



A Logic of "Black Box" Classifier Systems

Xinghan Liu, Emiliano Lorini

► To cite this version:

Xinghan Liu, Emiliano Lorini. A Logic of "Black Box" Classifier Systems. 28th Workshop on Logic, Language, Information and Computation (WOLLIC 2022), Sep 2022, Lasi, Romania. pp.158-174, <10.1007/978-3-031-15298-6_10>. <hal-03873355>

HAL Id: hal-03873355

<https://hal.science/hal-03873355v1>

Submitted on 26 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

A Logic of “Black Box” Classifier Systems

Xinghan Liu¹ and Emiliano Lorini²

¹ IRIT, University of Toulouse, France

² IRIT, CNRS, University of Toulouse, France

Abstract. Binary classifiers are traditionally studied by propositional logic (PL). PL can only represent them as white boxes, under the assumption that the underlying Boolean function is fully known. Binary classifiers used in practical applications and trained by machine learning are however opaque. They are usually described as black boxes. In this paper, we provide a product modal logic called PLC (Product modal Logic for binary input Classifier) in which the notion of “black box” is interpreted as the uncertainty over a set of classifiers. We give results about axiomatics and complexity of satisfiability checking for our logic. Moreover, we present a dynamic extension in which the process of acquiring new information about the actual classifier can be represented.

1 Introduction

The notions of explanation and explainability have been extensively investigated by philosophers [11, 14] and are key aspects of AI-based systems given the importance of explaining the behavior and prediction of an artificial intelligent system. A variety of notions of explanation for classifier systems have been discussed in the area of explainable AI (XAI). Since systems trained by machine learning are increasingly opaque, instead of studying specific models, the *model-agnostic* approach comes into focus. Namely, given a black box system or algorithm, we know nothing about how it works inside. Without opening the black box, we can query some (but not all) inputs and have some partial information about the system. Initially there were *global* model-agnostic explanations like partial dependence plots and global surrogate models. Recently LIME [21] and its followers e.g. SHAP [16] and Anchors [22] have raised a *local* model-agnostic explanation approach, namely explaining why a given input is classified in a certain way. For a comprehensive overview of the research in this area see, e.g., [18].

At the mathematical level, a binary classifier can be viewed as a Boolean function and is traditionally studied by propositional logic. Recent years have witnessed several logic-based approaches to local explanation of classifier systems [23, 6, 13, 12, 2], e.g., computing prime implicants and abductive explanations of a given classification, and detecting biases in the classification process by means of the notion counterfactual explanation. But, all these logic-based approaches deal with “white box” classifiers, i.e., specific transparent models representable by propositional formulas. A limitation is that given a Boolean function f and a propositional formula φ , φ either fully expresses f or does not express f at all.

This all-or-nothing nature makes it impossible to give a *partial* description of f , which is a natural way to represent a black box classifier.

The central idea of this paper is that a product modal logic is the proper way to represent a “black box” classifier. As we have shown in [15], it is natural to think of a classifier with binary inputs as a partition of an $S5$ Kripke model, where each possible state stands for an input instance. However, this only represents “white box” classifiers. We extend this semantics with a second dimension universally ranging over a set of possible classifiers, which results in a proper extension of the product modal logic $S5 \times S5 = S5^2$ [8] we call PLC (Product modal Logic of binary-input Classifiers). The notion of black box is interpreted as an agent’s uncertainty among those (white box) classifiers, as illustrated in Figure 1.




Fig. 1. A classifier associating color labels in $\{\text{red, yellow, green}\}$ to input instances. We do not know its Boolean formula, since f_0, f_1, f_2, f_3 are all compatible with our partial knowledge of it. However, we know that the two input instances s_0 and s_3 are both classified as green.

The paper is structured as follows. Section 2 introduces the modal language and semantic model of PLC which we name multi-classifier model (MCM). Its axiomatics along with the completeness and complexity results for the satisfiability checking problem are given in Section 3. In Section 4, we will exemplify the logic’s application by using it to represent the notion of black box and to formalize different notions of classifier explanation. A dynamic extension is given in Section 6 to capture the process of acquisition of new knowledge about the classifier. Some non-routine proofs are given in a technical annex at the end of the paper.

2 Language and Semantics

Let $Atm_0 = \{p, q, \dots\}$ be a countable set of atomic propositions which intend to denote input variables (features) of a classifier. We introduce a finite set Val to denote the possible output values (classifications, decisions) of the classifier.

Elements of Val are noted x, y, \dots . For any $x \in Val$, we call $\mathbf{t}(x)$ a decision atom, and have $Dec = \{\mathbf{t}(x) : x \in Val\}$.³ Finally let $Atm = Atm_0 \cup Dec$.

The modal language \mathcal{L} is defined by the following grammar:

$$\varphi ::= p \mid \mathbf{t}(x) \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_I \varphi \mid \Box_F \varphi,$$

where p ranges over Atm_0 and x ranges over Val .

Definition 1. A multi-classifier model (MCM) is a pair $\Gamma = (S, \Phi)$ where $S \subseteq 2^{Atm_0}$ and $\Phi \subseteq F_S$, with $F_S = Val^S$ the set of functions with domain S and codomain Val . A pointed MCM is a triple (Γ, s, f) where $\Gamma = (S, \Phi)$ is an MCM, $s \in S$ and $f \in \Phi$. The class of all multi-classifier models is noted **MCM**.

Formulas in \mathcal{L} are interpreted relative to a pointed MCM as follows.

