Advanced Modal Logic VIII

Yanjing Wang

Department of Philosophy, Peking University

Mar. 14th, 2024

Advanced Modal Logic (2024 Spring)

Filtration

Ultrafilter extension

Filtration

Finite model property via filtration (a *family* of constructions)

Definition (A filtration of a model through Σ)

Let Σ be a *subformula-closed* set of modal formulas. Let $\equiv_\Sigma$ be the logical equivalence w.r.t. the formulas in Σ. A filtration of a given Kripke model $\mathcal{M}$ of a unary similarity type is a tuple $(W_{\equiv_\Sigma}, \{\xrightarrow{a}_f\}_{a \in \mathbf{O}}, V_f)$ satisfying:

1. If $w \xrightarrow{a} v$ then $[w] \xrightarrow{a}_f [v]$.
2. If $[w] \xrightarrow{a}_f [v]$ then for all $\Box_a \varphi \in \Sigma$, $\mathcal{M}, w \vDash \Box_a \varphi$ implies $\mathcal{M}, v \vDash \varphi$.
3. $V_f([w]) = V(w)|_{\mathbf{P}'}$ where $\mathbf{P}' \subseteq \mathbf{P}$ is the set of proposition letters in Σ.

Compare the above definition with the $\Diamond$-based definition in the textbook. Note that if $\Diamond$ is not the *primitive* modality in our language, it can happen that $\Box\varphi \in \Sigma$ but $\Diamond\varphi$ $(\neg\Box\neg\varphi)$ is not.

Finite model property via filtration

Proposition

If $\mathcal{M}^f$ is a filtration of $\mathcal{M}$ through a finite subformula-closed set Σ then $\mathcal{M}^f$ has at most 2^n states where $n = |\Sigma|$.

Theorem (filtration theorem)

If $\mathcal{M}^f$ is a filtration of $\mathcal{M}$ through a finite subformula-closed set Σ then for all $\varphi \in \Sigma$ and w in $\mathcal{M}$ we have
$\mathcal{M}, w \vDash \varphi \iff \mathcal{M}^f, [w] \vDash \varphi$

Finite model property via filtration

By induction on φ.

The case of basic proposition is trivial. IH: for all the subformula ψ of φ: $\mathcal{M}, w \vDash \psi \iff \mathcal{M}^f, [w] \vDash \psi$
Boolean cases are trivial. Now consider $\Box_a \psi$: Suppose $\mathcal{M}, w \nvDash \Box_a \psi$ then there is a v such that $w \xrightarrow{a} v$ and $\mathcal{M}, v \nvDash \psi$.By condition (1) of $\xrightarrow{a}_f$, IH, and the fact that Σ is subformula-closed, $\mathcal{M}^f, [w] \nvDash \Box_a \psi$. Suppose $\mathcal{M}^f, [w] \nvDash \Box_a \psi$ then there is a $[v]$ such that $[w] \xrightarrow{a}_f [v]$ and $\mathcal{M}^f, [v] \nvDash \psi$. By IH, $\mathcal{M}, v \nvDash \psi$. Therefore by condition (2) of $\xrightarrow{a}_f$, $\mathcal{M}, w \nvDash \Box_a \psi$. $\Box$

Q: Do you see why we define the second condition on $\rightarrow_f$?

Q: Can we relax the first condition on $\rightarrow_f$?

Finite model property via filtration

Filtrations do exist:

- $\xrightarrow{a}_1 = \{([w],[v]) \mid \text{there exist } w' \in [w], v' \in [v] : w' \xrightarrow{a} v'\}$
- $\xrightarrow{a}_2 = \{([w],[v]) \mid \text{for all } \Box_a\varphi \in \Sigma\ \mathcal{M}, w \vDash \Box_a\varphi \implies \mathcal{M}, v \vDash \varphi.\}$

The second relation is well-defined: the exact choices of w, v in their equivalence classes do not matter. Any $\xrightarrow{a}_f : \xrightarrow{a}_1 \subseteq \xrightarrow{a}_f \subseteq \xrightarrow{a}_2$ (flexible choices). Q: Are filtrations always (restricted) bisimilar to the original model? No!

Theorem (ML has the *strong* finite model property)

If $\varphi \in \mathbf{ML}$ is satisfiable then it is satisfiable in a finite model with the size bounded by 2^n where n is the length of φ.

How many subformulas are there for a given φ? Less or equal than the length of the formula. Given a special class of models, you need to revise the def. of filtration.

Testing Validity and Satisfiability

Entscheidungsproblem (Hilbert 1928, dates back to Leibniz): find an algorithm such that given a mathematical statement expressed by *first-order logic* as an input, it can output true or false correctly. Church and Turing: it is impossible!

The satisfiability problem of a semantically given logic $(L, \mathbb{C}, \models)$ is the problem of testing whether L formulas are satisfiable in some model in $\mathbb{C}$.

