An Introduction to Logic

Madhavan Mukund
SPIC Mathematical Institute
92 G.N. Chetty Road
Madras 600 017, India
E-mail: madhavan@smi.ernet.in

Abstract

These are lecture notes for an introductory course on logic aimed at graduate students in Computer Science. The notes cover techniques and results from propositional logic, modal logic, propositional dynamic logic and first-order logic. The notes are based on a course taught to first year PhD students at SPIC Mathematical Institute, Madras, during August–December, 1997.

At the moment, these notes only cover propositional logic and modal logic. Notes on the remaining topics will be ready shortly.

Contents

1 Propositional Logic

1.1 Syntax

We begin with a countably infinite set of atomic propositions $\mathcal{P} = \{p_0, p_1, \ldots\}$ and two logical connectives $\neg$ (read as *not*) and $\vee$ (read as *or*).

The set Φ of formulas of propositional logic is the smallest set satisfying the following conditions:

- Every atomic proposition p is a member of Φ.
- If α is a member of Φ, so is $(\neg\alpha)$.
- If α and β are members of Φ, so is $(\alpha \vee \beta)$.

We shall normally omit parentheses unless we need to explicitly clarify the structure of a formula. We follow the convention that $\neg$ binds more tightly than $\vee$. For instance, $\neg\alpha \vee \beta$ stands for $((\neg\alpha) \vee \beta)$.

Exercise 1.1 Show that Φ is a countably infinite set. $\square$

The fact that Φ is the *smallest* set satisfying this inductive definition provides us with the principle of *structural induction.*

Structural induction principle Let S be a set such that:

- Every atomic proposition p is a member of S.
- If α is a member of S, so is $(\neg\alpha)$.
- If α and β are members of S, so is $(\alpha \vee \beta)$.

Then, $\Phi \subseteq S$.

1.2 Semantics

To assign meaning to formulas, we begin by assigning meaning to the atomic propositions. Let $\top$ denote the truth value *true* and $\perp$ the truth value *false.*

- A *valuation* v is a function $v : \mathcal{P} \to \{\top, \perp\}$.

We can also think of a valuation as a subset of $\mathcal{P}$—if $v : \mathcal{P} \to \{\top, \perp\}$, then $v \subseteq \mathcal{P} = \{p \mid v(p) = \top\}$. Thus, the set of all valuations is $2^{\mathcal{P}}$, the set of all subsets of $\mathcal{P}$.

We extend each valuation $v : \mathcal{P} \to \{\top, \perp\}$ to a map $\hat{v} : \Phi \to \{\top, \perp\}$ as follows:

- For $p \in \mathcal{P}$, $\hat{v}(p) = v(p)$.

- For α of the form $\neg\beta$, $\hat{v}(\alpha) = \begin{cases} \top & \text{if } \hat{v}(\beta) = \bot \\ \bot & \text{otherwise} \end{cases}$

- For α of the form $\beta \vee \gamma$, $\hat{v}(\alpha) = \begin{cases} \bot & \text{if } \hat{v}(\beta) = \hat{v}(\gamma) = \bot \\ \top & \text{otherwise} \end{cases}$

The principle of structural induction can be used to formally argue that $\hat{v}$ is well-defined (that is, $\hat{v}$ is indeed a function and is defined for all formulas).

Just as v can be defined as a subset of $\mathcal{P}$, $\hat{v}$ can be defined as a subset of Φ—namely, $\hat{v} = \{\alpha \mid \hat{v}(\alpha) = \top\}$.

Exercise 1.2 We saw that every subset of $\mathcal{P}$ defines a valuation v. Does every subset of Φ define an extended valuation $\hat{V}$? □

Since every valuation v gives rise to a unique extension $\hat{v}$, we shall always denote $\hat{v}$ as just v.

Derived connectives It will be convenient to introduce some additional connectives when discussing propositional logic.

$$\begin{aligned} \alpha \wedge \beta &\stackrel{\text{def}}{=} \neg(\neg\alpha \vee \neg\beta) \\ \alpha \supset \beta &\stackrel{\text{def}}{=} \neg\alpha \vee \beta \\ \alpha \equiv \beta &\stackrel{\text{def}}{=} (\alpha \supset \beta) \wedge (\beta \supset \alpha) \end{aligned}$$

The connective $\wedge$ is read as *and*, $\supset$ as *implies* and $\equiv$ as *if and only if*.

Exercise 1.3 Express $v(\alpha \wedge \beta)$, $v(\alpha \supset \beta)$ and $v(\alpha \equiv \beta)$ in terms of $v(\alpha)$ and $v(\beta)$. □

Exercise 1.4 According to the Pigeonhole Principle, if we try to place $n+1$ pigeons in n pigeonholes, then at least one pigeonhole must have two or more pigeons. For $i \in \{1, 2, \ldots, n+1\}$ and $j \in \{1, 2, \ldots, n\}$, let the atomic proposition p_{ij} denote that the i^{th} pigeon is placed in the the j^{th} pigeonhole. Write down a formula expressing the Pigeonhole Principle. What is the length of your formula as a function of n? □

Satisfiability and validity A formula α is said to be *satisfiable* if there is a valuation v such that $v(\alpha) = \top$. We write $v \models \alpha$ to indicate that $v(\alpha) = \top$.

The formula α is said to be *valid* if $v \models \alpha$ for every valuation v. We write $\models \alpha$ to indicate that α is valid. We also refer to valid formulas of propositional logic as *tautologies*.

Example 1.5 Let p be an atomic proposition. The formula p is satisfiable. The formula $p \vee \neg p$ is valid. The formula $p \wedge \neg p$ is not satisfiable.

The following observation connects the notions of satisfiability and validity.

Proposition 1.6 *Let α be a formula. α is valid iff $\neg\alpha$ is not satisfiable.*

In applications of logic to computer science, a central concern is to develop algorithms to check for satisfiability and validity of formulas. The preceding remark shows that the two notions are dual: an algorithm which tests validity of formulas can be converted into one for testing satisfiability and vice versa.

In principle, testing the validity of a formula α involves checking its truth value across an uncountable number of valuations. However, it is sufficient to look at the effect of valuations on the atomic propositions mentioned in α.

Let us define $\text{Voc}(\alpha)$, the *vocabulary* of α, as follows:

- For $p \in \mathcal{P}$, $\text{Voc}(p) = \{p\}$.
- If $\alpha = \neg\beta$, then $\text{Voc}(\alpha) = \text{Voc}(\beta)$.
- If $\alpha = \beta \vee \gamma$, then $\text{Voc}(\alpha) = \text{Voc}(\beta) \cup \text{Voc}(\gamma)$.

Proposition 1.7 *Let α be a formula and v_1, v_2 be valuations. If v_1 and v_2 agree on $\text{Voc}(\alpha)$ then $v_1(\alpha) = v_2(\alpha)$.*

This justifies the familiar algorithm for testing validity: build a truth-table for the propositions mentioned in α and check if all rows yield the value $\top$.

1.3 Axiomatisations

Though we have a straightforward algorithm for testing validity of formulas in propositional logic, such algorithms do not exist for more complicated logical systems. In particular, there is no such algorithm for first-order logic.

However, it is still possible to effectively enumerate all the valid formulas of first-order logic. One way of presenting such an enumeration is through an axiomatisation of the logic. To prepare the ground for studying axiomatisations of more complex logics, we begin with an axiomatisation for propositional logic.

Axiom System *AX* The axiom system AX consists of three axioms and one inference rule.

$$\begin{array}{rl}
\text{(A1)} & \alpha \supset (\beta \supset \alpha) \\
\text{(A2)} & (\alpha \supset (\beta \supset \gamma)) \supset ((\alpha \supset \beta) \supset (\alpha \supset \gamma)) \\
\text{(A3)} & (\neg\beta \supset \neg\alpha) \supset ((\neg\beta \supset \alpha) \supset \beta) \\
\text{(Modus Ponens, or MP)} & \dfrac{\alpha, \ \alpha \supset \beta}{\beta}
\end{array}$$

The rule MP is read as follows—from α and $\alpha \supset \beta$, infer β. It is important to note that these are *axiom schemes*—that is, they are not actual formulas but templates which can be instantiated into real formulas by consistently substituting concrete formulas for α, β and γ. For instance, if $p, q \in \mathcal{P}$, $p \supset (q \supset p)$ is an instance of axiom (A1). An alternate way to present such an axiomatisation is to list the axioms as concrete formulas and have an additional inference rule to permit uniform substitution of new formulas into an existing formula.

Derivations A *derivation* of α using the axiom system AX is a finite sequence of formulas $\beta_1, \beta_2, \ldots \beta_n$ such that:

- $\beta_n = \alpha$
- For each $i \in \{1, 2, \ldots, n\}$, β_i is either an instance of one of the axioms (A1)–(A3), or is obtained by applying the rule (MP) to formulas β_j, β_k, where $j, k < i$—that is, β_k is of the form $\beta_j \supset \beta_i$.

We write $\vdash_{AX} \alpha$ to denote that α is derivable using the axiom system AX and say that α is a *thesis* of the system. We will normally omit the subscript AX.

Here is an example of a derivation using our axiom system.

1.	$(p \supset ((p \supset p) \supset p)) \supset ((p \supset (p \supset p)) \supset (p \supset p))$	Instance of (A2)
2.	$p \supset ((p \supset p) \supset p)$	Instance of (A1)
3.	$(p \supset (p \supset p)) \supset (p \supset p))$	From 1 and 2 by MP
4.	$p \supset (p \supset p)$	Instance of (A1)
5.	$p \supset p$	From 3 and 4 by MP

Exercise 1.8 Show that $(\neg\beta \supset \neg\alpha) \equiv (\alpha \supset \beta)$ is a thesis of AX. □

The axiom system we have presented is called a Hilbert-style axiomatisation. There are several other ways of presenting axiomatisations. One common alternative to Hilbert-style systems is the *sequent calculus* notation due to Gentzen. Typically, Hilbert-style axiomatisations have a large number of axioms and very few inference rules, while sequent calculi have very few axioms and a large number of inference rules. Sequent calculi are often easier to work with when searching for derivations, but are also more complicated from a technical point of view. We shall look at sequent calculi later, when we come to first-order logic.

Another fact worth remembering is that the axiom system AX defined here is just one of many possible Hilbert-style axiom systems for propositional logic.

The main technical result we would like to establish is that the set of formulas derivable using AX is precisely the set of valid formulas of propositional logic.

Theorem 1.9 *For all formulas α, $\vdash \alpha$ iff $\models \alpha$.*

We break up the proof of this theorem into two parts. The first half is to show that every thesis of AX is valid. This establishes the *soundness* of the axiom system,

Lemma 1.10 (Soundness) *For all formulas α, if $\vdash \alpha$ then $\models \alpha$.*

Proof: If $\vdash \alpha$, then we can exhibit a derivation $\beta_1, \beta_2, \ldots, \beta_n$ of α. Formally, the proof of the lemma is by induction on the length of this derivation. Since every formula in the sequence $\beta_1, \beta_2, \ldots, \beta_n$ is either an instance of one of the axioms or is obtained by applying the rule (MP), it suffices to show that all the axioms define valid formulas and that (MP) preserves validity—in other words, if α is valid and $\alpha \supset \beta$ is valid, then β is valid. This is straightforward and we omit the details. □

The other half of Theorem 1.9 is more difficult to establish. We have to argue that every valid formula is derivable. Formally, this would show that our axiomatisation is *complete*.

We follow the approach of the logician Leon Henkin and attack the problem indirectly. Consider the contrapositive of the statement we want to prove—that is, if a formula α is *not* a thesis, then it is *not* valid.

Consistency We write $\not\vdash \alpha$ to denote that α is not a thesis. We say that α is *consistent* (with respect to AX) if $\not\vdash \neg\alpha$.

Exercise 1.11

(i) Show that $\alpha \vee \beta$ is consistent iff either α is consistent or β is consistent.

(ii) Show that if $\alpha \wedge \beta$ is consistent then both α and β are consistent. Is the converse true?

(iii) Suppose that $\vdash \alpha \supset \beta$. Which of the following is true?

  (a) If α is consistent then β is consistent.

  (b) If β is consistent then α is consistent.

□

By Proposition 1.6 we know that α is not valid iff $\neg\alpha$ is satisfiable. Suppose we can show the following.

Lemma 1.12 (Henkin) *For all formulas β, if β is consistent then β is satisfiable.*

We can then argue that our axiomatisation is complete. Consider a formula β which is not derivable. It can be shown that $\neg\neg\beta \supset \beta$ is a thesis. If β is not derivable, neither is $\neg\neg\beta$—otherwise, we can use the rule MP to derive β from $\neg\neg\beta \supset \beta$. Since $\not\vdash \neg(\neg\beta)$, $\neg\beta$ is consistent. By Lemma 1.12, $\neg\beta$ is satisfiable. Hence, by Proposition 1.6, β is not valid.

1.4 Maximal Consistent Sets and Completeness

To prove Lemma 1.12, we extend the notion of consistency from a single formula to sets of formulas. A finite set of formulas $X = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ is consistent if the formula $\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n$ is consistent—that is, $\not\vdash \neg(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n)$. An arbitrary set of formulas $X \subseteq \Phi$ is consistent if every finite subset of X is consistent. (Henceforth, $Y \subseteq_{\mathsf{fin}} X$ denotes that Y is a finite subset of X.)