Definition 2 (Satisfaction relation). Let $\Gamma = (S, \Phi)$ be an MCM, $s \in S$ and $f \in \Phi$. Then,

$$\begin{aligned} (\Gamma, s, f) \models p &\iff p \in s, \\ (\Gamma, s, f) \models \mathbf{t}(x) &\iff f(s) = x, \\ (\Gamma, s, f) \models \neg\varphi &\iff (\Gamma, s, f) \not\models \varphi, \\ (\Gamma, s, f) \models \varphi \wedge \psi &\iff (\Gamma, s, f) \models \varphi \text{ and } (\Gamma, s, f) \models \psi, \\ (\Gamma, s, f) \models \Box_I \varphi &\iff \forall s' \in S : (\Gamma, s', f) \models \varphi, \\ (\Gamma, s, f) \models \Box_F \varphi &\iff \forall f' \in \Phi : (\Gamma, s, f') \models \varphi. \end{aligned}$$

Both $\Box_I \varphi$ and $\Box_F \varphi$ have standard modal reading but range over different sets. $\Box_I \varphi$ has to be read “ φ necessarily holds for the actual function, regardless of the input instance”, while its dual $\Diamond_I \varphi =_{def} \neg \Box_I \neg \varphi$ has to be read “ φ possibly holds for the actual function, regardless of the input instance”. Similarly, $\Box_F \varphi$ has to be read “ φ necessarily holds for the actual input instance, regardless of the function” and its dual $\Diamond_F \varphi$ has to be read “ φ possibly holds for the actual input instance, regardless of the function”.

Let X be a finite subset of Atm_0 . An important abbreviation is the following:

$$[X]\varphi =_{def} \bigwedge_{Y \subseteq X} ((\bigwedge_{p \in Y} \bigwedge_{p \in X \setminus Y} \neg p) \rightarrow \Box_I ((\bigwedge_{p \in Y} \bigwedge_{p \in X \setminus Y} \neg p) \rightarrow \varphi)).$$

Complex as it seems, $[X]\varphi$ means nothing but “ φ necessarily holds, regardless of the values of the input variables outside X ” or “ φ necessarily holds, if the values of the input variables in X are kept fixed”. It can be justified by checking that $(\Gamma, s, f) \models [X]\varphi$, if and only if $\forall s' \in S$, if $s \cap X = s' \cap X$ then $(\Gamma, s', f) \models \varphi$. Its dual $\langle X \rangle \varphi =_{def} \neg [X] \neg \varphi$ has to be read “ φ possibly holds, if the values of the input variables in X are kept fixed”. These modalities have a *ceteris paribus* reading and were first introduced in [9]. Similar modalities are used in existing logics of functional dependence [25, 3].

³ Notice x is an output value rather than the output variable, which makes sense of the symbolic difference between p and $\mathbf{t}(x)$.

A formula φ of \mathcal{L} is said to be satisfiable relative to the class **MCM** if there exists a pointed multi-classifier model (Γ, s, f) with $\Gamma \in \mathbf{MCM}$ such that $(\Gamma, s, f) \models \varphi$. We say that φ is valid in the multi-classifier model $\Gamma = (S, \Phi)$, noted $\Gamma \models \varphi$, if $(\Gamma, s, f) \models \varphi$ for every $s \in S, f \in \Phi$. It is said to be valid relative to **MCM**, noted $\models_{\mathbf{MCM}} \varphi$, if $\neg\varphi$ is not satisfiable relative to **MCM**.

3 Axiomatics and Complexity

In this section, we are going to present two axiomatics for the language \mathcal{L} by distinguishing the finite-variable from the infinite-variable case. We will moreover give complexity results for satisfiability checking. Before, we are going to introduce an alternative Kripke semantics for the interpretation of the language \mathcal{L} . It will allow us to use the standard canonical model technique for proving completeness. Indeed, this technique cannot be directly applied to MCMs.

3.1 Alternative Kripke Semantics

The crucial concept of the alternative semantics is multi-decision model (MDM).

Definition 3. *An MDM is a tuple $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ where:*

- W is a set of worlds,
- \sim_{\square_I} and \sim_{\square_F} are equivalence relations on W ,
- $V : W \longrightarrow 2^{Atm}$ is a valuation function,

and which satisfies the following constraints, $\forall w, v \in W, \forall x, y \in Val$:

- (C1) $\sim_{\square_I} \circ \sim_{\square_F} = \sim_{\square_F} \circ \sim_{\square_I}$,
- (C2) if $V_{Atm_0}(w) = V_{Atm_0}(v)$ and $w \sim_{\square_I} v$ then $V_{Dec}(w) = V_{Dec}(v)$,
- (C3) if $w \sim_{\square_F} v$ then $V_{Atm_0}(w) = V_{Atm_0}(v)$,
- (C4) if $\mathbf{t}(x) \in V(w)$ and $x \neq y$ then $\mathbf{t}(y) \notin V(w)$,
- (C5) $\exists x \in Val$ such that $\mathbf{t}(x) \in V(w)$,

with $V_Y(w) = (V(w) \cap Y)$ for every $w \in W$ and for every $Y \subseteq Atm$, and \circ the standard composition operator for binary relations.

The class of multi-decision models is noted **MDM**. An MDM $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ is called finite if W is finite. The class of finite MDMs is noted finite-**MDM**. Interpretation of formulas in \mathcal{L} relative to a pointed MDM goes as follows. (We omit interpretations for \neg and \wedge which are defined as usual.)

Definition 4 (Satisfaction Relation). *Let $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ be an MDM and let $w \in W$. Then,*

$$\begin{aligned}
 (M, w) \models q &\iff q \in V(w) \text{ for } q \in Atm, \\
 (M, w) \models \square_I \varphi &\iff \forall v \in W, \text{ if } w \sim_{\square_I} v \text{ then } v \models \varphi, \\
 (M, w) \models \square_F \varphi &\iff \forall v \in W, \text{ if } w \sim_{\square_F} v \text{ then } v \models \varphi.
 \end{aligned}$$

Validity and satisfiability of formulas in \mathcal{L} relative to class **MDM** (resp. finite-**MDM**) are defined in the usual way.

