Inference and Logic

Readings: Chapter 6 of Russell & Norvig.

Logics

A **logic** is a triple $\langle \mathcal{L}, \mathcal{S}, \mathcal{R} \rangle$ where

- $\mathcal{L}$, the logic's **language**, is a class of sentences described by a formal grammar.
- $\mathcal{S}$, the logic's **semantics** is a formal specification of how to assign *meaning* in the "real world" to the elements of $\mathcal{L}$.
- $\mathcal{R}$, the logic's **inference system**, is a set of formal derivation *rules* over $\mathcal{L}$.

There are *several* logics: propositional, first-order, higher-order, modal, temporal, intuitionistic, linear, equational, non-monotonic, fuzzy, ...

We will concentrate on *propositional logic* and *first-order logic*.

Propositional Logic

Each sentence is made of

- **propositional variables** $(A, B, \ldots, P, Q, \ldots)$
- **logical constants** $(\mathbf{True}, \mathbf{False})$.
- **logical connectives** $(\wedge, \vee, \Rightarrow, \ldots)$.

Every propositional variable stands for a basic *fact*.

Ex: I'm hungry, Apples are red, Bill and Hillary are married.

Propositional Logic

The Language

- Each propositional variable $(A, B, \ldots, P, Q, \ldots)$ is a sentence.
- Each logical constant (**True**, **False**) is a sentence.
- If φ and ψ are sentences, all of the following are also sentences.

$$(\varphi) \qquad \neg\varphi \qquad \varphi \wedge \psi \qquad \varphi \vee \psi \qquad \varphi \Rightarrow \psi \qquad \varphi \Leftrightarrow \psi$$

- Nothing else is a sentence.

The Language of Propositional Logic

Formally, it is the language generated by the following grammar.

- Symbols:
  - Propositional variables: $A, B, \ldots, P, Q, \ldots$
  - Logical constants:

True	(true)	$\wedge$	(and)	$\Rightarrow$	(implies)	$\neg$	(not)
False	(false)	$\vee$	(or)	$\Leftrightarrow$	(equivalent)		

- Grammar Rules:

$$
\begin{array}{lcl}
\mathit{Sentence} & ::= & \mathit{AtomicS} \mid \mathit{ComplexS} \\
\mathit{AtomicS} & ::= & \mathbf{True} \mid \mathbf{False} \mid A \mid B \mid \ldots \mid P \mid Q \mid \ldots \\
\mathit{ComplexS} & ::= & (\mathit{Sentence}) \mid \mathit{Sentence}\ \mathit{Connective}\ \mathit{Sentence} \mid \neg \mathit{Sentence} \\
\mathit{Connective} & ::= & \wedge \mid \vee \mid \Rightarrow \mid \Leftrightarrow
\end{array}
$$

Propositional Logic

Ontological Commitments

Propositional Logic is about *facts* in the world that are either true or false, nothing else.

Semantics of Propositional Logic

Since each propositional variable stands for a fact about the world, its meaning ranges over the Boolean values $\{True, False\}$.

Note: Do note confuse, as the textbook does, $True, False$, which are values (ie semantical entities) with $\mathbf{True}, \mathbf{False}$ which are logical constants (ie symbols).

Semantics of Propositional Logic

- The meaning (value) of **True** is always *True*. The meaning of **False** is always *False*.
- The meaning of the other sentences depends on the meaning of the propositional variables.
  - **Base cases:** Truth Tables

P	Q	$\neg P$	$P \wedge Q$	$P \vee Q$	$P \Rightarrow Q$	$P \Leftrightarrow Q$
False	*False*	*True*	*False*	*False*	*True*	*True*
False	*True*	*True*	*False*	*True*	*True*	*False*
True	*False*	*False*	*False*	*True*	*False*	*False*
True	*True*	*False*	*True*	*True*	*True*	*True*

  - **Non-base Cases:** Given by reduction to the base cases.
    Ex: the meaning of $(P \vee Q) \wedge R$ *is the same as the meaning of* $A \wedge R$ *where* A *has the same meaning as* $P \vee Q$.

The Meaning of Logical Connectives: A Warning

Disjunction

- $A \vee B$ is true when A or B or *or both* are true (inclusive or).
- $A \oplus B$ is sometimes used to mean "either A or B but not both" (exclusive or).