A *maximal consistent set (MCS)* is a consistent set which cannot be extended by adding any formulas. In other words, $X \subseteq \Phi$ is an MCS iff X is consistent and for each formula $\alpha \notin X$, $X \cup \{\alpha\}$ is inconsistent.

Lemma 1.13 (Lindenbaum) *Every consistent set can be extended to an MCS.*

Proof: Let X be an arbitrary consistent set. Let $\alpha_0, \alpha_1, \alpha_2, \ldots$ be an enumeration of Φ.
We define an infinite sequence of sets $X_0, X_1, X_2, \ldots$ as follows.

- $X_0 = X$
- For $i \geq 0$, $X_{i+1} = \begin{cases} X_i \cup \{\alpha_i\} & \text{if } X_i \cup \{\alpha_i\} \text{ is consistent} \\ X_i & \text{otherwise} \end{cases}$

Each set in this sequence is consistent, by construction, and $X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots$. Let $Y = \bigcup_{i \geq 0} X_i$. We claim that Y is an MCS extending X. To establish this, we have to show that Y is consistent and that it maximal.

If Y is not consistent, then there is a subset $Z \subseteq_{\mathsf{fin}} Y$ which is inconsistent. Let $Z = \{\beta_1, \beta_2, \ldots, \beta_n\}$. We can write Z as $\{\alpha_{i_1}, \alpha_{i_2}, \ldots, \alpha_{i_n}\}$ where the indices correspond to our enumeration of Φ. Let $j = \max(i_1, i_2, \ldots, i_n)$. Then it is clear that $Z \subseteq_{\mathsf{fin}} X_{j+1}$ in the sequence $X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots \subseteq Y$. This implies that X_{j+1} is inconsistent, which is a contradiction.

Having established that Y is consistent, we show that it is maximal. Suppose that $Y \cup \{\beta\}$ is consistent for some formula $\beta \notin Y$. Let $\beta = \alpha_j$ in our enumeration of Φ. Since $\alpha_j \notin Y$, α_j was not added at step $j+1$ in our construction. This means that $X_j \cup \{\alpha_j\}$ is inconsistent. In other words, there exists $Z \subseteq_{\mathsf{fin}} X_j$ such that $Z \cup \{\alpha_j\}$ is inconsistent. Since $X_j \subseteq Y$, we must have $Z \subseteq_{\mathsf{fin}} Y$ as well, which contradicts the assumption that $Y \cup \{\alpha_j\}$ is consistent.
$\square$

Maximal consistent sets have a rich structure which we shall exploit to prove completeness.

Lemma 1.14 *Let X be a maximal consistent set. Then:*

(i) For all formulas α, $\alpha \in X$ iff $\neg\alpha \notin X$.

(ii) For all formulas α, β, $\alpha \vee \beta \in X$ iff $\alpha \in X$ or $\beta \in X$.

We postpone the proof of these properties and first show how they lead to completeness.

Maximal consistent sets and valuations Let X be an MCS. Define the valuation v_X to be the set $\{p \in \mathcal{P} \mid p \in X\}$—in other words, $v_X(p) = \top$ iff $p \in X$.

Proposition 1.15 *Let X be an MCS. For all formulas α, $v_X \models \alpha$ iff $\alpha \in X$.*

Proof: The proof is by induction on the structure of α.

Basis: $\alpha = p$, where $p \in \mathcal{P}$. Then, $v_X \models p$ iff (by the definition of v_X) $p \in X$.

Induction step: There are two cases to consider—when α is of the form $\neg\beta$ and when α is of the form $\beta \vee \gamma$.

($\alpha = \neg\beta$) $v_X \models \neg\beta$ iff (by the definition of valuations) $v_X \not\models \beta$ iff (by the induction hypothesis) $\beta \notin X$ iff (by the properties satisfied by MCSs) $\neg\beta \in X$.

($\alpha = \beta \vee \gamma$) $v_X \models \beta \vee \gamma$ iff (by the definition of valuations) $v_X \models \beta$ or $v_X \models \gamma$ iff (by the induction hypothesis) $\beta \in X$ or $\gamma \in X$ iff (by the properties satisfied by MCSs) $\beta \vee \gamma \in X$. □

Thus, every MCS X defines a canonical valuation v_X which satisfies precisely those formulas that belong to X. (Conversely, every valuation also defines an MCS in a canonical way: given a valuation v, $X_v = \{\alpha \mid v \models \alpha\}$. It is not difficult to establish that the valuation v_{X_v} generated by X_v is exactly the same as v.)

Proposition 1.15 immediately yields a proof of Henkin's lemma.

Proof: (of Lemma 1.12)
Let α be a consistent formula. By Lindenbaum's Lemma, $\{\alpha\}$ can be extended to an MCS X. By Proposition 1.15, $v_X \models \alpha$ since $\alpha \in X$. Thus, α is satisfiable. □

To complete our argument, we have to prove Lemma 1.14.

Proof Sketch: (of Lemma 1.14)
Let X be an MCS.

(i) For every formula α, we have to show that $\alpha \in X$ iff $\neg\alpha \notin X$.

We first show that $\{\alpha, \neg\alpha\} \not\subseteq X$. For this, we need the fact that $\alpha \supset \neg\neg\alpha$ and $\neg\neg\alpha \supset \alpha$ are both derivable using AX. We omit these derivations.

We know that $\alpha \supset \alpha$, or, equivalently, $\neg\alpha \vee \alpha$ is a thesis. From this, we can derive $\neg\neg(\neg\alpha \vee \alpha)$. But $\neg(\neg\alpha \vee \alpha)$ is just $\alpha \wedge \neg\alpha$, so we have $\neg(\alpha \wedge \neg\alpha)$ as a thesis. This means that $\{\alpha, \neg\alpha\}$ is inconsistent, which is a contradiction.

Next we show that at least one of α and $\neg\alpha$ is in X. Suppose neither formula belongs to X. Since X is an MCS, there must be sets $B \subseteq_{\text{fin}} X$ and $C \subseteq_{\text{fin}} X$ such that $B \cup \{\alpha\}$ is inconsistent and $C \cup \{\neg\alpha\}$ are inconsistent. Let $B = \{\beta_1, \beta_2, \ldots, \beta_n\}$ and $C = \{\gamma_1, \gamma_2, \ldots, \gamma_m\}$. Let $\hat{\beta}$ abbreviate the formula $\beta_1 \wedge \beta_2 \wedge \ldots \wedge \beta_n$ and $\hat{\gamma}$ abbreviate the formula $\gamma_1 \wedge \gamma_2 \wedge \ldots \wedge \gamma_m$. Then, we have $\vdash \neg(\alpha \wedge \hat{\beta})$ and $\vdash \neg(\neg\alpha \wedge \hat{\gamma})$. Rewriting

$\wedge$ in terms of $\vee$, this is equivalent to $\vdash \neg\alpha \vee \neg\hat{\beta}$ and $\vdash \neg\neg\alpha \vee \neg\hat{\gamma}$. From this, we can conclude that $\vdash \alpha \supset \neg\hat{\beta}$ and $\vdash \neg\alpha \supset \neg\hat{\gamma}$.

We now use that fact that $(\alpha \supset \beta) \supset ((\delta \supset \gamma) \supset ((\alpha \vee \delta) \supset (\beta \vee \gamma)))$ is a thesis. (Once again, we omit the derivation). Instantiating this with $\alpha = \alpha$, $\delta = \neg\alpha$, $\beta = \neg\hat{\beta}$ and $\gamma = \neg\hat{\gamma}$ we can derive $(\alpha \vee \neg\alpha) \supset (\neg\hat{\beta} \vee \neg\hat{\gamma})$. Since $\vdash \alpha \vee \neg\alpha$, we get $\vdash \neg\hat{\beta} \vee \neg\hat{\gamma}$. By rewriting $\vee$ in terms of $\wedge$, we can derive $\neg(\hat{\beta} \wedge \hat{\gamma})$. But this implies that $(B \cup C) \subseteq_{\mathrm{fin}} X$ is inconsistent, which is a contradiction.

(ii) The proof of the second part follows in a similar manner, assuming the derivability of appropriate formulas. We omit the details.

$\square$

1.5 Compactness and Strong Completeness

Often, we are not interested in absolute validity, but in restricted validity. Rather than asking whether a formula α is *always* true, we ask whether α is true in all valuations which satisfy certain properties. One way of restricting the class of valuations under consideration is to specify a set of formulas X and only look at those valuations where X is true. If α is true wherever the formulas from X are true, then α is a logical consequence of X.

Logical consequence Let X be a set of formulas and v a valuation. We write $v \models X$ to denote that $v \models \beta$ for every formula $\beta \in X$. A formula α is a *logical consequence of* X, written $X \models \alpha$, if for every valuation v such that $v \models X$ it is also the case that $v \models \alpha$.

The notion of logical consequence is central to the way we formalise mathematics. For instance, when we study algebraic structures such as groups, we first formulate axioms which characterise groups. Any theorem we prove about groups can be rephrased as a statement which is a logical consequence of these axioms: in other words, the theorem is true whenever the group axioms are also true.

As with validity, we now look at a syntactic approach to logical consequence.

Derivability Let X be a set of formulas. We say that a formula α is *derivable from* X, written $X \vdash \alpha$ if there exists a sequence $\alpha_1, \alpha_2, \ldots, \alpha_n$ of formulas such that $\alpha_n = \alpha$ and for $i \in \{1, 2, \ldots, n\}$, α_i is either a member of X, or an instance of one of the axioms (A1)–(A3) of AX, or is derived from α_j, α_k, $j, k < i$, using the inference rule MP. (Notice that unlike axioms, we cannot use the formulas in X as templates to generate new formulas for use in a derivation. The formulas in X are concrete formulas and must be used "as is".)

The theorem we would like to prove is the following.

Theorem 1.16 (Strong Completeness) *Let $X \subseteq \Phi$ and $\alpha \in \Phi$. Then, $X \models \alpha$ iff $X \vdash \alpha$.*

It is possible to prove this directly using a technique similar to the one used to prove the soundness and completeness of AX (see Exercise 1.22). However, we will prove it indirectly using two auxiliary results which are of independent interest—the Deduction Theorem and the Compactness Theorem.

We begin with the Deduction Theorem, which is a statement about derivability.

Theorem 1.17 *(Deduction) Let $X \subseteq \Phi$ and $\alpha, \beta \in \Phi$. Then, $X \cup \{\alpha\} \vdash \beta$ iff $X \vdash \alpha \supset \beta$.*

Proof: ($\Leftarrow$) Suppose that $X \vdash \alpha \supset \beta$. Then, by the definition of derivability, $X \cup \{\alpha\} \vdash \alpha \supset \beta$ as well. Since $\alpha \in X \cup \{\alpha\}$, $X \cup \{\alpha\} \vdash \alpha$. Applying MP, we get $X \cup \{\alpha\} \vdash \beta$.

($\Rightarrow$) Suppose that $X \cup \{\alpha\} \vdash \beta$. Then, there is a derivation $\beta_1, \beta_2, \ldots, \beta_n$ of β. The proof is by induction on n.

If $n = 1$, then β is either an instance of an axiom or a member of $X \cup \{\alpha\}$. If β is an instance of an axiom, then $X \vdash \beta$ as well. Further, from axiom (A1), $X \vdash \beta \supset (\alpha \supset \beta)$. Applying MP, we get $X \vdash \alpha \supset \beta$.

If $\beta \in X$, there are two cases to consider. If $\beta \in X \setminus \{\alpha\}$, then $X \vdash \beta$. Once again we have $X \vdash \beta \supset (\alpha \supset \beta)$ and hence $X \vdash \alpha \supset \beta$. On the other hand, if $\beta = \alpha$, we have $X \vdash \alpha \supset \alpha$ from the fact that $\alpha \supset \alpha$ is derivable in AX.

If $n > 1$, we look the justification for adding $\beta_n = \beta$ to the derivation. If β_n is an instance of an axiom or a member of $X \cup \{\alpha\}$, we can use the same argument as in the base case to show $X \vdash \alpha \supset \beta$.

On the other hand, if β_n was derived using MP, there exist β_i and β_j, with $i, j < n$ such that β_j is of the form $\beta_i \supset \beta_n$. By axiom (A2), $X \vdash (\alpha \supset (\beta_i \supset \beta_n)) \supset ((\alpha \supset \beta_i) \supset (\alpha \supset \beta_n))$. By the induction hypothesis, we know that $X \vdash \alpha \supset (\beta_i \supset \beta_n)$ and $X \vdash \alpha \supset \beta_i$. Applying MP twice, we get $X \vdash \alpha \supset \beta_n$.

$\square$

The Deduction Theorem reflects a method of proof which is common in mathematics—proving that property x implies property y is equivalent to assuming x and inferring y.

The second step in proving Strong Completeness is the Compactness Theorem, which is a statement about logical consequence. To prove this we need the following lemma about trees, due to König.