The most important result in this subsection is the semantic equivalence between **MCM** and **MDM**, regardless of Atm_0 being finite or infinite. Although a pointed MDM (M, w) looks like a pointed MCM (Γ, s, f) , it only approximates it. Indeed, unlike an MCM, an MDM M may be redundant, that is, (i) a classifier in M (i.e., a \sim_{\square_I} -equivalence class) may include multiple copies of the same input instance (i.e., of the same valuation for the atoms in Atm_0), or (ii) M may contain multiple copies of the same classifier (i.e., two identical \sim_{\square_F} -equivalence classes). Moreover, an MDM M may be “defective” insofar as (iii) the intersection between a classifier in M (i.e., a \sim_{\square_I} -equivalence class) and the set of all possible classifications of a given input instance by the classifiers in M (i.e., a \sim_{\square_F} -equivalence class) is not a singleton. What makes the proof of the following theorem non-trivial is transforming a possibly redundant or defective MDM into a non-redundant and non-defective one by preserving truth of formulas. A non-redundant and non-defective MDM is then isomorphic to an MCM.

Theorem 1. *Let $\varphi \in \mathcal{L}$. Then, φ is satisfiable relative to the class **MCM** if and only if it is satisfiable relative to the class **MDM**.*

3.2 Finite-Variable Case

We first consider the variant of the logic with finitely many propositional atoms in Atm_0 . For every finite $X, Y \subseteq Atm_0$ we define:

$$cn_{X,Y} =_{def} \bigwedge_{p \in X} p \wedge \bigwedge_{p \in (Y \setminus X)} \neg p.$$

Definition 5 (Logic PLC). *Let Atm_0 be finite. We define PLC as the extension of classical propositional logic given by axioms and rules of inference in Table 1.*

Axioms **AtLeast**_{t(x)}, **AtMost**_{t(x)} and **Funct** guarantee that every input $Y \subseteq Atm_0$, whose syntactic counterpart is cn_{Y, Atm_0} , has only one decision atom as output. Axioms **K**_■, **T**_■, **4**_■ and **5**_■ together with the rule of inference **Nec**_■ indicate that both modal operators \square_F and \square_I satisfy the principles of the modal logic S5. According to Axiom **Comm**, they moreover commute. This makes the logic meet the requirement of a product of two S5 modal logics, i.e., S5² [8]. Nevertheless, the existence of the two “independence” Axioms **Indep** _{\square_F, p} and **Indep** _{$\square_F, \neg p$} indicates that PLC is stronger than S5² in general.

Soundness of PLC relative to **MCM** is a simple exercise. To prove the completeness result, we first need to show that PLC is complete relative to **MDM**, which is proven by the canonical model construction.

Theorem 2. *Let Atm_0 be finite. Then, the logic PLC is sound and complete relative to the class **MDM**.*

Our main result of this subsection becomes a corollary of Theorems 1 and 2.

Corollary 1. *Let Atm_0 be finite. Then, the logic PLC is sound and complete relative to the class **MCM**.*

$(\blacksquare\varphi \wedge \blacksquare(\varphi \rightarrow \psi)) \rightarrow \blacksquare\psi$	(K $_{\blacksquare}$)
$\blacksquare\varphi \rightarrow \varphi$	(T $_{\blacksquare}$)
$\blacksquare\varphi \rightarrow \blacksquare\blacksquare\varphi$	(4 $_{\blacksquare}$)
$\neg\blacksquare\varphi \rightarrow \blacksquare\neg\blacksquare\varphi$	(5 $_{\blacksquare}$)
$\Box_F\Box_I\varphi \leftrightarrow \Box_I\Box_F\varphi$	(Comm)
$\bigvee_{x \in Val} \mathbf{t}(x)$	(AtLeast $_{\mathbf{t}(x)}$)
$\mathbf{t}(x) \rightarrow \neg\mathbf{t}(y) \text{ if } x \neq y$	(AtMost $_{\mathbf{t}(x)}$)
$(\mathbf{cn}_{X, Atm_0} \wedge \mathbf{t}(x)) \rightarrow \Box_I(\mathbf{cn}_{X, Atm_0} \rightarrow \mathbf{t}(x))$	(Func t)
$p \rightarrow \Box_F p$	(Indep $_{\Box_F, p}$)
$\neg p \rightarrow \Box_F \neg p$	(Indep $_{\Box_F, \neg p}$)
$\frac{\varphi}{\blacksquare\varphi}$	(Nec $_{\blacksquare}$)

Table 1. Axioms and rules of inference, with $\blacksquare \in \{\Box_I, \Box_F\}$

3.3 Infinite-Variable Case

We now move to the infinite-variable variant of our logic, under the assumption that the set Atm_0 is countably infinite. In order to obtain an axiomatics we just need to drop the “functionality” Axiom **Func**t of Table 1. Indeed, when Atm_0 is infinite, the construction \mathbf{cn}_{X, Atm_0} cannot be expressed in a finitary way.

Definition 6 (Logic WPLC). We define WPLC (Weak PLC) to be the extension of classical propositional logic given by Axioms **K** $_{\blacksquare}$, **T** $_{\blacksquare}$, **4** $_{\blacksquare}$, **5** $_{\blacksquare}$, **Comm**, **AtLeast** $_{\mathbf{t}(x)}$, **AtMost** $_{\mathbf{t}(x)}$, **Indep** $_{\Box_F, p}$ and **Indep** $_{\Box_F, \neg p}$, and the rule of inference **Nec** $_{\blacksquare}$ in Table 1.

Soundness of the logic WPLC is a straightforward exercise. For completeness, we need to distinguish MDMs from quasi-MDMs that are obtained by removing the “functionality” Constraint **C2** from Definition 3.