Implication

- $A \Rightarrow B$ does not require a causal connection between A and B.

  Ex: Sky-is-blue $\Rightarrow$ *Snow-is-white*

- When A is false, $A \Rightarrow B$ is always true *regardless* of the value of B.

  Ex: Two-equals-four $\Rightarrow$ *Apples-are-red*

- Beware of negations in implications.

  Ex: (¬*Has-blue-seal*) $\Rightarrow$ (¬*Chiquita*)

Semantics of Propositional Logic

- An assignment of Boolean values to the propositional variables of a sentence is an **interpretation** of the sentence.

P	H	$P \vee H$	$(P \vee H) \wedge \neg H$	$((P \vee H) \wedge \neg H) \Rightarrow P$
False	*False*	*False*	*False*	*True*
False	*True*	*True*	*False*	*True*
True	*False*	*True*	*True*	*True*
True	*True*	*True*	*False*	*True*

- The semantics of Propositional logic is **compositional**:
  - The meaning of a sentence is given recursively in terms of the meaning of the sentence's components (all the way down to its propositional variables).

Semantics of Propositional Logic

The meaning of a sentence in general depends on its interpretation. Some sentences, however, have always the same meaning.

P	H	$P \vee H$	$(P \vee H) \wedge \neg H$	$((P \vee H) \wedge \neg H) \Rightarrow P$
False	*False*	*False*	*False*	*True*
False	*True*	*True*	*False*	*True*
True	*False*	*True*	*True*	*True*
True	*True*	*True*	*False*	*True*

A sentence is

- **satisfiable** if it is true in *some* interpretation,
- **valid** if it is true in *every* possible interpretation.

Entailment in Propositional Logic

Given a set Γ of sentences and a sentence φ, we write

$$\Gamma \models \varphi$$

iff every interpretation that makes all sentences in Γ true makes φ also true.

$\Gamma \models \varphi$ is read as "Γ *entails* φ" or "φ *logically follows from* Γ."

Entailment in Propositional Logic: Examples

$$\begin{aligned}
\{A, A \Rightarrow B\} &\models B \\
\{A\} &\models A \lor B \\
\{A, B\} &\models A \land B \\
\{\} &\models A \lor \neg A \\
\{A\} &\not\models A \land B \\
\{A \lor \neg A\} &\not\models A
\end{aligned}$$

	A	B	$A \Rightarrow B$	$A \lor B$	$A \land B$	$A \lor \neg A$
1.	*False*	*False*	*True*	*False*	*False*	*True*
2.	*False*	*True*	*True*	*True*	*False*	*True*
3.	*True*	*False*	*False*	*True*	*False*	*True*
4.	*True*	*True*	*True*	*True*	*True*	*True*

Entailment in Propositional Logic

Note:

- $\Gamma \models \varphi$, for all $\varphi \in \Gamma$ (**inclusion property** of PL)
- if $\Gamma \models \varphi$, then $\Gamma' \models \varphi$ for all $\Gamma' \supseteq \Gamma$ (**monotonicity** of PL)
- φ is valid iff **True** $\models \varphi$ (also written as $\models \varphi$)
- φ is unsatisfiable iff $\varphi \models$ **False**
- $\Gamma \models \varphi$ iff the set $\Gamma \cup \{\neg\varphi\}$ is unsatisfiable

Logical Equivalence

Two sentences φ_1 and φ_2 are **logically equivalent**

$$\varphi_1 \equiv \varphi_2$$

if $\varphi_1 \models \varphi_2$ and $\varphi_2 \models \varphi_1$.