`ML` over the class of all the models has the *strong* finite model property which implies that it is decidable: just try these finite models (with finitely many propositions and labels of relations, modulo isomorphism) one by one (based on the model checking problem is decidable).

Is it possible that satisfiability is decidable but model checking is *undecidable*? Yes!

What about a logic that has the (strong) finite model property in general? It still can be undecidable e.g., first-order logic on finite models, Ex 6.2.4 and Ex 6.2.5

- Strong finite model property w.r.t. recursive class of models implies decidability (you can check whether a model is in $\mathbb{C}$, assuming model checking is decidable).
- There are decidable logics which do not have finite model property.

Definition (Another definition of finite model property)

A logic L (a set of formulas) has f.m.p. iff every non-theorem of L has a finite counter L-model iff L is characterized by a class of finite models.

The above finite model property and finite axiomatization implies decidability of the logic.

Ultrafilter extension

The last yet important model construction method

Not all the models are m-saturated. How to turn a model into an m-saturated one?

We need to add some successors such that every finitely satisfiable set of formulas is satisfiable in one of the successors.

Just adding worlds literally does not work. How to do it?

- A world "is" a (maximal consistent) set of formulas
- A formula "is" a set of worlds where it is true
- Then a world "is" a set of sets of worlds
- Enough worlds (应有尽有)
- Enough accessibility relations (应连尽连)

Ultrafilter

Given a set W, a *filter* F over W is a subset of $\mathcal{P}(W)$ s.t.:

- $W \in F$
- $X, Y \in F$ implies $X \cap Y \in F$
- $X \in F$ and $X \subseteq Y$ implies $Y \in F$

A proper filter is a filter such that $\emptyset \notin F$. An *ultrafilter* is a proper filter such that either $X \in F$ or $W \backslash X \in F$.

Example

Given the set of natural numbers $\mathbb{N}$, the set

$$\{X \mid X \text{ is a co-finite subset of } \mathbb{N}\} = \{X \mid \mathbb{N} \backslash X \text{ is finite }\}$$

is a proper filter.

$\mathbb{N}$ can be replaced by any infinite set.

Ultrafilter

Another intuition: a subset of W can be viewed as (the extension) of a formula which holds exactly on the states in this subset. From this point of view, a filter is a set of formulas which is closed under $\wedge$ and $\rightarrow$. A proper filter is a consistent set and an ultrafilter is a maximal consistent set of formulas.

A *principal ultrafilter* π_w is an ultrafilter generated by a singleton set $\{w\} : \pi_w = \{X \mid w \in X \subseteq W\}$ (check that it is indeed an ultrafilter).
Then it is not hard to see that a non-principal ultrafilter (if exists) contains only infinite subsets and all the co-finite subsets of W.
It also means that there is no non-principal ultrafilter over a finite W.

Ultrafilter

Theorem (Ultrafilter Theorem)

Any proper filter can be extended into an ultrafilter.

Proof.

By Zorn lemma (a version of the axiom of choice). □

Corollary

Any non-empty set E of $\mathcal{P}(W)$ can be extended into an ultrafilter iff E has the finite intersection property (any finite intersection of elements in E is non-empty).

Therefore to construct an ultrafilter from a non-empty set $E \subseteq \mathcal{P}(W)$, we just need to verify whether E has finite intersection property. To build a non-principal ultrafilter over an infinite set W, we can start from the proper filter of all the co-finite subsets of W, and apply the ultrafilter theorem.

Definition (Ultrafilter extension)

Given a model $\mathcal{M} = \langle W, \to, V \rangle$, its ultrafilter extension $\mathcal{M}^{\mathfrak{ue}} = \langle W^{\mathfrak{ue}}, \to^{\mathfrak{ue}}, V^{\mathfrak{ue}} \rangle$ where:

- $W^{\mathfrak{ue}} = \{u \mid u \text{ is an ultrafilter over } W\}$
- $u \to^{\mathfrak{ue}} u' \iff (\forall X : X \in u' \implies m_R(X) \in u)$
- $V^{\mathfrak{ue}}(u) = \{p \mid \{w \mid p \in V(w)\} \in u\}$

where $m_R(X) = \{w \mid \exists v \in X \text{ such that } wRv\}$

Proposition (alternative definition of $\to^{\mathfrak{ue}}$ which is more useful)

$u \to^{\mathfrak{ue}} u' \iff (\forall Y : l_R(Y) \in u \implies Y \in u')$

The intuition behind two definitions of $u \to^{\mathfrak{ue}} u'$:
$(\forall \varphi : \varphi \in u' \implies \Diamond\varphi \in u)$ and $(\forall \varphi : \Box\varphi \in u \implies \varphi \in u')$.

Ultrafilter extension

It is not hard to show that: $\pi_w \to^{\mathfrak{ue}} \pi_v \iff w \to v$ (left to right: take $\{v\} \in \pi_v$).