Definition 7 (Quasi-MDM). A quasi-MDM is a tuple $M = (W, \sim_{\Box_I}, \sim_{\Box_F}, V)$ where W , \sim_{\Box_I} , \sim_{\Box_F} and V are defined as in Definition 3 and which satisfies all constraints of Definition 3 except **C2**.

The class of quasi-MDMs is noted **QMDM**. A quasi-MDM $M = (W, \sim_{\Box_I}, \sim_{\Box_F}, V)$ is said to be finite if W is finite. The class of finite quasi-MDMs is noted finite-**QMDM**. Semantic interpretation of formulas in \mathcal{L} relative to quasi-MDMs is analogous to semantic interpretation relative to MDMs given in Definition 4. Moreover, validity and satisfiability of formulas in \mathcal{L} relative to class **QMDM** (resp. finite-**QMDM**) is again defined in the usual way.

The first crucial result of this subsection is that when Atm_0 is infinite the language \mathcal{L} cannot distinguish finite MDMs from finite quasi-MDMs.

Theorem 3. *Let $\varphi \in \mathcal{L}$ with Atm_0 infinite. Then, φ is satisfiable relative to the class **finite-MDM** if and only if it is satisfiable relative to the class **finite-QMDM**.*

The second result is that satisfiability for formulas in \mathcal{L} relative to the class **QMDM** is equivalent to satisfiability relative to the class **finite-QMDM**.

Theorem 4. *Let $\varphi \in \mathcal{L}$. Then, φ is satisfiable relative to the class **QMDM** if and only if it is satisfiable relative to the class **finite-QMDM**.*

The following theorem is provable by standard canonical model argument. Note that like Theorems 1 and 4, it does not rely on Atm_0 being infinite or finite.

Theorem 5. *The logic **WPLC** is sound and complete relative to the class **QMDM**.*

The fact that the logic **WPLC** is sound and complete relative to the class **MCM** is a direct corollary of Theorems 1, 3, 4 and 5.

Corollary 2. *Let Atm_0 be infinite. Then, the logic **WPLC** is sound and complete relative to the class **MCM**.*

3.4 Complexity Results

We now move to complexity of satisfiability checking. As for the axiomatics, we distinguish the finite-variable from the infinite-variable case. When Atm_0 is finite, the problem of verifying whether a formula is satisfiable is polynomial. The latter problem mirrors the satisfiability checking problem for the finite-variable modal logic **S5** which is also known to be polynomial [10].

Theorem 6. *Let Atm_0 be finite. Then, checking satisfiability of \mathcal{L} -formulas relative to the class **MCM** can be done in polynomial time.*

When moving from the finite-variable to the infinite-variable case complexity of satisfiability checking increases, but the problem remains in NP.

Theorem 7. *Let Atm_0 be infinite. Then, checking satisfiability of \mathcal{L} -formulas relative to the class **MCM** is NP-complete.*

4 Application

As mentioned, the \Box_F operator is interpreted as partial knowledge about the classifier properties. In this section, we are going to exemplify how to use it for representing abductive explanations of a black box classifier.

4.1 An Example of Classification Task

Consider a selection function which specifies whether a paper submitted to a conference is acceptable for presentation (1) or not (0) depending on its feature profile composed of four input features: significance (si), originality (or), clarity of the presentation (cl) and fulfillment of the anonymity requirement (an). For the sake of simplicity, we assume each feature in a paper profile is binary: si means the paper is significant while \neg si means the paper is not significant, or means the paper is original while \neg or means the paper is not original, and so on. We say that a first paper profile dominates a second paper profile, if all conditions satisfied by the second profile are satisfied by the first profile, and there exists a condition satisfied by the first profile which is not satisfied by the second profile. For example if the first profile is $si \wedge \neg or \wedge cl \wedge an$ and the second profile is $si \wedge \neg or \wedge \neg cl \wedge an$, then the first dominates the second.

The selection function is implemented in a classifier system that has to automatically split papers into two sets, the set of acceptable papers and the set of non-acceptable ones. We assume a certain agent (e.g., the author of a paper submitted to the conference) has only partial knowledge of the classifier system. In particular, she only knows that the classifier complies with the following three constraints: (1) submissions that satisfy the four conditions should be automatically accepted, (2) if a first paper profile dominates a second paper profile and the second paper profile is acceptable, then the first paper profile should also be acceptable, and (3) submissions that violate the anonymity requirement should be automatically rejected. In this case, the classifier is a black box for the agent.

Example 1. The multi-classifier model (MCM) representing the previous situation is the tuple $\Gamma = (S, \Phi)$ such that $S = 2^{\{si, or, cl, an\}}$ and

$$\begin{aligned} \forall f \in F_S, f \in \Phi \text{ iff } & (i) \forall s \in S, \text{ if } \{si, or, cl, an\} \subseteq s \text{ then } f(s) = 1, \\ & (ii) \forall s, s' \in S, \text{ if } s \subset s' \text{ and } f(s) = 1 \text{ then } f(s') = 1. \\ & (iii) \forall s \in S, \text{ if } an \notin s \text{ then } f(s) = 0. \end{aligned}$$

The agent does not know which function in Φ corresponds to the actual classifier, i.e., they are epistemically indistinguishable for her.

4.2 Explanations

Given space constraints, we exemplify explanations for white and black box classifiers by showing the dichotomy global vs. local explanation and the notion of *abductive explanation* based on *prime implicant*. Some notations and abbreviations are needed to formally represent them. Let λ denote a conjunction of literals, where a literal is an atom p or its negation $\neg p$. We write $\lambda \subseteq \lambda'$, call λ a part (subset) of λ' , if all literals in λ also occur in λ' ; and $\lambda \subset \lambda'$ if $\lambda \subseteq \lambda'$ but not $\lambda' \subseteq \lambda$. In the glossary of Boolean classifiers, s is called an *instance*, λ is called a *term* or *property* (of the instance). The set of terms is noted *Term*. Moreover, let $Atm(\varphi)$ denote the atoms occurring in φ . Finally, notice that the abbreviations $[X]\varphi$ and $\langle X \rangle \varphi$ introduced in Section 2 will be used.