Note:

- If $\varphi_1 \equiv \varphi_2$, every interpretation of their propositional variables will assign the same Boolean value to φ_1 and φ_2.
- Implication and equivalence ($\Rightarrow, \Leftrightarrow$), which are *syntactical entities*, are intimately related to entailment and logical equivalence ($\models, \equiv$), which are *semantical notions*:

$$\varphi_1 \models \varphi_2 \text{ iff } \models \varphi_1 \Rightarrow \varphi_2$$
$$\varphi_1 \equiv \varphi_2 \text{ iff } \models \varphi_1 \Leftrightarrow \varphi_2$$

Properties of Logical Connectives

- $\wedge$ and $\vee$ are *commutative*

$$\varphi_1 \wedge \varphi_2 \equiv \varphi_2 \wedge \varphi_1$$
$$\varphi_1 \vee \varphi_2 \equiv \varphi_2 \vee \varphi_1$$

- $\wedge$ and $\vee$ are *associative*

$$\varphi_1 \wedge (\varphi_2 \wedge \varphi_3) \equiv (\varphi_1 \wedge \varphi_2) \wedge \varphi_3$$
$$\varphi_1 \vee (\varphi_2 \vee \varphi_3) \equiv (\varphi_1 \vee \varphi_2) \vee \varphi_3$$

- $\wedge$ and $\vee$ are mutually *distributive*

$$\varphi_1 \wedge (\varphi_2 \vee \varphi_3) \equiv (\varphi_1 \wedge \varphi_2) \vee (\varphi_1 \wedge \varphi_3)$$
$$\varphi_1 \vee (\varphi_2 \wedge \varphi_3) \equiv (\varphi_1 \vee \varphi_2) \wedge (\varphi_1 \vee \varphi_3)$$

- $\wedge$ and $\vee$ are related by $\neg$ (DeMorgan's Laws)

$$\neg(\varphi_1 \wedge \varphi_2) \equiv \neg\varphi_1 \vee \neg\varphi_2$$
$$\neg(\varphi_1 \vee \varphi_2) \equiv \neg\varphi_1 \wedge \neg\varphi_2$$

Properties of Logical Connectives

$\wedge$, $\Rightarrow$, and $\Leftrightarrow$ are actually redundant:

$$\begin{aligned}
\varphi_1 \wedge \varphi_2 &\equiv \neg(\neg\varphi_1 \vee \neg\varphi_2) \\
\varphi_1 \Rightarrow \varphi_2 &\equiv \neg\varphi_1 \vee \varphi_2 \\
\varphi_1 \Leftrightarrow \varphi_2 &\equiv (\varphi_1 \Rightarrow \varphi_2) \wedge (\varphi_2 \Rightarrow \varphi_1)
\end{aligned}$$

We keep them all mainly for convenience.

Exercise

Use the truth tables to prove all the logical equivalences seen so far.

Inference Systems for Propositional Logic

- In practice, an inference system $\mathcal{I}$ for PL is a procedure that *given a set* $\Gamma = \{\alpha_1, \ldots, \alpha_m\}$ *of sentences and a sentence* φ*, may reply "yes", "no", or run forever.*

- If $\mathcal{I}$ replies positively on input (Γ, φ), we say that Γ *derives* φ *in* $\mathcal{I}$,[1] and write

$$\Gamma \vdash_{\mathcal{I}} \varphi$$

- Intuitively, $\mathcal{I}$ should be such that it replies "yes" on input (Γ, φ) only if φ is in fact entailed by Γ.

[1] Or, $\mathcal{I}$ derives φ from Γ, or, φ derives from Γ in $\mathcal{I}$.

All These Fancy Symbols!

Before we continue recall that

- $A \wedge B \Rightarrow C$

  is a sentence, a bunch of *symbols* manipulated by an *inference system* $\mathcal{I}$.

- $A \wedge B \models C$

  is a mathematical abbreviation standing for the statement:
  "every interpretation that makes $A \wedge B$ true, makes C also true."

- $A \wedge B \vdash_{\mathcal{I}} C$

  is a mathematical abbreviation standing for the statement:
  "$\mathcal{I}$ returns yes on input $(A \wedge B, C)$" [C derives from $A \wedge B$ in $\mathcal{I}$].

In other words,

- $\Rightarrow$ is a formal symbol of the logic, which is used by the inference system.
- $\models$ is a shorthand we use to talk about the meaning of formal sentences.
- $\vdash_{\mathcal{I}}$ is a shorthand we use to talk about the output of the inference system $\mathcal{I}$.

The formal symbol $\Rightarrow$ and the shorthands $\models$, $\vdash_{\mathcal{I}}$ are related as follows.

- The sentence $\varphi_1 \Rightarrow \varphi_2$ is valid (always true) if and only if $\varphi_1 \models \varphi_2$.