Lemma 1.18 (König) *Let T be a finitely branching tree—that is, every node has a finite number of children (though this number may be unbounded). If T has infinitely many nodes, then T has an infinite path.*

Proof: Let T be a finitely branching tree with infinitely many nodes. Call a node x in T *bad* if the subtree rooted at x has infinitely many nodes. Clearly, if a node x is bad, at least one of its children must be bad: x has only finitely many children and if all of them were good, the subtree rooted at x would be finite.

We now construct an infinite path $x_0 x_1 x_2 \ldots$ in T. Since T has an infinite number of nodes, the root of T is a bad node. Let x_0 be the root of T. It has at least one bad successor.

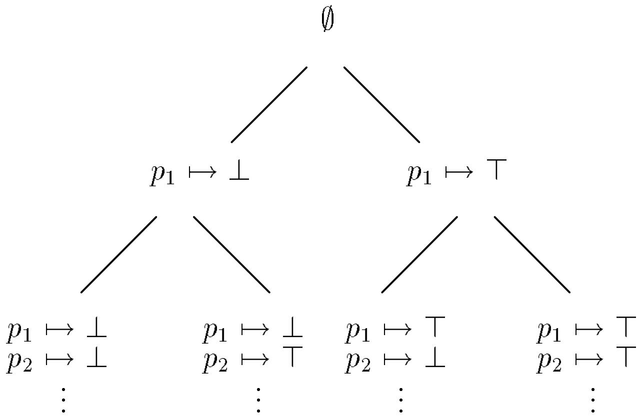

Figure 1: The tree T in the proof of Lemma 1.20

Pick one of the bad successors of x_0 and designate it x_1. Pick one of the bad successors of x_1 and designate it x_2, and so on. $\square$

Theorem 1.19 (Compactness) *Let $X \subseteq \Phi$ and $\alpha \in \Phi$. Then $X \models \alpha$ iff there exists $Y \subseteq_{\text{fin}} X$, $Y \models \alpha$.*

We shall first prove the following related result. Let X be a set of formulas. We say that X is satisfiable if there exists a valuation v such that $v \models X$.

Lemma 1.20 (Finite satisfiability) *Let $X \subseteq \Phi$. Then, X is satisfiable iff every $Y \subseteq_{\text{fin}} X$ is satisfiable.*

Proof: ($\Rightarrow$) Suppose X is satisfiable. Then, there is a valuation v such that $v \models X$. Clearly, $v \models Y$ for each $Y \subseteq_{\text{fin}} X$ as well.

($\Leftarrow$) Suppose X is *not* satisfiable. We have to show that there exists $Y \subseteq_{\text{fin}} X$ which is not satisfiable.

Assume that our set of atomic propositions $\mathcal{P}$ is enumerated $\{p_1, p_2, \ldots\}$. Let $\mathcal{P}_0 = \emptyset$ and for $i \in \{1, 2, \ldots\}$, let $\mathcal{P}_i = \{p_1, p_2, \ldots, p_i\}$. For $i \in \{1, 2, \ldots\}$, let Φ_i be the set of formulas generated using only atomic propositions from $\mathcal{P}_i$ and let $X_i = X \cap \Phi_i$.

We construct a tree T whose nodes are valuations over the sets $\mathcal{P}_i$, $i \in \{0, 1, 2, \ldots\}$. More formally, the set of nodes is given by $\{v \mid \exists i \in \{0, 1, 2, \ldots\}.\ v : \mathcal{P}_i \to \{\top, \perp\}\}$. The root of T is the unique function $\emptyset \to \{\top, \perp\}$.

The relation between nodes is given as follows. Let $v : \mathcal{P}_i \to \{\top, \perp\}$. Then v has two children $v', v'' : \mathcal{P}_{i+1} \to \{\top, \perp\}$, where v' extends v to $\mathcal{P}_{i+1}$ by setting p_{i+1} to $\top$ and v'' extends v to $\mathcal{P}_{i+1}$ by setting p_{i+1} to $\perp$. More formally, for each $p \in \mathcal{P}_i$, $v'(p) = v''(p) = v(p)$ and $v'(p_{i+1}) = \top$ and $v''(p_{i+1}) = \perp$. (See Figure 1).

Observe that T is a complete infinite binary tree. The nodes at level i of the tree consist of all possible valuations over $\mathcal{P}_i$—there are precisely 2^i such valuations for each i. Notice that if v_j at level j is an ancestor of v_i at level i then v_i agrees with v_j on the atomic propositions in $\mathcal{P}_j$.

The infinite paths in T are in 1-1 correspondence with valuations over $\mathcal{P}$. Let $\pi = v_0 v_1 v_2 \ldots$ be an infinite path in the tree. The valuation $v_\pi : \mathcal{P} \rightarrow \{\top, \bot\}$ is given by $p_i \mapsto v_i(p_i)$ for $i \in \{1, 2, \ldots\}$. Conversely, given a valuation $v : \mathcal{P} \rightarrow \{\top, \bot\}$, we can define a unique path $\pi_v = v_0 v_1 v_2 \ldots$ by setting v_0 to be the root of T and $v_i : \mathcal{P}_i \rightarrow \{\top, \bot\}$ to be the restriction of v to $\mathcal{P}_i$—that is, for all $p \in \mathcal{P}_i$, $v_i(p) = v(p)$. It is easy to verify that these two maps are inverses of each other.

Let us call a node v in T bad if $v(\beta) = \bot$ for some $\beta \in X$. Clearly, if v is bad, then so is every valuation in the subtree rooted at v. We prune T by deleting all bad nodes which also have bad ancestors. (Equivalently, along any path in T, we retain only those nodes upto and including the first bad node along the path.) It is not difficult to verify that the set of nodes which remains forms a subtree T' of T all of whose leaf nodes are bad and all of whose non-leaf nodes are not bad.

We claim that T' has only a finite number of nodes. Assuming that this is true, let $\{v_1, v_2, \ldots, v_m\}$ be the leaf nodes of T'. Since each v_i is bad, there is a corresponding formula $\beta_i \in X$ such that $v_i(\beta_i) = \bot$. We claim that $\{\beta_1, \beta_2, \ldots, \beta_m\} \subseteq_{\text{fin}} X$ is not satisfiable. Consider any valuation v. The corresponding path v_π must pass through one of the nodes in $\{v_1, v_2, \ldots, v_m\}$, say v_j. But then, $v_\pi(\beta_j) = v_j(\beta_j) = \bot$. Thus, $v \not\models \{\beta_1, \beta_2, \ldots, \beta_m\}$.

To see why T' must be finite, suppose instead that it has an infinite set of nodes. Then, by König's Lemma, it contains an infinite path $\pi = v_0 v_1 v_2 \ldots$ such that none of the nodes along this path is bad. The path π is also an infinite path in T. We know that π defines a valuation v_π. Consider any formula $\beta \in X$. Then $\beta \in X_j$ for some $j \in \{1, 2, \ldots\}$, so $v_\pi(\beta) = v_j(\beta) = \top$. Thus, $v_\pi \models X$, which contradicts our assumption that X is not satisfiable. $\square$

We can now complete our proof of compactness.

Proof: [of Theorem 1.19 (Compactness)] ($\Leftarrow$) If $Y \subseteq_{\text{fin}} X$ and $Y \models \alpha$ then it is clear that $X \models \alpha$. For, if $v \models X$, then $v \models Y$ as well and, by the assumption that $Y \models \alpha$, $v \models \alpha$ as required.

($\Rightarrow$) For all $Z \subseteq \Phi$ and all $\beta \in \Phi$, it is clear that $Z \models \beta$ iff $Z \cup \{\neg\beta\}$ is not satisfiable.

Suppose $X \models \alpha$. Then, $X \cup \{\neg\alpha\}$ is not satisfiable. By Lemma 1.20, there is a subset $Y \subseteq_{\text{fin}} X \cup \{\neg\alpha\}$ such that Y is not satisfiable. Thus, $(Y \setminus \{\neg\alpha\}) \cup \{\neg\alpha\}$ is not satisfiable either, where $(Y \setminus \{\neg\alpha\}) \subseteq_{\text{fin}} X$. This implies that $Y \setminus \{\neg\alpha\} \models \alpha$. $\square$

With the Deduction Theorem and the Compactness Theorem behind us, we can prove Strong Completeness.

Proof: [Theorem 1.16 (Strong Completeness)]
To show that $X \vdash \alpha$ implies $X \models \alpha$ is routine. Conversely, suppose that $X \models \alpha$. By compactness, there is a finite subset $Y \subseteq_{\text{fin}} X$ such that $Y \models \alpha$. Let $Y = \{\beta_1, \beta_2, \ldots, \beta_m\}$. It is then easy to see that $\beta_1(\supset (\beta_2(\supset \cdots (\beta_m \supset \alpha) \cdots)$ is valid. Hence, by the completeness

theorem for propositional logic, $\vdash \beta_1 (\supset (\beta_2 (\supset \cdots (\beta_m \supset \alpha) \cdots)$. Applying the Deduction Theorem m times we get $\{\beta_1, \beta_2, \ldots, \beta_m\} \vdash \alpha$. Since $\{\beta_1, \beta_2, \ldots, \beta_m\} \subseteq X$, it follows that $X \vdash \alpha$. □

Observe that we could alternatively derive compactness from strong completeness. If $X \models \alpha$ then, by strong completeness, $X \vdash \alpha$. We let $Y \subseteq_{\text{fin}} X$ be the subset of formulas actually used in the derivation of α. Thus, $Y \vdash \alpha$ as well. By the other half of strong completeness, $Y \models \alpha$.

We conclude our discussion of propositional logic with two exercises. The first leads to an alternative proof of compactness which is more along the lines of the completeness proof for propositional logic. The second exercise leads to a direct proof of strong completeness.

Exercise 1.21 (Compactness)
Let X be a set of formulas. X is said to be a *finitely satisfiable set (FSS)* if every $Y \subseteq_{\text{fin}} X$ is satisfiable.

Equivalently, X is an FSS if there is no finite subset $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ of X such that $\neg(\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n)$ is valid.

(Note that if X is an FSS we are not promised a single valuation v which satisfies every finite subset of X. Each finite subset could be satisfied by a *different* valuation).

Show that:

(i) Every FSS can be extended to a maximal FSS.

(ii) If X is a maximal FSS then:

  (a) For every formula α, $\alpha \in X$ iff $\neg\alpha \notin X$.

  (b) For all formulas α, β, $(\alpha \vee \beta) \in X$ iff $(\alpha \in X$ or $\beta \in X)$.

(iii) Every maximal FSS X generates a valuation v_X such that for every formula α, $v_X \models \alpha$ iff $\alpha \in X$.

From these facts conclude that:

(iv) Any FSS X is simultaneously satisfiable (that is, for any FSS X, there exists v_X such that $v_X \models X$).

(v) For all X and all α, $X \models \alpha$ iff there exists $Y \subseteq_{\text{fin}} X$ such that $Y \models \alpha$.

□

Exercise 1.22 (Strong Completeness)
We define a new notion of consistency. A set X is said to be *consistent* if there is no formula α such that $X \vdash \alpha$ and $X \vdash \neg\alpha$.

Show that:

(i) X is consistent iff every finite subset of X is consistent.

(ii) Every consistent set X can be extended to a maximal consistent set (MCS).

(iii) Every MCS X generates a valuation v_X such that for all formulas α, $v_X \models \alpha$ iff $\alpha \in X$.

(iv) Every consistent set X is satisfiable: that is, there exists a valuation v_X such that $v_X \models X$.

(v) If $X \models \alpha$ then $X \cup \{\neg\alpha\}$ is not consistent.

(vi) Use the Deduction Theorem to show that that if $X \models \alpha$ then $X \vdash \neg\alpha \supset (\beta \wedge \neg\beta)$ for some formula β.

Conclude that if $X \models \alpha$ then $X \vdash \alpha$.

$\square$

2 Modal Logic

In propositional logic, a valuation is a static assignment of truth values to atomic propositions. In computer science applications, atomic propositions describe properties of the current state of a program. It is natural to expect that the truth of an atomic proposition varies as the state changes. Modal logic is a framework to describe such a situation.[1]

The basic idea in modal logic is to look at a collection of possible valuations simultaneously. Each valuation represents a possible state of the world. Separately, we specify how these "possible worlds" are connected to each other. We then enrich our logical language with a way of referring to truth across possible worlds.

2.1 Syntax

As in propositional logic, we begin with a countably infinite set of atomic propositions $\mathcal{P} = \{p_0, p_1, \ldots\}$ and two logical connectives $\neg$ (read as *not*) and $\vee$ (read as *or*). We add a unary *modality* $\Box$ (read as *box*).

The set Φ of formulas of modal logic is the smallest set satisfying the following:

- Every atomic proposition p is a member of Φ.

[1] Traditional modal logic arose out of philosophical enquiries into the nature of necessary and conditional truth. We shall concentrate on the technical aspects of the subject and avoid all discussion of the philosophical foundations of modal logic.

- If α is a member of Φ, so is $(\neg\alpha)$.
- If α and β are members of Φ, so is $(\alpha \vee \beta)$.
- If α is a member of Φ, so is $(\Box\alpha)$.