Let us start with *prime implicant*, a key concept in the theory of Boolean functions since [20]. It can be presented in the language $\mathcal{L}(Atm)$ as follows:

$$\text{PImp}(\lambda, x) =_{\text{def}} \Box_I \left(\lambda \rightarrow (\mathbf{t}(x) \wedge \bigwedge_{p \in Atm(\lambda)} \langle Atm(\lambda) \setminus \{p\} \rangle \neg \mathbf{t}(x)) \right).$$

The abbreviation $\text{PImp}(\lambda, x)$ has to be read “ λ is a prime implicant for the classification x ”. Roughly speaking, the latter means that (i) λ necessarily leads to the classification x (why λ is an *implicant*), and (ii) for any of its proper subsets λ' , possibly there is a state where λ' holds but the classification is different from x (why λ is *prime*).

Prime implicant counts as a “global” explanation, in the sense that it is a property of the classifier and holds at *all* its input instances. Partially, as a response to the local approach in model-agnostic methods, researchers from logic-based approaches in XAI focus on the “localized” prime implicant namely *abductive explanation* (AXp).⁴ An abductive explanation is not only a prime implicant, but also a *property of the actual instance*. The notion of abductive explanation is expressed in \mathcal{L} as follows:

$$\text{AXp}(\lambda, x) =_{\text{def}} \lambda \wedge \text{PImp}(\lambda, x).$$

$\text{AXp}(\lambda, x)$ just means that λ is an abductive explanation of the actual classification x . Let us instantiate the notions of prime implicant and abductive explanation in the paper example we introduced in Section 4.1.

Example 2. Take the MCM $\Gamma = (S, \Phi)$ in Example 1, and let $s_1 = \{\text{si}, \text{or}, \text{an}\} \in S$. Consider the function f_1 s.t. $\forall s \in S : f_1(s) = 1$ iff $\text{an} \in s$ and $\{\text{or}, \text{cl}\} \cap s \neq \emptyset$. The function f_1 is syntactically expressed by the formula $\Box_I (\mathbf{t}(1) \leftrightarrow ((\text{or} \wedge \text{an}) \vee (\text{cl} \wedge \text{an})))$. Clearly $f_1 \in \Phi$ for it satisfies the three constraints. Hence, we have:

$$(\Gamma, s_1, f_1) \models \text{AXp}(\text{or} \wedge \text{an}, 1) \wedge \text{PImp}(\text{or} \wedge \text{an}, 1) \wedge \text{PImp}(\text{cl} \wedge \text{an}, 1).$$

Meanwhile $(\Gamma, s_1, f_1) \not\models \text{AXp}(\text{cl} \wedge \text{an}, 1)$, because $(\Gamma, s_1, f_1) \not\models \text{cl} \wedge \text{an}$. But consider $s_2 = \{\text{si}, \text{cl}, \text{an}\} \in S$. We have $(\Gamma, s_2, f_1) \models \text{AXp}(\text{cl} \wedge \text{an}, 1)$.

Now we investigate what happens when facing a black box model $\Gamma = (S, \Phi)$. The agent has uncertainty about the actual classifier’s properties. Therefore, it is interesting to draw the distinction between objective and subjective (or epistemic) explanation. Objective explanation coincides with the notion of explanation in the context of white box classifiers defined above. Subjective explanation refers to the agent’s interpretation of the classifier and her explanation of the classifier’s decision in the light of her partial knowledge.

We say the term λ is a *subjective* prime implicant for x , noted $\text{SubPImp}(\lambda, x)$, if the agent knows that λ is a prime implicant for x , that is:

$$\text{SubPImp}(\lambda, x) =_{\text{def}} \Box_F \text{PImp}(\lambda, x).$$

⁴ It has many names in literature: PI explanation [23], sufficient reason [6]. We adopt the one from [13] for its nice correspondence to contrastive explanation in [12].

Similarly, we say λ is a *subjective* abductive explanation of the actual classification x , noted $\text{SubAXp}(\lambda, x)$, if the agent knows that λ is an abductive explanation of the actual classification x , that is:

$$\text{SubAXp}(\lambda, x) =_{\text{def}} \Box_{\text{F}} \text{AXp}(\lambda, x).$$

It is worth noting that in the case of a white box classifier, if the set of input instances S is finite, we can always find an abductive explanation of the actual classification. That is, for every $\Gamma = (S, \Phi) \in \mathbf{MCM}$, $s \in S$ and $f \in \Phi$:

$$\text{if } S \text{ is finite then } \exists \lambda \in \text{Term such that } (\Gamma, s, f) \models \text{AXp}(\lambda, f(s)).$$

Nonetheless, this result cannot be generalized to the black box case. Indeed, as the following example shows, there is no guarantee for the existence of a subjective explanation of the actual classification. The problem is that the minimality condition can collapse when moving from objective to subjective explanation, since the agent can have more than one classifier in her epistemic state.