  Example: $A \Rightarrow (A \vee B)$ *is valid and* $A \models (A \vee B)$

	A	B	$A \vee B$	$A \Rightarrow (A \vee B)$
1.	*False*	*False*	*False*	*True*
2.	*False*	*True*	*True*	*True*
3.	<u>*True*</u>	*False*	*True*	*True*
4.	<u>*True*</u>	*True*	*True*	*True*

- A **sound** inference system can derive *only* sentences that logically follow from a given set of sentences:

$$\text{if} \;\; \Gamma \vdash_{\mathcal{I}} \varphi \;\; \text{then} \;\; \Gamma \models \varphi.$$

- A **complete** inference system can derive *all* sentences that logically follow from a given set of sentences:

$$\text{if} \;\; \Gamma \models \varphi \;\; \text{then} \;\; \Gamma \vdash_{\mathcal{I}} \varphi.$$

Inference in Propositional Logic

There are two (equivalent) types of inference systems of Propositional Logic:

- one based on truth tables ($\mathcal{TT}$)
- one based on derivation rules ($\mathcal{R}$)

Truth Tables

The inference system $\mathcal{TT}$ is specified as follows:

$\{\alpha_1, \ldots, \alpha_m\} \vdash_{\mathcal{TT}} \varphi$ *iff* *all the values in the truth table of* $(\alpha_1 \wedge \cdots \wedge \alpha_m) \Rightarrow \varphi$ *are True.*

Inference by Truth Tables

- The truth-tables-based inference system is sound:

$$
\begin{array}{rll}
\alpha_1, \ldots, \alpha_m \vdash_{\mathcal{TT}} \varphi & \text{implies} & \text{truth table of } (\alpha_1 \wedge \cdots \wedge \alpha_m) \Rightarrow \varphi \text{ all true} \\
& \text{implies} & (\alpha_1 \wedge \cdots \wedge \alpha_m) \Rightarrow \varphi \text{ is valid} \\
& \text{implies} & \models (\alpha_1 \wedge \cdots \wedge \alpha_m) \Rightarrow \varphi \\
& \text{implies} & (\alpha_1 \wedge \cdots \wedge \alpha_m) \models \varphi \\
& \text{implies} & \alpha_1, \ldots, \alpha_m \models \varphi
\end{array}
$$

- It is also complete (exercise: prove it).
- Its time complexity is $O(2^n)$ where n is the number of propositional variables in $\alpha_1, \ldots, \alpha_m, \varphi$.
- We cannot hope to do better because a related, simpler problem (determining the satisfiability of a sentence) is NP-complete.
- However, really hard cases of propositional inference are somewhat rare.

Rule-Based Inference in Propositional Logic

An inference system in Propositional Logic can also be specified as a set $\mathcal{R}$ of inference (or derivation) rules.

Each rule is actually a *pattern* premises/conclusion.

A rule *applies* to Γ and *derives* φ if

- some of the sentences in Γ match with the premises of the rule and
- φ matches with the conclusion.

A rule is **sound** it the set of its premises entails its conclusion.

Rule-Based Inference in Propositional Logic

Inference Rules

- **And-Introduction**

$$\frac{\alpha \qquad \beta}{\alpha \wedge \beta}$$

- **And-Elimination**

$$\frac{\alpha \wedge \beta}{\alpha} \qquad\qquad \frac{\alpha \wedge \beta}{\beta}$$

- **Or-Introduction**

$$\frac{\alpha}{\alpha \vee \beta} \qquad\qquad \frac{\alpha}{\beta \vee \alpha}$$

Rule-Based Inference in Propositional Logic

Inference Rules (cont'd)

- **Implication-Elimination** (aka **Modus Ponens**)

$$\frac{\alpha \Rightarrow \beta \qquad \alpha}{\beta}$$

- **Unit Resolution**

$$\frac{\alpha \lor \beta \qquad \neg\beta}{\alpha}$$

- **Resolution**

$$\frac{\alpha \lor \beta \qquad \neg\beta \lor \gamma}{\alpha \lor \gamma} \quad \text{or, equivalently,} \quad \frac{\neg\alpha \Rightarrow \beta, \qquad \beta \Rightarrow \gamma}{\neg\alpha \Rightarrow \gamma}$$

Rule-Based Inference in Propositional Logic

Inference Rules (cont'd.)