As before, we omit parentheses if there is no ambiguity. The derived propositional connectives $\wedge$, $\supset$ and $\equiv$ are defined as before. In addition, we have a derived modality $\Diamond$ (read *diamond*) which is *dual* to the modality $\Box$, defined as follows: $\Diamond\alpha \stackrel{\text{def}}{=} \neg\Box\neg\alpha$.

2.2 Semantics

Frames A *frame* is a structure $F = (W, R)$, where W is a set of *possible worlds* and $R \subseteq W \times W$ is the *accessibility relation*. If $w\ R\ w'$, we say that w' is an R-neighbour of w.

In more familiar terms, a frame is just a directed graph over the set of nodes W. We do not make any assumptions about the set W—not even the fact that it is countable.

Models A model is a pair $M = (F, V)$ where $F = (W, R)$ is a frame and $V : W \to 2^{\mathcal{P}}$ is a *valuation*.[2]

Recall that a propositional valuation $v : \mathcal{P} \to \{\top, \bot\}$ can also be viewed as a set $v \subseteq \mathcal{P}$ consisting of those atomic propositions p such that $v(p) = \top$. We have implicitly used this when defining valuations in modal logic. Formally, V is a function which assigns a propositional valuation to each world in W—in other words, for each $w \in W$, $V(w) : \mathcal{P} \to \{\top, \bot\}$. Thus, V is actually a function of the form $W \to (\mathcal{P} \to \{\top, \bot\})$, which we abbreviate as $V : W \to 2^{\mathcal{P}}$.

Satisfaction The notion of truth is localised to each world in a model. We write $M, w \models \alpha$ to denote that α is true at the world w in the model M. The satisfaction relation is defined inductively as follows.

$$
\begin{array}{ll}
M, w \models p & \text{iff } p \in V(w) \text{ for } p \in \mathcal{P} \\
M, w \models \neg\alpha & \text{iff } M, w \not\models \alpha \\
M, w \models \alpha \vee \beta & \text{iff } M, w \models \alpha \text{ or } M, w \models \beta \\
M, w \models \Box\alpha & \text{iff for each } w' \in W, \text{ if } w\ R\ w' \text{ then } M, w' \models \alpha
\end{array}
$$

Thus, $M, w \models \Box\alpha$ if every world accessible from w satisfies α. Notice that if w is isolated—that is, there is no world w' such that $w\ R\ w'$—then $M, w \models \Box\alpha$ for *every* formula α.

Exercise 2.1 Verify that $M, w \models \Diamond\alpha$ iff there exists w', $w\ R\ w'$ and $M, w' \models \alpha$. $\Box$

[2]The semantics we describe here was first formalised by Saul Kripke, so these models are often called Kripke models in the literature.

Satisfiability and validity As usual, we say that α is *satisfiable* if there exists a frame $F = (W, R)$ and a model $M = (F, V)$ such that $M, w \models \alpha$ for some $w \in W$. The formula α is *valid*, written $\models \alpha$, if for every frame $F = (W, R)$, for every model $M = (F, V)$ and for every $w \in W$, $M, w \models \alpha$.

Example 2.2 Here are some examples of valid formulas in modal logic.

(i) Every tautology of propositional logic is valid. Consider a tautology α and a world w in a model $M = ((W, R), V)$. Since the truth of α depends only on $V(w)$, and α is true under all propositional valuations, $M, w \models \alpha$.

(ii) The formula $\Box(\alpha \supset \beta) \supset (\Box\alpha \supset \Box\beta)$ is valid. Consider a model $M = ((W, R), V)$ and a world $w \in W$. Suppose that $M, w \models \Box(\alpha \supset \beta)$. We must argue that $M, w \models \Box\alpha \supset \Box\beta$. Let $M, w \models \Box\alpha$. Then we must show that $M, w \models \Box\beta$. In other words, we must show that every R-neighbour w' of w satisfies β. Since we assumed $M, w \models \Box(\alpha \supset \beta)$, we know that $M, w' \models \alpha \supset \beta$. Moreover, since $M, w \models \Box\alpha$, $M, w' \models \alpha$. By the semantics of the connective $\supset$, it follows that $M, w' \models \beta$, as required.

(iii) Suppose that α is valid. Then, $\Box\alpha$ must also be valid. Consider any model $M = ((W, R), V)$ and any $w \in W$. To check that $M, w \models \Box\alpha$ we have to verify that every R-neighbour of w satisfies α. Since α is valid, $M, w' \models \alpha$ for all $w' \in W$. So, every R-neighbour of w does satisfy α and $M, w \models \Box\alpha$.

Exercise 2.3 The argument given in part (i) of Exercise 2.2 applies only to non-modal instances of propositional tautologies—for instance, the explanation does not justify the validity of the formula $\Box\alpha \vee \neg\Box\alpha$. Show that *all* substitution instances of propositional tautologies are valid formulas in modal logic. $\Box$

As in propositional logic, one of our central concerns in modal logic is to be able to decide when formulas are satisfiable (or, dually, valid). Notice that unlike the truth-table based algorithm for propositional logic, there is no obvious decision procedure for satisfiability in modal logic. To check satisfiability of a formula α, though it suffices to look at valuations over the vocabulary of α, we also have to specify an underlying frame. There is no a priori bound on the size of this frame.

Later in this section we will describe a sound and complete axiomatisation for modal logic. This will give us an effective way of enumerating all valid formulas. After that, we will encounter a technique by which we can bound the size of the underlying frame required to satisfy a formula α. But, we first examine an aspect of modal logic which does not have any counterpart in propositional logic.

2.3 Correspondence Theory

The modalities $\Box$ and $\Diamond$ can be used to describe interesting properties of the accessibility relation R of a frame $F = (W, R)$. This area of modal logic is called *correspondence theory.*

Let α be a formula of modal logic. With α, we identify a class of frames $\mathcal{C}_\alpha$ as follows:

> $F = (W, R) \in \mathcal{C}_\alpha$ iff for every valuation V over W, for every world $w \in W$ and for every substitution instance β of α, $((W, R), V), w \models \beta$.

In other words, when defining $\mathcal{C}_\alpha$, we interpret α as a template, much like an axiom scheme. Notice that for any frame $F = (W, R)$ which does not belong to $\mathcal{C}_\alpha$, we can find a valuation V, a world w and a substitution instance β of α such that $((W, R), V), w \not\models \beta$.

Characterising classes of frames We say a class of frames $\mathcal{C}$ is *characterised* by the formula α if $\mathcal{C} = \mathcal{C}_\alpha$.

We now look at some examples of frame conditions which can characterised by formulas of modal logic.

Proposition 2.4 *The class of reflexive frames is characterised by the formula* $\Box\alpha \supset \alpha$.

Proof: We first show that every reflexive frame belongs to $\mathcal{C}_{\Box\alpha\supset\alpha}$. Let $M = ((W, R), V)$ be a model where R is reflexive. Consider any world $w \in W$. Suppose that $M, w \models \Box\alpha$. We have to show that $M, w \models \alpha$ as well. Since $M, w \models \Box\alpha$, every R-neighbour of w satisfies α. But R is reflexive, so w is an R-neighbour of itself. Hence, $M, w \models \alpha$.

Conversely, we show that every non-reflexive frame does not belong to $\mathcal{C}_{\Box\alpha\supset\alpha}$. Let $F = (W, R)$ be a frame where for some $w \in W$, it is not the case that $w\ R\ w$. Choose a proposition p and define a valuation V as follows: $V(w) = \emptyset$ and $V(w') = \{p\}$ for all $w' \neq w$. Clearly, $(F, V), w \models \Box p$ but $(F, V), w \not\models p$. Hence w fails to satisfy the substitution instance $\Box p \supset p$ of the formula $\Box\alpha \supset \alpha$. $\Box$

Proposition 2.5 *The class of transitive frames is characterised by the formula* $\Box\alpha \supset \Box\Box\alpha$.

Proof: We first show that every transitive frame belongs to $\mathcal{C}_{\Box\alpha\supset\Box\Box\alpha}$. Let $M = ((W, R), V)$ be a model where R is transitive. Consider any world $w \in W$. Suppose that $M, w \models \Box\alpha$. We have to show that $M, w \models \Box\Box\alpha$ as well.

For this, we have to show that every R-neighbour w' of w satisfies $\Box\alpha$. Consider any R-neighbour w' of w. If w' has no R-neighbours, then it is trivially the case that $M, w' \models \Box\alpha$. On the other hand, if w' has R-neighbours, then we must show that each R-neighbour of w' satisfies α. Let w'' be an R-neighbour of w'. Since $w\ R\ w'$ and $w'\ R\ w''$, by transitivity w'' is also an R-neighbour of w. Since we assumed that $M, w \models \Box\alpha$, it must be the case that $M, w'' \models \alpha$, as required.

Conversely, we show that every non-transitive frame does not belong to $\mathcal{C}_{\Box\alpha\supset\Box\Box\alpha}$. Let $F = (W, R)$ be a frame where for some $w, w', w'' \in W$, $w\ R\ w'$ and $w'\ R\ w''$ but it is not the case that that $w\ R\ w''$. Choose a proposition p and define a valuation V as follows:

$$V(\hat{w}) = \begin{cases} \{p\} & \text{if } w\ R\ \hat{w} \\ \emptyset & \text{otherwise} \end{cases}$$

Since w'' is not an R-neighbour of w, $V(w'') = \emptyset$. This means that $M, w' \not\models \Box p$, for w'' is an R-neighbour of w' and $M, w'' \not\models p$. Therefore, $M, w \not\models \Box\Box p$, since w' is an R-neighbour of w. On the other hand, $M, w \models \Box p$ by the definition of V. Hence, $M, w \not\models \Box p \supset \Box\Box p$, which is an instance of $\Box\alpha \supset \Box\Box\alpha$. $\Box$

The characteristic formula for transitivity can dually be written $\Diamond\Diamond\alpha \supset \Diamond\alpha$. This form represents transitivity more intuitively—the formula says that if $w\ R\ w'\ R\ w''$ and w'' satisfies α, there exists an R-neighbour $\hat{w}$ of w satisfying α. If R is transitive, w'' is a natural candidate for $\hat{w}$. Similarly, $\alpha \supset \Diamond\alpha$ is the dual (and more appealing) form of the characteristic formula for reflexivity. We have used the $\Box$ forms of these formulas because they are more standard in the literature.

Proposition 2.6 *The class of symmetric frames is characterised by the formula* $\alpha \supset \Box\Diamond\alpha$.

Proof: We first show that every symmetric frame belongs to $\mathcal{C}_{\alpha\supset\Box\Diamond\alpha}$. Let $M = ((W, R), V)$ be a model where R is symmetric. Consider any world $w \in W$. Suppose that $M, w \models \alpha$. We have to show that $M, w \models \Box\Diamond\alpha$ as well.

For this, we have to show that every R-neighbour w' of w satisfies $\Diamond\alpha$. Consider any R-neighbour w' of w. Since R is symmetric, w is an R-neighbour of w'. We assumed that $M, w \models \alpha$ so $M, w' \models \Diamond\alpha$, as required.

Conversely, we show that every non-symmetric frame does not belong to $\mathcal{C}_{\alpha\supset\Box\Diamond\alpha}$. Let $F = (W, R)$ be a frame where for some $w, w' \in W$, $w\ R\ w'$ but it is not the case that that $w'\ R\ w$. Choose a proposition p and define a valuation V as follows:

$$V(\hat{w}) = \begin{cases} \emptyset & \text{if } w'\ R\ \hat{w} \\ \{p\} & \text{otherwise} \end{cases}$$

By construction $M, w' \not\models \Diamond p$. Hence, since wRw', $M, w \not\models \Box\Diamond p$. On the other hand, $M, w \models p$ by the definition of V, so $M, w \not\models p \supset \Box\Diamond p$, which is an instance of the formula $\alpha \supset \Box\Diamond\alpha$. $\Box$

We say that an accessibility relation R over W is *Euclidean* if for all $w, w', w'' \in W$, if $w\ R\ w'$ and $w\ R\ w''$ then $w'\ R\ w''$ and $w''\ R\ w'$ (see Figure 2).

Proposition 2.7 *The class of Euclidean frames is characterised by the formula* $\Diamond\alpha \supset \Box\Diamond\alpha$.

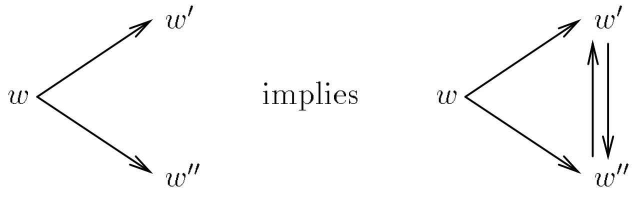

Figure 2: The Euclidean condition

Proof: We first show that every Euclidean frame belongs to $\mathcal{C}_{\Diamond\alpha\supset\Box\Diamond\alpha}$. Let $M = ((W, R), V)$ be a model where R is Euclidean. Consider any world $w \in W$. Suppose that $M, w \models \Diamond\alpha$. We have to show that $M, w \models \Box\Diamond\alpha$ as well.

Let w' be an R-neighbour of w. We must show that $M, w' \models \Diamond\alpha$. Since $M, w \models \Diamond\alpha$, there must exist w_α such that $w\ R\ w_\alpha$ and $M, w_\alpha \models \alpha$. Since R is Euclidean, $w'\ R\ w_\alpha$ as well, so $M, w' \models \Diamond\alpha$ as required.