Example 3. Let $\Gamma = (S, \Phi)$, f_1 and s_1 be the same as in Example 2. There is no λ such that $(\Gamma, s_1, f_1) \models \Box_{\text{F}} \text{AXp}(\lambda, 1)$. To see this, consider f_2 s.t. $\forall s \in S : f_2(s) = 1$ iff $\{\text{si}, \text{an}\} \subseteq s$. The function f_2 is syntactically expressed by the formula $\Box_{\text{I}}(\text{t}(1) \leftrightarrow (\text{si} \wedge \text{an}))$. Clearly $f_2 \in \Phi$ for it satisfies the three constraints. We have $(\Gamma, s_1, f_2) \models \text{AXp}(\text{si} \wedge \text{an}, 1)$. But there is no term which minimally explains both $f_1(s_1)$ and $f_2(s_1)$. Indeed, $\text{or} \wedge \text{an}$ is not enough for explaining $f_2(s_1)$, $\text{si} \wedge \text{an}$ is not enough for explaining $f_1(s_1)$, and $\text{si} \wedge \text{or} \wedge \text{an}$ fails the minimality condition for both. Therefore, we have

$$(\Gamma, s_1, f_1) \models \text{AXp}(\text{or} \wedge \text{an}, 1) \wedge \bigwedge_{\lambda \in \text{Term}(\{\text{si}, \text{or}, \text{cl}, \text{an}\})} \neg \text{SubAXp}(\lambda, 1).$$

However, this does not mean that the agent knows nothing about the classifier. For instance, she knows that violating the anonymity requirement is a prime implicant for rejection, that is, $(\Gamma, s_1, f_1) \models \text{SubAXp}(\neg \text{an}, 0)$.

To sum up, the four notions of explanation we introduced can be organized in Table 2 along the two dimensions objective vs subjective and local vs global.

	Local	Global
Objective	$\text{AXp}(\lambda, x)$	$\text{PImp}(\lambda, x)$
Subjective	$\text{SubAXp}(\lambda, x)$	$\text{SubPImp}(\lambda, x)$

Table 2. Notions of prime implicant and abductive explanation

5 Dynamic Extension

Before concluding, we are going to present a simple dynamic extension of the language \mathcal{L} by operators of the form $[\varphi]$. They describe the consequences of

removing from the actual model all classifiers that do not *globally* satisfy the constraint φ . More generally, they allow us to model the process of gaining new knowledge about the classifier’s properties. The extended modal language \mathcal{L}^{dyn} is defined by the following grammar:

$$\varphi ::= p \mid \mathbf{t}(x) \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_I \varphi \mid \Box_F \varphi \mid [\varphi]\psi,$$

where p ranges over Atm_0 and x ranges over Val .

The new formula $[\varphi]\psi$ has to be read “ ψ holds after having discarded all classifiers that do not globally satisfy the property φ ”. Notice the similar but different notations $[X]$ and $[\varphi]$. For example, $[\{p\}]$, $[\{p, q\}]$ are abbreviations with ceteris paribus meaning, while $[p]$, $[p \wedge \neg q]$ are dynamic operators.

The interpretation of the operators $[\varphi]$ relative to a pointed MCM (Γ, s, f) with $\Gamma = (S, \Phi)$, $s \in S$ and $f \in \Phi$ goes as follows:

$$(\Gamma, s, f) \models [\varphi]\psi \iff \text{if } (\Gamma, s, f) \models \Box_I \varphi \text{ then } (\Gamma^\varphi, s, f) \models \psi,$$

where $\Gamma^\varphi = (S^\varphi, \Phi^\varphi)$ is the MCM such that:

$$\begin{aligned} S^\varphi &= S, \\ \Phi^\varphi &= \{f' \in \Phi : \forall s' \in S, (\Gamma, s', f') \models \varphi\}. \end{aligned}$$

The previous update semantics for the operator $[\varphi]$ is reminiscent of the semantics of public announcement logic (PAL) [19, 7]. However, there is an important difference. While PAL has a one-dimensional state elimination semantics, our update semantics operates on a single dimension of the product in an MCM. In particular, it only removes classifiers that do not globally satisfy the constraint φ , without modifying the set S of input instances.

The logics D-PLC and D-WPLC (Dynamic PLC and D-WPLC) extend the logic PLC and WPLC by the dynamic operators $[\varphi]$. They are defined as follows.

Definition 8 (Logics D-PLC and D-WPLC). *We define D-PLC (resp. D-WPLC) to be the extension of PLC (resp. WPLC) of Definition 5 (resp. Definition 6) generated by the following reduction axioms for the dynamic operators $[\varphi]$:*

$$\begin{aligned} [\varphi]p &\leftrightarrow (\Box_I \varphi \rightarrow p) \\ [\varphi]\mathbf{t}(x) &\leftrightarrow (\Box_I \varphi \rightarrow \mathbf{t}(x)) \\ [\varphi]\neg\psi &\leftrightarrow (\Box_I \varphi \rightarrow \neg[\varphi]\psi) \\ [\varphi](\psi_1 \wedge \psi_2) &\leftrightarrow ([\varphi]\psi_1 \wedge [\varphi]\psi_2) \\ [\varphi]\Box_I \psi &\leftrightarrow (\Box_I \varphi \rightarrow \Box_I [\varphi]\psi) \\ [\varphi]\Box_F \psi &\leftrightarrow (\Box_I \varphi \rightarrow \Box_F [\varphi]\psi) \end{aligned}$$

and the following rule of inference:

$$\frac{\varphi_1 \leftrightarrow \varphi_2}{\psi \leftrightarrow \psi[\varphi_1/\varphi_2]} \quad (\mathbf{RE})$$

It is a routine exercise to verify that the equivalences in Definition 8 are valid for the class **MCM** and that the rule of replacement of equivalents (**RE**) preserves validity. The completeness of D-PLC and D-WPLC for this class of models follows from Theorem 2 and Corollary 1, in view of the fact that the reduction axioms and the rule of replacement of proved equivalents can be used to find, for any \mathcal{L}^{dyn} -formula, a provably equivalent \mathcal{L} -formula.

Theorem 8. *Let Atm_0 be finite. Then, the logic D-PLC is sound and complete relative to the class **MCM**.*

Theorem 9. *Let Atm_0 be infinite. Then, the logic D-WPLC is sound and complete relative to the class **MCM**.*

The following decidability result is a consequence of Theorem 7 and the fact that via the reduction axioms in Definition 8 we can find a reduction of satisfiability checking of \mathcal{L}^{dyn} -formulas to satisfiability checking of \mathcal{L} -formulas.