Therefore the submodel of $\mathcal{M}^{\mathfrak{ue}}$ obtained by restricting to the principal ultrafilters is an isomorphic copy of $\mathcal{M}$. The extra worlds in $\mathcal{M}^{\mathfrak{ue}}$ are non-principal ultrafilters. By the ultrafilter theorem and its corollary, such non-principal ultrafilters exist if W is infinite. This justifies the name: ultrafilter **extension**.

What is the ulrafilter extension of $(\mathbb{N}, <)$?

Ultrafilter extension

Given $\mathcal{M}$, we abuse the notation of $V_{\mathcal{M}}$ and let $V_{\mathcal{M}}(\varphi)$ be the set of worlds in $\mathcal{M}$ where φ is true.

Theorem

Given a pointed model $\mathcal{M}, w$, $\mathcal{M}, w \equiv_{ML} \mathcal{M}^{\mathfrak{ue}}, \pi_w$.

It is a bit hard to prove this theorem directly since the induction hypothesis would be only about principal ultrafilters in $\mathcal{M}^{\mathfrak{ue}}$, but clearly a principal ultrafilter π_w may have a successor which is a non-principal ultrafilter given that W is infinite. Thus we prove the following more general result first:

Theorem

Given a pointed model $\mathcal{M}, w$, $\mathcal{M}^{\mathfrak{ue}}, u \vDash \varphi \iff V_{\mathcal{M}}(\varphi) \in u$.

by induction on φ.

To handle the Boolean cases, you need to use the properties of ultrafilters. They are *not* that trivial.

The case of $\Box\psi$: if $\mathcal{M}^{\mathfrak{ue}}, u \not\models \Box\psi$ then there is a u' such that $u \to^{\mathfrak{ue}} u'$ and $u' \not\models \psi$. By IH, $V(\psi) \notin u'$. By the definition of $\to^{\mathfrak{ue}}$, for all X $l_R(X) \in u$ implies $X \in u'$. Then by IH $l_R(V(\psi)) \notin u$, namely $V(\Box\psi) \notin u$.

Now suppose $V(\Box\psi) \notin u$, we need to show $\mathcal{M}^{\mathfrak{ue}}, u \not\models \Box\psi$. The proof strategy is that we **construct** a successor u' of u such that $V(\psi) \notin u'$. Let $u'' = \{V(\neg\psi)\} \cup \{Y \mid l_R(Y) \in u\}$, we just need to show that it has finite intersection property. First note that $V(\neg\Box\psi) \cap l_R(Y) \neq \emptyset$ (since u is an ultrafilter and $V(\Box\psi) \notin u$ but $l_R(Y) \in u$), thus $V(\neg\psi) \cap Y$ is not empty. Moreover, based on the fact that $\{Y \mid l_R(Y) \in u\}$ is closed under intersection we can prove the theorem. □

Ultrafilter extension

Theorem

Given a pointed model $\mathcal{M}$, $\mathcal{M}^{\mathfrak{ue}}$ is m-saturated.

Proof.

The idea: given a world u in $\mathcal{M}^{\mathfrak{ue}}$ and a set of formulas which is finitely satisfiable in the set of successors of u, we construct a u' such that $u \rightarrow^{\mathfrak{ue}} u'$ and for all $\varphi \in \Sigma : V(\varphi) \in u'$. Let $u'' = \{Y \mid l_R(Y) \in u\} \cup \{V(\varphi) \mid \varphi = \psi_0 \wedge \cdots \wedge \psi_n, n \geq 0, \psi_k \in \Sigma\}$, we need to show that u'' has the finite intersection property. Since such a φ is satisfiable at some successor of u thus $V(\Diamond\varphi) \in u$. Similar as before we can show that $V(\varphi) \cap Y \neq \emptyset$.

□

Bisimilarity somewhere else

Based on the previous results we can show:

Theorem

For any pointed models $\mathcal{M}, w$ and $\mathcal{N}, v$:

$\mathcal{M}, w \equiv_{ML} \mathcal{N}, v \implies \mathcal{M}^{\mathfrak{ue}}, \pi_w \leftrightarrows \mathcal{N}^{\mathfrak{ue}}, \pi_v$.

Proof.

By a detour:

$$\begin{array}{ccc} \mathcal{M}, w & \equiv_{\mathtt{ML}} & \mathcal{N}, v \\ \big\downarrow {\scriptstyle \mathtt{ML}\equiv} & & \big\downarrow {\scriptstyle \mathtt{ML}\equiv} \\ \mathcal{M}^{\mathfrak{ue}}, \pi_w & \equiv_{\mathtt{ML}} = \leftrightarrows & \mathcal{N}^{\mathfrak{ue}}, \pi_v \end{array}$$

m-saturated m-saturated

□