Conversely, we show that every non-Euclidean frame does not belong to $\mathcal{C}_{\Diamond\alpha\supset\Box\Diamond\alpha}$. Let $F = (W, R)$ be a frame where for some $w, w', w'' \in W$, $w\ R\ w'$ and $w\ R\ w''$ but one of $w'\ R\ w''$ and $w''\ R\ w'$ fails to hold. Without loss of generality, assume that it is not the case that $w''\ R\ w'$.

Choose a proposition p and define a valuation V such that $V(w') = \{p\}$ and $V(\hat{w}) = \emptyset$ for all $\hat{w} \neq w'$. Then, since $w\ R\ w'$, $M, w \models \Diamond p$ by the definition of V. On the other hand, by construction $M, w'' \not\models \Diamond p$, so $M, w \not\models \Box\Diamond p$. So, $M, w \not\models \Diamond p \supset \Box\Diamond p$, which is an instance of $\Diamond\alpha \supset \Box\Diamond\alpha$. $\Box$

Notice that if R is Euclidean, for all w', if there exists w such that $w\ R\ w'$, then $w'\ R\ w'$. It is not difficult to verify that if R is reflexive and Euclidean then R is in fact an equivalence relation.

Exercise 2.8 What classes of frames are characterised by the following formulas?

(i) $\Diamond\alpha \supset \Box\alpha$.

(ii) $\Diamond\alpha \supset \Diamond\Diamond\alpha$.

(iii) $\alpha \supset \Box\alpha$.

$\Box$

Are there natural classes of frames which *cannot* be characterised in modal logic? We will see later that irreflexive frames form one such class. But first, we return to the notions of satisfiability and validity and look for a completeness result.

2.4 Axiomatising valid formulas

Validity revisited We said earlier that a formula α is valid if for every frame $F = (W, R)$, every model $M = (F, V)$ and every world w, $M, w \models \alpha$. In light of our discussion of correspondence theory we can refine this notion by restricting the range over which we consider frames.

Let $\mathcal{C}$ be a class of frames. We say that a formula α is $\mathcal{C}$-valid if for every frame $F = (W, R)$ from the class $\mathcal{C}$, for every model $M = (F, V)$ and for every world w, $M, w \models \alpha$. We denote the fact that α is $\mathcal{C}$-valid by $\models_{\mathcal{C}} \alpha$.

Let $\mathcal{F}$ represents the class of all frames. Then, the set of $\mathcal{F}$-valid formulas is the same as the set of valid formulas according to our earlier definition. In other words, the notions $\models_{\mathcal{F}} \alpha$ and $\models \alpha$ are equivalent.

Dually, we say that a formula α is $\mathcal{C}$-satisfiable if there is a frame $F = (W, R)$ in the class $\mathcal{C}$, a model $M = (F, V)$ and a world w, such that $M, w \models \alpha$. Once again, a formula is $\mathcal{F}$-satisfiable iff it is satisfiable according to our earlier definition.

Completeness for the class $\mathcal{F}$

Consider the following axiom system.

Axiom System $\boldsymbol{K}$

Axioms

(A0) All tautologies of propositional logic.

(K) $\Box(\alpha \supset \beta) \supset (\Box\alpha \supset \Box\beta)$.

Inference Rules

(MP) $\dfrac{\alpha, \alpha \supset \beta}{\beta}$ (G) $\dfrac{\alpha}{\Box\alpha}$

The axiom (A0) is an abbreviation for any set of axioms which are sound and complete for Propositional Logic—in particular, we could instantiate (A0) with the axioms (A1)–(A3) of the system AX discussed in the previous section.

As usual, we say that α is a *thesis* of System K[3], denoted $\vdash_K \alpha$, if we can derive α using the axioms (A0) and (K) and the inference rules (MP) and (G). Once again, we will omit the subscript and write $\vdash \alpha$ if there is no confusion about which axiom system we are referring to.

The result we want to establish is the following.

Theorem 2.9 *For all formulas α, $\vdash_K \alpha$ iff $\models_{\mathcal{F}} \alpha$.*

As usual, one direction of the proof is easy.

Lemma 2.10 (Soundness of System K) *If $\vdash_K \alpha$ then $\models_{\mathcal{F}} \alpha$.*

[3]The name K is derived from Saul Kripke.

Proof: As we observed in the previous section, it suffices to show that each axiom is $\mathcal{F}$-valid and that the inference rules preserve $\mathcal{F}$-validity. This is precisely what we exhibited in Example 2.2 and Exercise 2.3. $\square$

As in Propositional Logic, we use a Henkin-style argument to show that every $\mathcal{F}$-valid formula is derivable using System K.

Consistency As before, we say that a formula α is *consistent* with respect to System K if $\not\vdash_K \neg\alpha$. A finite set of formulas $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ is consistent if the conjunction $\alpha_1 \wedge \alpha_2 \wedge \cdots \wedge \alpha_n$ is consistent. Finally, an arbitrary set of formulas X is consistent if every finite subset of X is consistent.

Our goal is to prove the following.

Lemma 2.11 *Let α be a formula which is consistent with respect to System K. Then, α is $\mathcal{F}$-satisfiable.*

As we saw in the case of Propositional Logic, this will yield as an immediate corollary the result which we seek:

Corollary 2.12 (Completeness for System *K*) *Let α be a formula which is $\mathcal{F}$-valid. Then, $\vdash_K \alpha$.*

Maximal Consistent Sets

As before, we say that a set of formulas X is a *maximal consistent set* or MCS if X is consistent and for all $\alpha \notin X$, $X \cup \{\alpha\}$ is inconsistent. As we saw earlier, by Lindenbaum's Lemma, every consistent set of formulas can be extended to an MCS.

We will once again use the properties of MCSs established in Lemma 1.14. In addition, the following properties of MCSs will prove useful.

Lemma 2.13 *Let X be a maximal consistent set.*

(i) If β is a substitution instance of an axiom, then $\beta \in X$.

(ii) If $\alpha \supset \beta \in X$ and $\alpha \in X$, then $\beta \in X$.

Proof: The proof is routine and is left as an exercise. $\square$

The canonical model

When we studied propositional logic, we saw that each maximal consistent set defines a "propositional world". In modal logic, we have to construct frames with many propositional worlds. In fact, we generate a frame with *all* possible worlds, with a suitable accessibility relation.

Canonical model The *canonical frame* for System K is the pair $F_K = (W_K, R_K)$ where:

- $W_K = \{X \mid X \text{ is an MCS}\}$.
- If X and Y are MCSs, then $X\ R_K\ Y$ iff $\{\alpha \mid \Box\alpha \in X\} \subseteq Y$.

The *canonical model* for System K is given by $M_K = (F_K, V_K)$ where for each $X \in W_K$, $V_K(X) = X \cap \mathcal{P}$.

Exercise 2.14 We can dually define R_K using the modality $\Diamond$ rather than $\Box$. Verify that $X\ R_K\ Y$ iff $\{\Diamond\alpha \mid \alpha \in Y\} \subseteq X$. $\Box$

The heart of the completeness proof is the following lemma.

Lemma 2.15 *For each MCS $X \in W_K$ and for each formula $\alpha \in \Phi$, $M_K, X \models \alpha$ iff $\alpha \in X$.*

Proof: As usual, the proof is by induction on the structure of α.
Basis: If $\alpha = p \in \mathcal{P}$, then $M_K, X \models p$ iff $p \in V_K(X)$ iff $p \in X$, by the definition of V_K.

Induction step:

$\alpha = \neg\beta$: Then $M_K, X \models \neg\beta$ iff $M_K, X \not\models \beta$ iff (by the induction hypothesis) $\beta \notin X$ iff (by the fact that X is an MCS) $\neg\beta \in X$.

$\alpha = \beta \vee \gamma$: Then $M_K, X \models \beta \vee \gamma$ iff $M_K, X \models \beta$ or $M_K, X \models \gamma$ iff (by the induction hypothesis) $\beta \in X$ or $\gamma \in X$ iff (by the fact that X is an MCS) $\beta \vee \gamma \in X$.

$\alpha = \Box\beta$: We analyse this case in two parts:

($\Leftarrow$) Suppose that $\Box\beta \in X$. We have to show that $M_K, X \models \Box\beta$. Consider any MCS Y such that $X\ R_K\ Y$. Since $\Box\beta \in X$, from the definition of R_K it follows that $\beta \in Y$. By the induction hypothesis $M_K, Y \models \beta$. Since the choice of Y was arbitrary, $M_K, X \models \Box\beta$.

($\Rightarrow$) Suppose that $M_K, X \models \Box\beta$. We have to show that $\Box\beta \in X$. Suppose that $\Box\beta \notin X$. Then, since X is an MCS, $\neg\Box\beta \in X$. We show that this leads to a contradiction.

Claim $Y_0 = \{\gamma \mid \Box\gamma \in X\} \cup \{\neg\beta\}$ is consistent.

If we assume the claim, we can extend Y_0 to an MCS Y. Clearly, $X\ R_K\ Y$. Since $\neg\beta \in Y$, $\beta \notin Y$. By the induction hypothesis, $M_K, Y \not\models \beta$. This means that $M_K, X \not\models \Box\beta$ which contradicts our initial assumption that $M_K, X \models \Box\beta$.

To complete the proof, we must verify the claim.

Proof of claim Suppose that Y_0 is not consistent. Then, there exists a finite subset $\{\gamma_1, \gamma_2, \ldots, \gamma_n\}$ of Y_0 such that $\gamma_1 \wedge \gamma_2 \wedge \cdots \wedge \gamma_n \wedge \neg\beta$ is inconsistent. Let us denote $\gamma_1 \wedge \gamma_2 \wedge \cdots \wedge \gamma_n$ by $\tilde{\gamma}$.

We then have the following sequence of derivations:

$\vdash \neg(\tilde{\gamma} \wedge \neg\beta)$	By the definition of consistency
$\vdash \neg\tilde{\gamma} \vee \beta$	Tautology of propositional logic (Axiom A0)
$\vdash \tilde{\gamma} \supset \beta$	Definition of $\supset$
$\vdash \Box(\tilde{\gamma} \supset \beta)$	Inference rule G
$\vdash \Box\tilde{\gamma} \supset \Box\beta$	Axiom K plus one application of MP
$\vdash \neg(\Box\tilde{\gamma} \wedge \neg\Box\beta)$	Tautology of propositional logic (Axiom A0)

We can easily show that $\vdash \Box(\gamma \wedge \delta) \equiv (\Box\gamma \wedge \Box\delta)$.

In one direction, since $\vdash \gamma \wedge \delta \supset \gamma$ is a tautology of propositional logic, we can use the rule G to get $\vdash \Box(\gamma \wedge \delta \supset \gamma)$. From axiom K and one application of MP, $\vdash \Box(\gamma \wedge \delta) \supset \Box\gamma$. Symmetrically, it follows that $\vdash \Box(\gamma \wedge \delta) \supset \Box\delta$. So, $\vdash \Box(\gamma \wedge \delta) \supset (\Box\gamma \wedge \Box\delta)$.

Conversely, $\vdash \gamma \supset (\delta \supset (\gamma \wedge \delta))$ from propositional logic. By applying axiom K and MP a couple of times, we obtain $\vdash \Box\gamma \supset (\Box\delta \supset \Box(\gamma \wedge \delta))$, from which it follows that $\vdash (\Box\gamma \wedge \Box\delta) \supset \Box(\gamma \wedge \delta)$.

We can extend this argument to show that $\vdash \Box(\delta_1 \wedge \delta_2 \wedge \cdots \wedge \delta_n) \equiv (\Box\delta_1 \wedge \Box\delta_2 \wedge \cdots \Box\delta_n)$ for all n.

From the last line in our derivation above, it then follows that $\vdash \neg(\Box\gamma_1 \wedge \Box\gamma_2 \wedge \cdots \Box\gamma_n \wedge \neg\Box\beta)$. Thus the set $\{\Box\gamma_1, \Box\gamma_2, \ldots, \Box\gamma_n, \neg\Box\beta\}$ is inconsistent. But this is a finite subset of X, which means that X is itself inconsistent, contradicting the fact that X is an MCS.

$\Box$

From the preceding result, the proof of Lemma 2.11 is immediate.

Proof: (of Lemma 2.11) Let α be a formula which is consistent with respect to System K. By Lindenbaum's Lemma, α can be extended to a maximal consistent set X_α. By the preceding result $M, X_\alpha \models \alpha$, so α is $\mathcal{F}$-satisfiable. $\Box$

Once we have proved Lemma 2.11, we immediately obtain a proof of completeness (Corollary 2.12) using exactly the same argument as in propositional logic.

It is worth pointing out one important difference between the canonical model constructed for System K and the models constructed when proving completeness for propositional logic. In propositional logic, to satisfy a consistent formula α, we build a valuation v which depends on α. On the other hand, the construction of the canonical model for System K is *independent* of the choice of α. Thus, *every* consistent formula α is satisfied within the model M_K.

Completeness for other classes of frames

Can we axiomatise the set of $\mathcal{C}$-valid formulas for a class of frames $\mathcal{C}$ which is properly included in $\mathcal{F}$? To do this, we use the characteristic formulas which we looked at when discussing correspondence theory.

