mathcal{I}$, if and only if:

$$\mathcal{I} \models \top, \quad \mathcal{I} \not\models \bot$$

$$\begin{array}{rcl}
\mathcal{I} \models a & \text{iff} & a^{\mathcal{I}} = \mathtt{T} \\
\mathcal{I} \models \neg\phi & \text{iff} & \mathcal{I} \not\models \phi \\
\mathcal{I} \models \phi \wedge \psi & \text{iff} & \mathcal{I} \models \phi \text{ and } \mathcal{I} \models \psi \\
\mathcal{I} \models \phi \vee \psi & \text{iff} & \mathcal{I} \models \phi \text{ or } \mathcal{I} \models \psi \\
\mathcal{I} \models \phi \to \psi & \text{iff} & \text{if } \mathcal{I} \models \phi, \text{ then } \mathcal{I} \models \psi \\
\mathcal{I} \models \phi \leftrightarrow \psi & \text{iff} & \mathcal{I} \models \phi, \text{ if and only if } \mathcal{I} \models \psi
\end{array}$$

Exercises

Let:

$$\mathcal{I}: \begin{cases} a & \mapsto & \mathtt{T} \\ b & \mapsto & \mathtt{F} \\ c & \mapsto & \mathtt{F} \\ d & \mapsto & \mathtt{T} \end{cases}$$

Check the truth value under $\mathcal{I}$ of the following formulas:

- $b \rightarrow c \vee d$
- $c \vee d \rightarrow b$
- $b \leftrightarrow c \vee d$
- $((a \vee b) \leftrightarrow (c \vee d)) \wedge (\neg(a \wedge b) \vee (c \wedge \neg d))$

Exercises

- Find an interpretation and a formula such that the formula is true in that interpretation (or: the interpretation satisfies the formula).
- Find an interpretation and a formula such that the formula is not true in that interpretation (or: the interpretation does not satisfy the formula).
- Find a formula which can't be true in any interpretation (or: no interpretation can satisfy the formula).

An interpretation $\mathcal{I}$ is said to be a **model** of ϕ when:

$$\mathcal{I} \models \phi$$

Definition. A formula ϕ is

- **satisfiable**, if there is some $\mathcal{I}$ that is a model of ϕ,
- **unsatisfiable**, if ϕ is not satisfiable,
- **valid** (i.e., a **tautology**), if every $\mathcal{I}$ is a model of ϕ.
- **falsifiable**, if there is some $\mathcal{I}$ that does not satisfy ϕ,
- Two formulas are **logically equivalent** ($\phi \equiv \psi$), if **for all** $\mathcal{I}$:

$$\mathcal{I} \models \phi \text{ iff } \mathcal{I} \models \psi$$

Truth Tables

A truth table is a convenient format for displaying the semantics of a formula by showing its truth value for *every possible* interpretation of the formula.

Definition. Let ϕ be a formula with n atoms. A *truth table* is a table with $n+1$ columns and 2^n rows. There is a column for each atom in ϕ, plus a column for the formula ϕ. The first n columns specify the interpretation $\mathcal{I}$ that maps atoms in ϕ to $\{\mathtt{T}, \mathtt{F}\}$. The last column shows the truth value of ϕ under $\mathcal{I}$.

Example: Truth Table for the formula $A \wedge B$:

A	B	$A \wedge B$
False	*False*	*False*
False	*True*	*False*
True	*False*	*False*
True	*True*	*True*

Exercise

Satisfiable, tautology?

$$(((a \wedge b) \leftrightarrow a) \rightarrow b)$$
$$((\neg\phi \rightarrow \neg\psi) \rightarrow (\psi \rightarrow \phi))$$
$$(a \vee b \vee \neg c) \wedge (\neg a \vee \neg b \vee d) \wedge (\neg a \vee b \vee \neg d)$$

Equivalent?

$$\begin{array}{rcl} (\phi \vee (\psi \wedge \chi)) & \equiv & ((\phi \vee \psi) \wedge (\psi \wedge \chi)) \\ \neg(\phi \vee \psi) & \equiv & \neg\phi \wedge \neg\psi \end{array}$$

Try to use **truth tables** to support your conclusions.

Important Facts

Theorem.

- ϕ is valid iff $\neg\phi$ is unsatisfiable.
- ϕ is unsatisfiable iff $\neg\phi$ is valid.
- ϕ is satisfiable iff $\neg\phi$ is falsifiable.

Relationship between $\leftrightarrow$ and $\equiv$

Theorem.

- $\phi \equiv \psi$ iff $\phi \leftrightarrow \psi$ is a tautology.

Logical equivalence justifies substitution of one formula for another.

Substitution Theorem: If ϕ and ψ are equivalent, and χ' results from replacing ϕ in χ by ψ, then χ and χ' are equivalent.