Theorem 10. *Checking satisfiability of \mathcal{L}^{dyn} -formulas relative to **MCM** is decidable.*

Let us end up with the paper example to illustrate the expressive power of our dynamic extension.

Example 4. Let $\Gamma = (S, \Phi)$, f_1 and s_1 be the same as in Example 2. We have

$$(\Gamma, s_1, f_1) \models [(\text{or} \wedge \text{an}) \rightarrow \text{t}(1)]\Box_F \bigvee_{\lambda \subseteq (\text{or} \wedge \text{an})} \text{AXp}(\lambda, 1).$$

This means that after having discarded all classifiers which do not take $(\text{or} \wedge \text{an})$ as an implicant for acceptance of a paper, the agent knows that there must be a part of $\text{or} \wedge \text{an}$ that abductively explains the acceptance of the paper s_1 .

6 Conclusion

We have presented a product modal logic which supports reasoning about (i) partial knowledge and uncertainty of a classifier's properties and, (ii) objective and subjective explanations of a classifier's decision. Moreover, we have studied a dynamic extension of the logic which allows us to represent the event of gaining new knowledge about the classifier's properties.

Our logic is intrinsically single-agent: it models the uncertainty of one agent about the actual classifier's properties. In future work, we plan to generalize our framework to the multi-agent setting. The extension would result in a multi-relational product semantics in which every agent has her own epistemic indistinguishability relation which commutes with the input instance dimension (the equivalence relation \sim_{\Box_i} in Definition 3 of MDM). We also plan to enrich this semantics with a knowledge update mechanism in the spirit of Section 5. This would allow us to represent exchange of information between agents with an explanatory purpose, which is named dialogical explanation by philosophers [24] and interactive explanation by researchers in the XAI domain [1, 17].

References

1. Amershi, S., Cakmak, M., Knox, W., Kulesza, T.: Power to the people: The role of humans in interactive machine learning. *AI Magazine* **35**(4), 105–120 (2014)
2. Audemard, G., Bellart, S., Bounia, L., Koriche, F., Lagniez, J.M., Marquis, P.: On the computational intelligibility of boolean classifiers. In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, vol. 18, pp. 74–86 (2021)
3. Baltag, A., van Benthem, J.: A simple logic of functional dependence. *Journal of Philosophical Logic* **50**(5), 939–1005 (2021)
4. Bezhanishvili, N., Hodkinson, I.M.: All normal extensions of $S5$ -squared are finitely axiomatizable. *Studia Logica* **78**(3), 443–457 (2004)
5. Bezhanishvili, N., Marx, M.: All proper normal extensions of $S5$ -square have the polynomial size model property. *Studia Logica* **73**(3), 367–382 (2003)
6. Darwiche, A., Hirth, A.: On the reasons behind decisions. In: *ECAI 2020 - 24th European Conference on Artificial Intelligence*. *Frontiers in Artificial Intelligence and Applications*, vol. 325, pp. 712–720. IOS Press (2020)
7. van Ditmarsch, H., van Der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*, Synthese Library, vol. 337. Springer (2007)
8. Gabbay, D.M., Kurucz, A., Wolter, F., Zakharyashev, M.: *Many-dimensional modal logics: theory and applications*. Elsevier (2003)
9. Grossi, D., Lorini, E., Schwarzentruher, F.: The ceteris paribus structure of logics of game forms. *Journal of Artificial Intelligence Research* **53**, 91–126 (2015)
10. Halpern, J.Y.: The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence* **75**(2), 361–372 (1995)
11. Hempel, C.G., Oppenheim, P.: Studies in the logic of explanation. *Philosophy of science* **15**(2), 135–175 (1948)
12. Ignatiev, A., Narodytska, N., Asher, N., Marques-Silva, J.: From contrastive to abductive explanations and back again. In: *International Conference of the Italian Association for Artificial Intelligence*. pp. 335–355. Springer (2020)
13. Ignatiev, A., Narodytska, N., Marques-Silva, J.: Abduction-based explanations for machine learning models. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 1511–1519 (2019)
14. Kment, B.: Counterfactuals and explanation. *Mind* **115**(458), 261–310 (2006)
15. Liu, X., Lorini, E.: A logic for binary classifiers and their explanation. In: *Proceedings of the 4th International Conference on Logic and Argumentation (CLAR 2021)*. LNCS, vol. 13040, pp. 302–321. Springer (2021)
16. Lundberg, S., Lee, S.I.: A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874* (2017)
17. Miller, T.: Explanation in artificial intelligence: insights from the social sciences. *Artificial Intelligence* **267**, 1–38 (2019)
18. Molnar, C.: *Interpretable machine learning*. Lulu. com (2020)
19. Plaza, J.: Logics of public communications. *Synthese* **158**(2), 165–179 (2007)
20. Quine, W.V.: A way to simplify truth functions. *The American mathematical monthly* **62**(9), 627–631 (1955)
21. Ribeiro, M.T., Singh, S., Guestrin, C.: "Why should i trust you?" Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. pp. 1135–1144 (2016)

22. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18). vol. 32 (2018)
23. Shih, A., Choi, A., Darwiche, A.: Formal verification of bayesian network classifiers. In: International Conference on Probabilistic Graphical Models. pp. 427–438. PMLR (2018)
24. Walton, D.: A new dialectical theory of explanation. *Philosophical Explorations* **7**(1), 71–89 (2004)
25. Yang, F., Väänänen, J.: Propositional logics of dependence. *Annals of Pure and Applied Logic* **167**(7), 557–589 (2016)

A Technical Annex

A.1 Proof of Theorem 1

Proof. We start with the left-to-right direction of the proof. Let (Γ, s_0, f_0) be a pointed MCM with $\Gamma = (S, \Phi)$, $S \subseteq 2^{Atm_0}$ and $\Phi \subseteq F_S$ such that $(\Gamma, s_0, f_0) \models \varphi$. We define the tuple $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ as follows:

- $W = \{(s, f) : s \in S \text{ and } f \in \Phi\}$,
- $\forall (s, f), (s', f') \in W, (s, f) \sim_{\square_I} (s', f') \text{ iff } f = f'$
- $\forall (s, f), (s', f') \in W, (s, f) \sim_{\square_F} (s', f') \text{ iff } s = s'$,
- $\forall (s, f) \in W, V(s, f) = s \cup \{t(f(s))\}$.