Reflexive frames

System T is the set of axioms obtained by adding the following axiom scheme to System K.

(T) $\Box\alpha \supset \alpha$

Lemma 2.16 *System T is sound and complete with respect to the class of reflexive frames.*

Proof: To show that System T is sound with respect to reflexive frames, we only need to verify that the new axiom (T) is sound for this class of frames—the other axioms and rules from System K continue to be sound. The soundness of axiom (T) follows from Proposition 2.4.

To show completeness, we must argue that every formula which is consistent with respect to System T can be satisfied in a model based on a reflexive frame. To establish this, we follow the proof of completeness for System K and build a canonical model $M_T = ((W_T, R_T), V_T)$ for System T which satisfies the property described in Lemma 2.11. We just need to verify that the resulting frame (W_T, R_T) is reflexive.

For any MCS X, we need to verify that $X \; R_T \; X$ or, in other words, that $\{\alpha \mid \Box\alpha \in X\} \subseteq X$. Consider any formula $\Box\alpha \in X$. Since $\Box\alpha \supset \alpha$ is an axiom of System T, $\Box\alpha \supset \alpha \in X$, by Lemma 2.13 (i). From Lemma 2.13 (ii), it then follows that $\alpha \in X$, as required. $\Box$

Transitive frames

System 4 is the set of axioms obtained by adding the following axiom scheme to System K.

(4) $\Box\alpha \supset \Box\Box\alpha$

Lemma 2.17 *System 4 is sound and complete with respect to the class of transitive frames.*

Proof: We know that the axiom (4) is sound for the class of transitive frames from Proposition 2.5. This establishes the soundness of System 4.

To show completeness, we must argue that every formula which is consistent with respect to System 4 can be satisfied in a model based on a transitive frame. Once again, we can build a canonical model $M_4 = ((W_4, R_4), V_4)$ for System 4 which satisfies the property described in Lemma 2.11. We just need to verify that the resulting frame (W_4, R_4) is transitive.

In other words, if X, Y, Z are MCSs such that $X \; R_4 \; Y$ and $Y \; R_4 \; Z$, we need to verify that $X \; R_T \; Z$—that is, we must show that $\{\alpha \mid \Box\alpha \in X\} \subseteq Z$. Consider any formula $\Box\alpha \in X$. Since $\Box\alpha \supset \Box\Box\alpha$ is an axiom of System 4, it follows from Lemma 2.13 that $\Box\Box\alpha \in X$. Since $X \; R_4 \; Y$, it must be the case that $\Box\alpha \in Y$. Further, since $Y \; R_4 \; Z$ it must be the case that $\alpha \in Z$, as required. $\Box$

Exercise 2.18 The System B is obtained by adding the following axiom to System K.

(B) $\alpha \supset \Box\Diamond\alpha$.

Verify that System B is sound and complete with respect to symmetric frames. $\Box$

Combinations of frame conditions

By combining the characteristic formulas for different frame conditions, we obtain completeness for smaller classes of frames.

Reflexive and transitive frames

The System *S4* is obtained by adding the axioms (T) (for reflexivity) and (4) (for transitivity) to System K.

Lemma 2.19 *System S4 is sound and complete with respect to the class of reflexive and transitive frames.*

Proof: Since System T is sound for the class of reflexive frames and System *4* is sound for the class of transitive frames, it follows that System *S4* is sound for the class of reflexive *and* transitive frames.

To show completeness, as usual we build a canonical model $M_{S4} = ((W_{S4}, R_{S4}), V_{S4})$ satisfying the property in Lemma 2.11. Using the argument in the proof of Lemma 2.16, it follows that R_{S4} is reflexive. Similarly, from the proof of Lemma 2.17 it follows that R_{S4} is transitive. $\Box$

Equivalence relations

The System *S5* is obtained by adding the following axioms to System K.

(T) $\Box\alpha \supset \alpha$

(5) $\Diamond\alpha \supset \Box\Diamond\alpha$.

We have already seen that (T) is the axiom for reflexivity, while (5) characterises Euclidean frames.

Exercise 2.20

(i) Show that System *S5* is sound and complete for the class of frames whose accessibility relation is an equivalence relation.

(ii) Show that the axioms (4) and (B) can be derived in System *S5*.

$\Box$

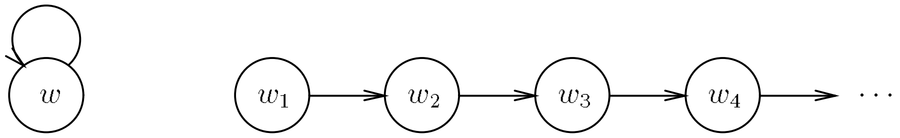

Figure 3: A pair of similar frames

2.5 Bisimulations and expressiveness

Intuitively, it is clear that models which have "similar" structure satisfy the same modal logic formulas. For instance, if we choose the same valuation for all worlds in the two frames shown in Figure 3, it seems evident that no formula can distinguish the resulting pair of models.

To formalise this notion, we introduce bisimulations.

Bisimulation Let $M_1 = ((W_1, R_1), V_1)$ and $M_2 = ((W_2, R_2), V_2)$ be a pair of models. A *bisimulation* is a relation $\sim\ \subseteq W_1 \times W_2$ satisfying the following conditions.

(i) If $w_1 \sim w_2$ and $w_1 R_1 w'_1$ then there exists w'_2 such that $w_2 R_2 w'_2$ and $w'_1 \sim w'_2$.

(ii) If $w_1 \sim w_2$ and $w_2 R_2 w'_2$ then there exists w'_1 such that $w_1 R_1 w'_1$ and $w'_1 \sim w'_2$.

(iii) If $w_1 \sim w_2$ then $V_1(w_1) = V_2(w_2)$.

Notice that the empty relation is a trivial example of a bisimulation. Two worlds which are related by a bisimulation satisfy exactly the same formulas.

Lemma 2.21 *Let $\sim$ be a bisimulation between $M_1 = ((W_1, R_1), V_1)$ and $M_2 = ((W_2, R_2), V_2)$. For all $w_1 \in W_1$ and $w_2 \in W_2$, if $w_1 \sim w_2$, then for all formulas α, $M_1, w_1 \models \alpha$ iff $M_2, w_2 \models \alpha$.*

Proof: As usual, the proof is by induction on the structure of α.

Basis: Suppose $\alpha = p \in \mathcal{P}$. By the definition of bisimulations, we know that $V_1(w_1) = V_2(w_2)$. Hence, $M_1, w_1 \models p$ iff $M_2, w_2 \models p$.

Induction step: The propositional cases $\alpha = \neg\beta$ and $\alpha = \beta \vee \gamma$ are easy, so we omit them and directly consider the case $\alpha = \Box\beta$.

($\Rightarrow$) Suppose that $M_1, w_1 \models \Box\beta$. We must show that $M_2, w_2 \models \Box\beta$ as well. For this, we must argue that $M_2, w'_2 \models \beta$ for each world w'_2 such that $w_2 R_2 w'_2$. Since $\sim$ is a bisimulation, for each such w'_2 there exists a world w'_1 such that $w_1 R_1 w'_1$ and $w'_1 \sim w'_2$. Since $M_1, w_1 \models \Box\beta$, it follows that $M_1, w'_1 \models \beta$. Since $w'_1 \sim w'_2$, by the induction hypothesis, it follows that $M_2, w'_2 \models \beta$. Since w'_2 was an arbitrarily chosen R_2-neighbour of w_2, we have $M_2, w_2 \models \Box\beta$, as required.

($\Leftarrow$) Suppose that $M_2, w_2 \models \Box\beta$. We must show that $M_1, w_1 \models \Box\beta$ as well. The argument is symmetric to the earlier one and we omit the details.

□

We can use bisimulations to show that certain classes of frames *cannot* be characterised in modal logic.

Lemma 2.22 *The class of irreflexive frames cannot be characterised in modal logic.*

Proof: Let α be a formula that characterises the class of irreflexive frames. Consider the pair of frames in Figure 3. Since the first frame is not irreflexive, there should be a valuation V and an instance β of α such that β is not satisfied at w under V.

Let us define a valuation V' on the second model such that for each w_i, $V'(w_i) = V(w)$. We can clearly set up a bisimulation between the two models by relating w to each of the worlds w_i. This means that w satisfies exactly the same formulas as each of the worlds w_i. In particular, β is not satisfied at each w_i. This is a contradiction because the second model is irreflexive and β is an instance of the formula α which we claimed was a characteristic formula for irreflexive frames. □

Exercise 2.23 We say that a frame (W, R) is "non-connected" if there are worlds w and w' such that it is not the case that $w(R \cup R^{-1})^* w'$. In other words, we convert (W, R) into an undirected graph by ignoring the orientation of edges in R. The frame is "non-connected" if there are two nodes in the resulting undirected graph which are not reachable from each other.

Show that there is no axiom which characterises the class of "non-connected" frames. □

Antisymmetry

We have seen that irreflexivity cannot be characterised in modal logic. Another natural frame condition which is beyond the expressive power of modal logic is *antisymmetry*. Recall that a relation R on W is antisymmetric if $w\ R\ w'$ and $w'\ R\ w$ imply that $w = w'$.

Lemma 2.24 *Let α be a formula which is satisfiable over the class of reflexive and transitive frames. Then, α is satisfiable in a model based on an reflexive, transitive and antisymmetric frame.*

Proof: Let $M = ((W, R), V)$ be a model where R is reflexive and transitive. We describe a technique called *bulldozing*, due to Krister Segerberg, for constructing a new model $\hat{M} = ((\hat{W}, \hat{R}), \hat{V})$, where $\hat{R}$ is reflexive, transitive and antisymmetric, such that $\hat{M}$ and M satisfy the *same* formulas.

Consider the frame (W, R). If R is not antisymmetric, there are two worlds w and w' in W such that $w\ R\ w'$ and $w'\ R\ w$. The idea is to break each loop of this kind by making infinitely many copies of w and w' and arranging these copies alternately in a chain. We then verify that the new model which we construct is bisimilar to the original model.

Formally, we say that $X \subseteq W$ is a *cluster* if $X \times X \subseteq R$—in a cluster, every world can "see" every other world.

Let Cl be the class of *maximal* clusters of W—that is, $X \in Cl$ if X is cluster and for each $w \notin X$, $(X \cup \{w\}) \times (X \cup \{w\}) \not\subseteq R$. Since R is reflexive, every singleton $\{w\}$ is a cluster. It follows that the set Cl of maximal clusters is not empty and that every world $w \in W$ belongs to some maximal cluster in Cl. In fact, W is partitioned into maximal clusters.

For each $X \in Cl$, define $W_X = X \times \mathbb{N}$, where $\mathbb{N}$ is the set $\{0, 1, 2, \ldots\}$ of natural numbers. Thus W_X contains infinitely many copies of each world from X. For each set W_X, we define an accessibility relation within W_X. For this, we first fix an arbitrary total order $\leq_X$ on X. For $X \in Cl$, $R_X \subseteq W_X \times W_X$ is then defined as follows:

$$\begin{array}{rrl} R_X = & & \{((w,i),(w,i)) \mid w \in X \text{ and } i \in \mathbb{N}\} \\ & \cup & \{((w,i),(w',i)) \mid w, w' \in X \text{ and } w \leq_X w'\} \\ & \cup & \{((w,i),(w',j)) \mid w, w' \in X \text{ and } i < j\} \end{array}$$

We then define a relation across maximal clusters based on the original accessibility relation R:

$$R' = \bigcup\{(W_X \times W_Y) \mid X \neq Y \text{ and for some } w \in X \text{ and } w' \in Y,\ w\ R\ w'\}$$

Finally, we can define the new frame $(\hat{W}, \hat{R})$ corresponding to (W, R).

- $\hat{W} = \bigcup_{X \in Cl} W_X$.
- $\hat{R} = R' \cup \bigcup_{X \in Cl} R_X$.

It can be verified that $\hat{R}$ is reflexive, transitive and antisymmetric (Exercise 2.25).

Each world in $\hat{W}$ is of the form (w, i) where $w \in X$ for some maximal cluster $X \in Cl$ and $i \in \mathbb{N}$. We extend $(\hat{W}, \hat{R})$ to a model by defining $\hat{V}((w,i)) = V(w)$ for all $w \in W$ and $i \in \mathbb{N}$.

We define a relation $\sim\ \subseteq \hat{W} \times W$ as follows:

$$\sim\ = \{((w,i), w) \mid w \in W,\ i \in \mathbb{N}\}$$

We claim that $\sim$ is a bisimulation between $\hat{M}$ and M. From the definition of $\hat{V}$, we have $\hat{V}((w,i)) = V(w)$ for all $w \in W$ and $i \in \mathbb{N}$, so the third condition in the definition of bisimulations is trivially satisfied.

Suppose that $(w,i) \sim w$ and $(w,i)\ \hat{R}\ (w',j)$. We must show that $w\ R\ w'$. If w and w' belong to the same maximal cluster X, then $w\ R\ w'$ because all elements in X are R-neighbours of each other. On the other hand, if $w \in X$ and $w' \in Y$ for distinct clusters X and Y, it must be the case that $(w,i)\ R'\ (w',j)$. This means that we have $w_1 \in X$ and $w_1' \in Y$ such that $w_1\ R\ w_1'$. Since $w\ R\ w_1$ and $w_1'\ R\ w'$, from the transitivity of R it follows that $w\ R\ w'$.