Entailment

- Extension of the interpretation relationship to sets of formulas Θ

$$\mathcal{I} \models \Theta \quad \text{iff} \quad \mathcal{I} \models \phi \ \text{ for all } \ \phi \in \Theta$$

- Extension of the entailment relationship to sets of formulas Θ.

$$\Theta \models \phi \quad \text{iff} \quad \mathcal{I} \models \phi \ \text{ for all models } \mathcal{I} \ \text{ of } \Theta$$

Note: we want the formula ϕ to be implied by a set Θ, if ϕ is true in all models of Θ (symbolically, $\Theta \models \phi$)

Propositional inference: Truth Table method

Let $\alpha = A \lor B$ and $KB = \{(A \lor C),\ (B \lor \neg C)\}$

Is it the case that $KB \models \alpha$?

Check all possible models – α must be true wherever KB is true

A	B	C	$A \lor C$	$B \lor \neg C$	KB	α
False	*False*	*False*				
False	*False*	*True*				
False	*True*	*False*				
False	*True*	*True*				
True	*False*	*False*				
True	*False*	*True*				
True	*True*	*False*				
True	*True*	*True*				

Let $\alpha = A \vee B$ and $KB = \{(A \vee C),\ (B \vee \neg C)\}$

Is it the case that $KB \models \alpha$?

Check all possible models – α must be true wherever KB is true

A	B	C	$A \vee C$	$B \vee \neg C$	KB	α
False	*False*	*False*	*False*			
False	*False*	*True*	*True*			
False	*True*	*False*	*False*			
False	*True*	*True*	*True*			
True	*False*	*False*	*True*			
True	*False*	*True*	*True*			
True	*True*	*False*	*True*			
True	*True*	*True*	*True*			

Propositional inference: Truth Table method

Let $\alpha = A \vee B$ and $KB = \{(A \vee C),\ (B \vee \neg C)\}$

Is it the case that $KB \models \alpha$?

Check all possible models – α must be true wherever KB is true

A	B	C	$A \vee C$	$B \vee \neg C$	KB	α
False	*False*	*False*	*False*	*True*		
False	*False*	*True*	*True*	*False*		
False	*True*	*False*	*False*	*True*		
False	*True*	*True*	*True*	*True*		
True	*False*	*False*	*True*	*True*		
True	*False*	*True*	*True*	*False*		
True	*True*	*False*	*True*	*True*		
True	*True*	*True*	*True*	*True*		

Propositional inference: Truth Table method

Let $\alpha = A \vee B$ and $KB = \{(A \vee C),\ (B \vee \neg C)\}$

Is it the case that $KB \models \alpha$?

Check all possible models – α must be true wherever KB is true

A	B	C	$A \vee C$	$B \vee \neg C$	KB	α
False	*False*	*False*	*False*	*True*	*False*	
False	*False*	*True*	*True*	*False*	*False*	
False	*True*	*False*	*False*	*True*	*False*	
False	*True*	*True*	*True*	*True*	*True*	
True	*False*	*False*	*True*	*True*	*True*	
True	*False*	*True*	*True*	*False*	*False*	
True	*True*	*False*	*True*	*True*	*True*	
True	*True*	*True*	*True*	*True*	*True*	

Propositional inference: Truth Table method

Let $\alpha = A \vee B$ and $KB = \{(A \vee C),\ (B \vee \neg C)\}$

Is it the case that $KB \models \alpha$?

Check all possible models – α must be true wherever KB is true

A	B	C	$A \vee C$	$B \vee \neg C$	KB	α
False	*False*	*False*	*False*	*True*	*False*	*False*
False	*False*	*True*	*True*	*False*	*False*	*False*
False	*True*	*False*	*False*	*True*	*False*	*True*
False	*True*	*True*	*True*	*True*	*True*	*True*
True	*False*	*False*	*True*	*True*	*True*	*True*
True	*False*	*True*	*True*	*False*	*False*	*True*
True	*True*	*False*	*True*	*True*	*True*	*True*
True	*True*	*True*	*True*	*True*	*True*	*True*

Thus $KB \models \alpha$

Note. The method of truth tables is a very inefficient since we need to evaluate a formula for each of 2^n possible interpretations, where n is the number of distinct atoms in the formula.