- **Double-Negation-Elimination**

$$\frac{\neg\neg\alpha}{\alpha}$$

- **False-Introduction**

$$\frac{\alpha \wedge \neg\alpha}{\textbf{False}}$$

- **False-Elimination**

$$\frac{\textbf{False}}{\beta}$$

Inference by Proof

We say there is a **proof** of φ from Γ in $\mathcal{R}$ if we can derive φ by applying the rules of $\mathcal{R}$ repeatedly to Γ and its derived sentences.

Example: a proof of P from $\{(P \vee H) \wedge \neg H\}$

1.	$(P \vee H) \wedge \neg H$	*by assumption*
2.	$P \vee H$	*by $\wedge$-elimination applied to (1)*
3.	$\neg H$	*by $\wedge$-elimination applied to (1)*
4.	P	*by unit resolution applied to (2),(3)*

We can represent a proof more visually as a *proof tree*:

Example:

$$\frac{\dfrac{(P \vee H) \wedge \neg H}{P \vee H} \qquad \dfrac{(P \vee H) \wedge \neg H}{\neg H}}{P}$$

Rule-Based Inference in Propositional Logic

More formally, there is a proof of φ from Γ in $\mathcal{R}$ if

1. $\varphi \in \Gamma$ or,
2. there is a rule in $\mathcal{R}$ that applies to Γ and produces φ or,
3. there is a proof of each $\varphi_1, \ldots, \varphi_m$ from Γ in $\mathcal{R}$ and a rule that applies to $\{\varphi_1, \ldots, \varphi_m\}$ and produces φ.

Then, the inference system $\mathcal{R}$ is specified as follows:

$$\Gamma \vdash_{\mathcal{R}} \varphi \;\; \textit{iff there is a proof of } \varphi \textit{ from } \Gamma \textit{ in } \mathcal{R}$$

$\mathcal{R}$ is sound because all of its rules are sound.

Example: the Resolution rule $\dfrac{\alpha \vee \beta, \quad \neg\beta \vee \gamma}{\alpha \vee \gamma}$

	α	β	γ	$\neg\beta$	$\alpha \vee \beta$	$\neg\beta \vee \gamma$	$\alpha \vee \gamma$
1.	*False*	*False*	*False*	*True*	*False*	*True*	*False*
2.	*False*	*False*	*True*	*True*	*False*	*True*	*True*
3.	*False*	*True*	*False*	*False*	*True*	*False*	*False*
4.	*False*	*True*	*True*	*False*	<u>*True*</u>	<u>*True*</u>	*True*
5.	*True*	*False*	*False*	*True*	<u>*True*</u>	<u>*True*</u>	*True*
6.	*True*	*False*	*True*	*True*	<u>*True*</u>	<u>*True*</u>	*True*
7.	*True*	*True*	*False*	*False*	*True*	*False*	*True*
8.	*True*	*True*	*True*	*False*	<u>*True*</u>	<u>*True*</u>	*True*

All the interpretations that make both $\alpha \vee \beta$ *and* $\neg\beta \vee \gamma$ *true (ie, 4,5,6,8) make* $\alpha \vee \gamma$ *also true.*

- Exercise: prove that the other inference rules are sound as well.
- Is $\mathcal{R}$ also complete?

The rules of $\mathcal{R}$

$$\frac{\alpha \quad \beta}{\alpha \wedge \beta} \qquad \frac{\alpha}{\alpha \vee \beta} \qquad \frac{\alpha}{\beta \vee \alpha}$$

$$\frac{\alpha \wedge \beta}{\alpha} \qquad \frac{\alpha \wedge \beta}{\beta}$$

$$\frac{\alpha \Rightarrow \beta \quad \alpha}{\beta} \qquad \frac{\alpha \vee \beta \quad \neg\beta}{\alpha} \qquad \frac{\alpha \vee \beta \quad \neg\beta \vee \gamma}{\alpha \vee \gamma}$$

$$\frac{\neg\neg\alpha}{\alpha} \qquad \frac{\alpha \wedge \neg\alpha}{\mathbf{False}} \qquad \frac{\mathbf{False}}{\beta}$$