Conversely, suppose that $(w,i) \sim w$ and $w \; R \; w'$. We must exhibit a world (w',j) such that $(w,i) \; \hat{R} \; (w',j)$. If w and w' belong to the same maximal cluster X, we just choose (w',j) such that $i < j$. Then, by the definition of R_X, $(w,i) \; R_X \; (w',j)$, so $(w,i) \; \hat{R} \; (w',j)$ as well. On the other hand, if $w \in X$ and $w' \in Y$ for distinct maximal clusters X and Y, then $(x,i) \; R' \; (y,j)$ for all $j \in \mathbb{N}$, so once again we can pick a (w',j) such that $(x,i) \; \hat{R} \; (y,j)$.

Thus, $\sim$ is a bisimulation between $\hat{M}$, whose frame is a antisymmetric and transitive, and M, whose frame is transitive. Hence, for any world $w \in W$ and any formula α, $M, w \models \alpha$ iff $\hat{M}, (w,i) \models \alpha$ for all $i \in \mathbb{N}$. In other words, every formula which is satisfiable in the class of transitive frames is also satisfiable in the class of antisymmetric and transitive frames. $\Box$

Exercise 2.25 Show that the relation $\hat{R}$ constructed in the proof of Lemma 2.24 is reflexive, transitive and antisymmetric. $\Box$

Corollary 2.26 *The class of antisymmetric frames cannot be characterised in modal logic.*

Proof: Let α be a formula characterizing the class of antisymmetric frames. Let (W,R) be a frame where R is reflexive and transitive but not antisymmetric. Then, there exists an instance β of α and a valuation V over (W,R) such that $M, w \models \neg\beta$ for some $w \in W$. By Lemma 2.24, we can convert M into a model $\hat{M} = ((\hat{W}, \hat{R}), \hat{V})$ where $\hat{R}$ is reflexive, transitive and antisymmetric, such that $M, \hat{w} \models \neg\beta$ for some $\hat{w} \in \hat{W}$. This is a contradiction, since β was assumed to be an instance of the formula α which characterises antisymmetric frames. $\Box$

We have already seen that the system *S4* is sound and complete for the class of reflexive, transitive frames. This class is very close to the class of partial orders, which are ubiquitous in computer science. The fact that antisymmetry cannot be characterised in modal logic means that modal logic cannot distinguish between reflexive and transitive frames (often called *preorders*) and reflexive, transitive and antisymmetric frames (or *partial orders*).

Corollary 2.27 *The system S4 is sound and complete for the class of partial orders.*

Proof: Since partial orders are reflexive and transitive, *S4* is certainly sound for this class of frames. We already know that every formula which is consistent with respect to *S4* is satisfiable in a preorder. The bulldozing construction described in the proof of Lemma 2.24 shows that every formula satisfiable over a preorder is also satisfiable over a partial order. $\Box$

2.6 Decidability: Filtrations and the finite model property

Though we have looked at sound and complete axiomatisations of different classes of frames, we have yet to establish any results concerning decidability. The basic technique for showing decidability is to prove that any formula which is satisfiable is in fact satisfiable in a *finite* model.

Finite model property Let A be an axiom system which is sound and complete with respect to a class of frames $\mathcal{C}$. The system A has the *finite model property* if for all formulas α, $\not\vdash_A \alpha$ implies there is a model $M = (F, V)$ based on a finite frame $F = (W, R) \in \mathcal{C}$ such that for some $w \in W$, $M, w \models \neg\alpha$.

Since A is sound and complete for the class $\mathcal{C}$, this is equivalent to demanding that any formula which is satisfiable in the class $\mathcal{C}$ is in fact satisfiable in a model based on a finite frame from the class $\mathcal{C}$.

Assume that we can effectively decide whether or not a given finite frame belongs to the class $\mathcal{C}$, we can then systematically enumerate all finite models built from the class $\mathcal{C}$. As a consequence, the finite model property allows us to enumerate the set of formulas satisfiable within the class $\mathcal{C}$. On the other hand, the completeness of the axiom system A allows us to enumerate the set of formulas which are valid in this class of frames.

To check whether a formula α is valid, we interleave these enumerations. If α is valid, it will be enumerated as a thesis of the system A. On the other hand, if α is not valid, its negation $\neg\alpha$ must be satisfiable, so $\neg\alpha$ will appear in the enumeration of formulas satisfiable over $\mathcal{C}$. Thus, the finite model property yields a decision procedure for validity (and, dually, satisfiability).

Subformulas Let α be formula. The set of *subformulas of* α, denoted $\mathrm{sf}(\alpha)$, is the smallest set of formulas such that:

- $\alpha \in \mathrm{sf}(\alpha)$.
- If $\neg\beta \in \mathrm{sf}(\alpha)$ then $\beta \in \mathrm{sf}(\alpha)$.
- If $\beta \vee \gamma \in \mathrm{sf}(\alpha)$ then $\beta \in \mathrm{sf}(\alpha)$ and $\gamma \in \mathrm{sf}(\alpha)$.
- If $\Box\beta \in \mathrm{sf}(\alpha)$ then $\beta \in \mathrm{sf}(\alpha)$.

Exercise 2.28 Show that the size of the set $\mathrm{sf}(\alpha)$ is bounded by the length of α. More formally, for a formula α, define $|\alpha|$, the length of α, to be the number of symbols in α. Show that if $|\alpha| = n$ then $|\mathrm{sf}(\alpha)| \leq n$. Give an example where $|\mathrm{sf}(\alpha)| < |\alpha|$. $\Box$

For a set X of formulas, we write $\mathrm{sf}(X)$ to denote the set $\bigcup_{\alpha \in X} \mathrm{sf}(\alpha)$. A set of formulas X is said to be *subformula-closed* (or just *sf-closed*) if $X = \mathrm{sf}(X)$.

Let $M = ((W, R), V)$ and $M' = ((W', R'), V')$ be a pair of models. We have already seen that if we can set up a bisimulation $\sim$ between M and M', then for each pair of worlds $(w, w') \in \ \sim$, the worlds w and w' satisfy the same formulas. Often, we are willing to settle for a weaker relationship between w and w'—we do not require them to agree on *all* formulas, but only on formulas from a fixed set X. For sf-closed subsets X, this can be achieved using filtrations.

Filtrations Let $M = ((W, R), V)$ and $M' = ((W', R'), V')$ be a pair of models and X an sf-closed set of formulas. An *X-filtration* from M to M' is a function $f : W \to W'$ such that:

(i) For all $w, w' \in W$, if $w \; R \; w'$ then $f(w) \; R \; f(w')$.

(ii) The map f is surjective.

(iii) For all $p \in \mathcal{P} \cap X$, $p \in V(w)$ iff $p \in V'(f(w))$.

(iv) If $(f(w), f(w')) \in R'$, then for each formula of the form $\Box\alpha$ in X, if $M, w \models \Box\alpha$ then $M, w' \models \alpha$.

In a filtration, we have a weaker requirement on the inverse image of f than in a bisimulation. If $f(w)R'f(w')$, we do not demand that $w \; R \; w'$. We only insist that w and w' be "semantically" related upto the formulas in X. It is quite possible that $(w, w') \notin R$ and hence for some $\Box\beta \notin X$, $M, w \models \Box\beta$ while $M, w' \not\models \beta$.

Lemma 2.29 *Let f be an X-filtration from $M = ((W, R), V)$ to $M' = ((W', R'), V')$ where X is an sf-closed set of formulas. Then, for all $\alpha \in X$ and for all $w \in W$, $M, w \models \alpha$ iff $M', f(w) \models \alpha$.*

Proof: The proof is by induction on the structure of α.
Basis If $\alpha = p \in \mathcal{P} \cap X$, then $M, w \models p$ iff $p \in V(w)$ iff (by the definition of X-filtrations) $p \in V'(f(w))$ iff $M, f(w) \models p$.

Induction step The propositional cases $\alpha = \neg\beta$ and $\alpha = \beta \vee \gamma$ are easy, so we omit them and directly consider the case $\alpha = \Box\beta$.

($\Rightarrow$) Suppose $M, w \models \Box\beta$. To show that $M', f(w) \models \Box\beta$, we must show that for each w' with $f(w)R'w'$, $M', w' \models \beta$. Fix an arbitrary w' such that $f(w)R'w'$. Since f is surjective, there is a world $w'' \in W$ such that $w' = f(w'')$. From the last clause in the definition of filtrations, it follows that $M, w'' \models \beta$. Since X is sf-closed, $\beta \in X$. From the induction hypothesis, we have $M', f(w'') \models \beta$ or, in other words, $M', w' \models \beta$. Since w' was an arbitrary R'-neighbour of $f(w)$, it follows that $M', f(w) \models \Box\beta$.

($\Leftarrow$) Suppose that $M', f(w) \models \Box\beta$. To show that $M, w \models \Box\beta$, we must show that for each w' with $w \; R \; w'$, $M, w' \models \beta$. Fix an arbitrary w' such that $w \; R \; w'$. From the first clause in the definition of filtrations, it follows that $f(w)R'f(w')$. Since $M', f(w) \models \Box\beta$, it must be the case that $M', f(w') \models \beta$. Since $\beta \in X$, from the induction hypothesis we have $M, w' \models \beta$. Since w' was an arbitrary R-neighbour of w, it follows that $M, w \models \Box\beta$. $\Box$

Recall that our goal is to establish the finite model property for a class of frames $\mathcal{C}$—whenever a formula α is satisfiable over $\mathcal{C}$, then there is a model for α based on a *finite* frame from the class $\mathcal{C}$.

Our strategy will be as follows: given a formula α and an arbitrary model M for α, define an sf-closed set of formulas X_α and a finite model M_α such that $\alpha \in X_\alpha$ and there is an X_α-filtration from M to M_α. Lemma 2.29 then tells us that α is satisfied in M_α. Since this procedure applies uniformly to all satisfiable formulas α over the given class of frames, it follows that this class of frames has the finite model property.

Defining X_α is easy—we set $X_\alpha = \mathrm{sf}(\alpha)$. To construct M_α, we have to define a frame (W_α, R_α) and a valuation $V_\alpha : W_\alpha \to 2^{\mathcal{P}}$.

We define W_α and V_α in a uniform manner for all classes of frames. To define W_α, we begin with the following equivalence relation $\simeq_\alpha$ on W: $w \simeq_\alpha w'$ if for each $\beta \in X_\alpha$, $M, w \models \beta$ iff $M, w' \models \beta$. In other words, $w \simeq_\alpha w'$ iff the worlds w and w' satisfy exactly the same formulas from the set X_α. We use $[w]$ represent the equivalence class of w with respect to the relation $\simeq$—that is, $[w] = \{w' \mid w' \simeq_\alpha w\}$.

Let $W_\alpha = \{[w] \mid w \in W\}$. Observe that W_α is finite whenever X_α is finite. Since $X_\alpha = \mathrm{sf}(\alpha)$, we know that X_α is finite (recall Exercise 2.28).

Defining V_α is simple: for each $[w] \in W_\alpha$, $V_\alpha([w]) = \bigcap_{w' \in [w]} V(w')$.

Defining R_α is more tricky: in general, this relation has to be defined taking into account the class of frames under consideration. We now show how to define "suitable" R_α for some of the classes of frames for which we have already shown complete axiomatisations.

Lemma 2.30 *The axiom system K has the finite model property.*

Proof: Recall that system K is sound and complete for the class $\mathcal{F}$ of *all* frames. From our discussion of the finite model property, it suffices to show that any formula satisfiable over $\mathcal{F}$ is in fact satisfiable over a finite frame in $\mathcal{F}$.

Let α be a satisfiable formula and let $M = ((W, R), V)$ be a model for α—-that is, for some $w_\alpha \in W$, $M, w_\alpha \models \alpha$. Let $X_\alpha = \mathrm{sf}(\alpha)$ and define W_α and V_α as described earlier. Define R_α as follows:

$$R_\alpha = \{[w], [w'] \mid \text{For each formula } \beta \in X_\alpha, \text{ if } M, w \models \Box\beta \text{ then } M, w' \models \beta\}$$

Let $M_\alpha = ((W_\alpha, R_\alpha), V_\alpha)$.

Fix the function $f : W \to W_\alpha$ such that $w \mapsto [w]$ for each $w \in W$. We claim that f is an X_α-filtration from M to M_α—for this, we have to verify that f satisfies properties (i)–(iv) in the definition of filtrations.

It is clear that f is surjective (property (ii)).

To verify property (iii) we have to show that for each $p \in \mathcal{P} \cap X_\alpha$ and for each $w \in W$, $p \in V(w)$ iff $p \in V_\alpha([w])$. Since the worlds in $[w]$ agree on all formulas in X_α, it follows that $p \in V(w)$ iff for each $w' \simeq_\alpha w$, $p \in V(w')$ iff $p \in \bigcap_{w' \in [w]} V(w')$ iff (by the definition of V_α) $p \in V_\alpha([w])$.