Equivalences (I)

Commutativity	$\phi \lor \psi$	$\equiv$	$\psi \lor \phi$
	$\phi \land \psi$	$\equiv$	$\psi \land \phi$
	$\phi \leftrightarrow \psi$	$\equiv$	$\psi \leftrightarrow \phi$
Associativity	$(\phi \lor \psi) \lor \chi$	$\equiv$	$\phi \lor (\psi \lor \chi)$
	$(\phi \land \psi) \land \chi$	$\equiv$	$\phi \land (\psi \land \chi)$
	$(\phi \leftrightarrow \psi) \leftrightarrow \chi$	$\equiv$	$\phi \leftrightarrow (\psi \leftrightarrow \chi)$
Idempotence	$\phi \lor \phi$	$\equiv$	ϕ
	$\phi \land \phi$	$\equiv$	ϕ
Absorption	$\phi \lor (\phi \land \psi)$	$\equiv$	ϕ
	$\phi \land (\phi \lor \psi)$	$\equiv$	ϕ
Distributivity	$\phi \land (\psi \lor \chi)$	$\equiv$	$(\phi \land \psi) \lor (\phi \land \chi)$
	$\phi \lor (\psi \land \chi)$	$\equiv$	$(\phi \lor \psi) \land (\phi \lor \chi)$

Equivalences (II)

Implication	$\phi \rightarrow \psi$	$\equiv$	$\neg\phi \vee \psi$
Tautology	$\phi \vee \top$	$\equiv$	$\top$
	$\phi \rightarrow \top$	$\equiv$	$\top$
	$\bot \rightarrow \phi$	$\equiv$	$\top$
	$\phi \rightarrow \phi$	$\equiv$	$\top$
Unsatisfiability	$\phi \wedge \bot$	$\equiv$	$\bot$
Constants	$\phi \vee \neg\phi$	$\equiv$	$\top$
	$\phi \wedge \neg\phi$	$\equiv$	$\bot$
Neutrality	$\phi \wedge \top$	$\equiv$	ϕ
	$\phi \vee \bot$	$\equiv$	ϕ
	$\top \rightarrow \phi$	$\equiv$	ϕ
Negation	$\phi \rightarrow \bot$	$\equiv$	$\neg\phi$
Double Negation	$\neg\neg\phi$	$\equiv$	ϕ
De Morgan	$\neg(\phi \vee \psi)$	$\equiv$	$\neg\phi \wedge \neg\psi$
De Morgan	$\neg(\phi \wedge \psi)$	$\equiv$	$\neg\phi \vee \neg\psi$

Minimal set of Logical Operators

From the above presented equivalences the following follows.

Theorem. The logical operators $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$ can be defined from negation, $\neg$, and one of $\wedge, \vee, \rightarrow$.

Properties of Entailment

- $\Theta \cup \{\phi\} \models \psi$ iff $\Theta \models \phi \rightarrow \psi$
  $\boxed{\text{(Deduction Theorem)}}$
- $\Theta \cup \{\phi\} \models \neg\psi$ iff $\Theta \cup \{\psi\} \models \neg\phi$
  (Contraposition Theorem)
- $\Theta \cup \{\phi\}$ is unsatisfiable iff $\Theta \models \neg\phi$
  (Contradiction Theorem)

Inference procedures use syntactic operations on sentences, often expressed in standardized forms.

Conjunctive Normal Form (CNF)

conjunction of $\underbrace{\textit{disjunctions} \text{ of } \textit{literals}}_{\textit{clauses}}$: $\bigwedge_{i=1}^{n}(\bigvee_{j=1}^{m} l_{i,j})$

E.g., $(A \vee \neg B) \wedge (B \vee \neg C \vee \neg D)$

Disjunctive Normal Form (DNF)

disjunction of $\underbrace{\textit{conjunctions} \text{ of } \textit{literals}}_{\textit{terms}}$: $\bigvee_{i=1}^{n}(\bigwedge_{j=1}^{m} l_{i,j})$

E.g., $(A \wedge B) \vee (A \wedge \neg C) \vee (A \wedge \neg D) \vee (\neg B \wedge \neg C) \vee (\neg B \wedge \neg D)$

Normal Forms, cont.

Horn Form (restricted)

conjunction of *Horn clauses* (clauses with ≤ 1 positive literal)

E.g., $(A \vee \neg B) \wedge (B \vee \neg C \vee \neg D)$

Often written as set of implications:

$B \implies A$ and $(C \wedge D) \implies B$

Theorem For every formula, there exists an equivalent formula in CNF and one in DNF.

Why Normal Forms?

- We can transform propositional formulas, in particular, we can construct their CNF and DNF.
- DNF tells us something as to whether a formula is satisfiable. If all disjuncts contain $\bot$ or complementary literals, then no model exists. Otherwise, the formula is satisfiable.
- CNF tells us something as to whether a formula is a tautology. If all clauses (= conjuncts) contain $\top$ or complementary literals, then the formula is a tautology. Otherwise, the formula is falsifiable.

But:

- the transformation into DNF or CNF is expensive (in time/space)

Summary: important notions

- Syntax: formula, atomic formula, literal, clause
- Semantics: truth value, assignment, interpretation
- Formula satisfied by an interpretation
- Logical implication, entailment
- Satisfiability, validity, tautology, logical equivalence
- Deduction theorem, Contraposition Theorem
- Conjunctive normal form, Disjunctive Normal form, Horn form