Property (i) demands that $(w, w') \in R$ implies $([w], [w']) \in R_\alpha$. By the definition of R_α, $([w], [w']) \in R_\alpha$ if for each $\beta \in X_\alpha$, whenever $M, w \models \Box\beta$, $M, w' \models \beta$ as well. This is immediate from the fact that $(w, w') \in R$.

Finally, property (iv) states that whenever $([w], [w']) \in R_\alpha$, for each formula $\Box\beta \in X_\alpha$, if $M, w \models \Box\beta$ then $M, w' \models \beta$. This follows directly from the definition of R_α.

Having established that f is an X_α-filtration from M to M_α, it follows that $M_\alpha, [w_\alpha] \models \alpha$. Thus M_α is a finite model for α, as required. $\Box$

Lemma 2.31 *The axiom system T has the finite model property.*

Proof: Recall that system T is sound and complete for the class of reflexive frames. Let α be a formula satisfiable at a world w_α in a model $M = ((W, R), V)$ where (W, R) is a reflexive frame. We have to exhibit a finite model for α based on a reflexive frame.

Define X_α and $M_\alpha = ((W_\alpha, R_\alpha), V_\alpha)$ as in the proof of Lemma 2.30. We have already seen that $f : w \mapsto [w]$ then defines an X_α-filtration from M to M_α. To complete the proof of the present lemma, it suffices to show that the frame (W_α, R_α) is reflexive.

Since R is reflexive, we have $(w, w) \in R$ for each $w \in W$. By property (i) of filtrations, $(w, w) \in R$ implies $([w], [w]) \in R_\alpha$. Since f is surjective, it then follows that R_α is reflexive as well. (Notice that this argument actually establishes that *any* filtration from a reflexive model M to a model M' preserves reflexivity.) $\Box$

Lemma 2.32 *The axiom system S4 has the finite model property.*

Proof: Recall that *S4* is sound and complete for the class of reflexive and transitive frames. Let α be a formula satisfiable at a world w_α in a model $M = ((W, R), V)$ where (W, R) is reflexive and transitive. We have to exhibit a finite model for α based on a reflexive and transitive frame.

Let $X_\alpha = \mathrm{sf}(\alpha)$ and define W_α and V_α in terms of $\simeq_\alpha$ as usual. Let R_α be defined as follows:

$$R_\alpha = \{[w], [w'] \mid \text{For each formula } \Box\beta \in X_\alpha. \text{ if } M, w \models \Box\beta \text{ then } M, w' \models \Box\beta.\}$$

Let $M_\alpha = ((W_\alpha, R_\alpha), V_\alpha)$.

As usual, we define $f : W \to W_\alpha$ by $w \mapsto [w]$. We have already seen that such a function satisfies properties (ii) and (iii) in the definition of a filtration.

We have to verify that f satisfies properties (i) and (iv) with the new definition of R_α. To show property (i), we have to verify that if $(w, w') \in R$ then $([w], [w']) \in R_\alpha$. Suppose that $M, w \models \Box\beta$. Since (W, R) is transitive, $M, w \models \Box\beta \supset \Box\Box\beta$, so $M, w \models \Box\Box\beta$ as well. Since (w, w'), $M, w' \models \Box\beta$. Thus $([w], [w']) \in R_\alpha$.

For property (iv), we have to show that if $([w], [w']) \in R_\alpha$ then for each formula of the form $\Box\beta$ in X_α, if $M, w \models \Box\beta$, then $M, w' \models \beta$. From the definition of R_α, we know that if $M, w \models \Box\beta$, then $M, w' \models \Box\beta$ as well. Since (W, R) is reflexive, $M, w' \models \Box\beta \supset \beta$, so $M, w' \models \beta$ as required.

Having established that f is an X_α-filtration from M to M_α, it remains to prove that (W_α, R_α) is a reflexive, transitive frame. Recall that (W, R) is assumed to be a reflexive and transitive frame. We have already remarked in the proof of the previous lemma that any filtration from a reflexive model preserves reflexivity, so it is immediate that (W_α, R_α) is a reflexive frame.

To show transitivity, suppose that $([w_1],[w_2])$ and $([w_2],[w_3])$ belong to R_α. We have to show that $([w_1],[w_3]) \in R_\alpha$ as well. This means that for each formula $\Box\beta$ in X_α, we have to show that if $M, w_1 \models \Box\beta$ then $M, w_3 \models \Box\beta$. Suppose that $M, w_1 \models \Box\beta$. Since $([w_1],[w_2]) \in R_\alpha$, we know that $M, w_2 \models \Box\beta$. Now, since $([w_2],[w_3]) \in R_\alpha$, it follows that $M, w_3 \models \Box\beta$ as well. $\Box$

Exercise 2.33

(i) Recall that the axiom system B is sound and complete for the class of symmetric frames. Show that B has the finite model property. Define R_α as follows:

$$R_\alpha = \{[w],[w'] \mid \text{For each formula } \Box\beta \in X_\alpha, \begin{array}{ll} \text{(i)} & \text{if } M, w \models \Box\beta \text{ then } M, w' \models \beta \\ \text{(ii)} & \text{if } M, w' \models \Box\beta \text{ then } M, w \models \beta\} \end{array}$$

(ii) Recall that the axiom system *S5* is sound and complete for the class of frames based on equivalence relations. Show that *S5* has the finite model property. Define R_α as follows:

$$R_\alpha = \{([w],[w'] \mid \text{For each formula } \Box\beta \in X_\alpha,\ M, w \models \Box\beta \text{ iff } M, w' \models \Box\beta\}$$

$\Box$

Small model property In all the finite models we have constructed, we have defined W_α to be the set of equivalence classes generated by the relation $\simeq_\alpha$. Since the size of $\mathrm{sf}(\alpha)$ is bounded by $|\alpha|$, it follows that $|W_\alpha|$ is bounded by $2^{|\alpha|}$. Thus, when we establish the finite model property using the equivalence relation $\simeq_\alpha$, we in fact derive a bound on the size of a finite model for α. As a result, we establish a stronger property, which we call the *small model property*.

More formally, we say that a class of frames $\mathcal{C}$ has the small model property if there is a function $f_\mathcal{C} : \mathbb{N} \to \mathbb{N}$ such that for each formula α satisfiable over the class $\mathcal{C}$, there is a model for α over $\mathcal{C}$ whose size is bounded by $f_\mathcal{C}(|\alpha|)$. For instance, in the examples we have seen, $f_\mathcal{C}(|\alpha|) = 2^{|\alpha|}$.

The small model property gives us a more direct decidability argument—to check if α is satisfiable, we just have to enumerate all models of size less than $f_\mathcal{C}(|\alpha|)$. To show that this is possible, we first observe that the number of frames in this subclass is bounded. To bound the number of models based on this finite set of frames, notice that it suffices to consider valuations restricted to the finite set of atomic propositions which occur in α. Thus given a finite frame, there are only finitely many different valuations possible over that frame.

This decision procedure has the advantage of giving us a bound on the complexity of the decision problem. This bound is just the bound on the number of different models which can be generated whose size is less than $f_{\mathcal{C}}(|\alpha|)$.

Exercise 2.34 In the examples we have seen (axiom systems K, T etc.) verify that the satisfiability of a formula α can be checked in time which is doubly exponential in $|\alpha|$. $\Box$

2.7 Labelled transition systems and multi-modal logic

Transition systems A *transition system* is a pair $(S, \rightarrow)$ where S is a set of *states* and $\rightarrow \subseteq S \times S$ is a *transition relation.* Transition systems are a general framework to describe computing systems. States describe configurations of the system—for instance, the contents of the disk, memory and registers of a computer at a particular instant. The transition relation then describes when one configuration can follow another—for instance the effect of executing a machine instruction which affects some of the memory, register or disk locations and leaves the rest of the configuration untouched.

It is clear that a transition system has exactly the same structure as a frame (W, R) in modal logic. Hence, we can use modal logic to describe properties of transition systems. This is one of the main reasons why modal logic is interesting to computer scientists.

Often, we are interested in a more structured representation of the configuration space of a computing system—in particular, we not only want to record that a transition is possible from a configuration s to a configuration s' but we also want to keep track of the "instruction" which caused this change of configuration. This leads us to the notion of labelled transition systems.

Labelled transition systems A *labelled transition system* is a triple $(S, \Sigma, \rightarrow)$ where S is a set of states, Σ is a set of actions and $\rightarrow \subseteq S \times \Sigma \times S$ is a *labelled transition relation.*

The underlying structure in a finite automaton is a familiar example of a labelled transition system, where the set of states is finite.

How can we reason about labelled transition systems in the framework of modal logic? One option is to ignore the labels and consider the derived transition relation $\Rightarrow = \{(s, s') \mid \exists a \in \Sigma : (s, a, s') \in \rightarrow\}$. We can then reason about the frame $(S, \Rightarrow)$ using the modalities $\Box$ and $\Diamond$. This approach is clearly not satisfactory because we have lost all information about the labels of actions within our logic. A more faithful translation involves the use of multi-modal logics.

Multi-modal logics A *multi-relational frame* is a structure $(W, R_1, R_2, \ldots, R_n)$ where each R_i, $i \in \{1, 2, \ldots, n\}$ is a binary relation on W. A multi-relational frame can be viewed as the superposition of n normal frames (W, R_1), (W, R_2), $\ldots$, (W, R_n), all defined with respect to the same set of worlds.

To reason about a multi-relational frame, we define a *multi-modal logic* whose syntax consists of a set $\mathcal{P}$ of atomic propositions, the boolean connectives $\neg$ and $\vee$ and a set of n modalities $\Box_1, \Box_2, \ldots, \Box_n$.

To define the semantics of multi-modal logic, we first fix a valuation $V : W \to 2^{\mathcal{P}}$ as before. We then define the satisfaction relation $M, w \models \alpha$. The propositional cases are the same as for standard modal logic. The only difference is in the semantics of the modalities. For each $i \in \{1, 2, \ldots, n\}$, we define

$$M, w \models \Box_i \alpha \text{ iff for each } w' \in W, \text{ if } w\ R_i\ w' \text{ then } M, w' \models \alpha$$

Thus, the modalities $\{\Box_i\}_{i \in \{1,2,\ldots,n\}}$ are used to "independently" reason about the relations $\{R_i\}_{i \in \{1,2,\ldots,n\}}$. We can then use the theory we have developed to describe properties of each of these relations. For instance, the multi-relational frames where the axioms $\Box_3 \alpha \supset \alpha$ and $\Box_7 \alpha \supset \Box_7 \Box_7 \alpha$ are valid correspond to the class where R_3 is reflexive and R_7 is transitive. We can express interdependencies between different relations using formulas which combine these modalities. For instance, the formula $\alpha \supset \Diamond_5 \Diamond_2 \beta$ indicates that a world which satisfies α has an R_5-neighbour which in turn has an R_2-neighbour where β holds.

We have see how to characterise classes of frames using formulas from modal logic. We can extend this idea in a natural way to characterise classes of multi-relational frames.

Exercise 2.35 Consider the class of multi-relational frames (W, R_1, R_2) where $R_2 = R_1^{-1}$. Describe axioms to characterise this class. (*Hint:* The combined relation $R_1 \cup R_2$ is a symmetric relation on W. Work with suitable modifications of axiom (B). You may use more than one axiom.) $\Box$

To reason about labelled transition systems in this framework, we have to massage the structure $(S, \Sigma, \to)$ into a multi-relational frame. To achieve this, we define a relation $\to_a \ \subseteq S \times S$ for each $a \in \Sigma$ as follows:

$$\to_a \,=\, \{(s, s') \mid (s, a, s') \in \,\to\}$$

It is then clear that the multi-relational frame $(S, \{\to_a\}_{a \in \Sigma})$ describes the same structure as the original labelled transition system $(S, \Sigma, \to)$.

To reason about the structure $(S, \{\to_a\}_{a \in \Sigma})$, we have modalities $\Box_a$ (read as *Box a*) and $\Diamond_a$ (read as *Diamond a*) for each $a \in \Sigma$. Traditionally, the modality $\Box_a$ is written $[a]$ and the modality $\Diamond_a$ is written $\langle a \rangle$.

When reasoning about labelled transition systems, the set of atomic propositions $\mathcal{P}$ corresponds to properties which distinguish one configuration of the system from each other. For instance, we could have an atomic proposition to denote that "memory location 27 is unused" or that "the printer is busy". In these notes, we will not go into the details of how to model a computing system in terms of such a logic.

Assuming we have an abstract encoding of system properties in terms of atomic propositions, we can now reason about the dynamic behavior of the system. For instance, we can assert $M, s \models [c]\langle b \rangle \alpha$ to denote that in the state s, any c-transition will lead to a state from where we can use a b-transition to realise the property described by α. In particular, if α is just the constant $\top$, this formula asserts that a b-transition is enabled after any c-transition.

Unfortunately, we still do not have the expressive power we need to make non-trivial statements about programs. For instance, we cannot say that after a c-transition, we can *eventually* reach a state where a b-transition is enabled. Or that we have reached a portion of the state space where *henceforth* only a and d transitions are possible.

For this, we need to move from modal logic to dynamic logic, which is the topic of discussion in the next section.