It is routine exercise to verify that M so defined is an MDM. Moreover, by induction on the structure of φ , it is easy to prove that “ $(\Gamma, s, f) \models \varphi$ iff $(M, (s, f)) \models \varphi$ ” for every $s \in S$ and $f \in \Phi$. Thus, $(M, (s_0, f_0)) \models \varphi$ since $(\Gamma, s_0, f_0) \models \varphi$.

Let us now prove the right-to-left direction. Let $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ be an MDM and $w_0 \in W$ such that $(M, w_0) \models \varphi$. Given $v \in W$, let $|v| = \{u \in W : v \sim_{\square_I} u \text{ and } V(v) = V(u)\}$. We transform the MDM M into a tuple $M' = (W', \sim'_{\square_I}, \sim'_{\square_F}, V')$ such that:

- $W' = \{|v| : v \in W\}$,
- $\forall |v|, |u| \in W', |v| \sim_{\square_I} |u| \text{ iff } \exists v' \in |v|, u' \in |u| \text{ such that } v' \sim_{\square_I} u'$,
- $\forall |v|, |u| \in W', |v| \sim'_{\square_F} |u| \text{ iff } \exists v' \in |v|, u' \in |u| \text{ such that } v' \sim_{\square_F} u'$,
- $\forall |v| \in W', V'(|v|) = V(v)$.

Like what we did for V , let $V'_Y(|v|) = V'(|v|) \cap Y$ for all $Y \subseteq Atm$.

It is a routine exercise to verify that M' is an MDM and, by induction on the structure of φ , that “ $(M, v) \models \varphi$ iff $(M', |v|) \models \varphi$ ” for every $v \in W$. Thus, $(M, |w_0|) \models \varphi$ since $(M, w_0) \models \varphi$. Finally, because of Constraints **C2** and **C3** in Definition 3, the following property holds:

$$(\mathbf{C6}) \quad (\sim'_{\square_I} \cap \sim'_{\square_F}) = id_{W'},$$

where $id_{W'}$ is the identity relation on W' .

Let W'/\sim'_{\square_I} be the quotient set of W' by the equivalence relation \sim'_{\square_I} . We note τ, τ', \dots its elements. Given $\tau, \tau' \in W'/\sim'_{\square_I}$, we write $\tau \approx_F \tau'$ if and only

if, $\forall |v| \in \tau, \forall |u| \in \tau'$, if $V'_{Atm_0}(|v|) = V'_{Atm_0}(|u|)$ then $V'_{Dec}(|v|) = V'_{Dec}(|u|)$. Given $|v|, |u| \in W'$, we write $|v| \simeq |u|$ if and only if $\exists \tau, \tau' \in W' / \sim'_{\square_I}$ such that $|v| \in \tau, |u| \in \tau', \tau \approx_F \tau'$ and $V'_{Atm_0}(|v|) = V'_{Atm_0}(|u|)$. Clearly, \approx_F and \simeq are equivalence relations.

We are going to transform the MDM M' into an MDM which does not contain multiple copies of the same function and which satisfies the same formulas as M' . We define it to be a tuple $M'' = (W'', \sim''_{\square_I}, \sim''_{\square_F}, V'')$ such that:

- $W'' = \{\simeq(|v|) : |v| \in W'\}$,
- $\forall \simeq(|v|), \simeq(|u|) \in W'', \simeq(|v|) \sim''_{\square_I} \simeq(|u|)$ iff $\exists |v'| \in \simeq(|v|), |u'| \in \simeq(|u|)$ such that $|v'| \sim'_{\square_I} |u'|$,
- $\forall \simeq(|v|), \simeq(|u|) \in W'', \simeq(|v|) \sim''_{\square_F} \simeq(|u|)$ iff $\exists |v'| \in \simeq(|v|), |u'| \in \simeq(|u|)$ such that $|v'| \sim'_{\square_F} |u'|$,
- $\forall \simeq(|v|) \in W'', V''(\simeq(|v|)) = V'(|v|)$.

Again, it is routine to verify that M'' is an MDM which satisfies the previous Constraint **C6**. Moreover, by induction on the structure of φ , it is easy to prove that “ $(M', |v|) \models \varphi$ iff $(M'', \simeq(|v|)) \models \varphi$ ” for every $|v| \in W'$. Thus, $(M'', \simeq(|w_0|)) \models \varphi$ since $(M', w_0) \models \varphi$.

We can easily build an MCM isomorphic to M'' .

A.2 Proof of Theorem 3

Proof. The left-to-right direction is trivial. We prove the right-to-left direction. Suppose Atm_0 is infinite. Moreover, let $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ be a finite quasi-MDM and $w_0 \in W$ such that $(M, w_0) \models \varphi$. Since Atm_0 is infinite and W is finite, we can define an injection $g : W \rightarrow Atm_0 \setminus Atm(\varphi)$. We define the tuple $M' = (W', \sim'_{\square_I}, \sim'_{\square_F}, V')$ as follows:

- $W' = W$;
- $\sim'_{\square_I} = \sim_{\square_I}$;
- $\sim'_{\square_F} = \sim_{\square_F}$;
- for every $w \in W$,

$$V'(w) = (V(w) \setminus \{g(v) : v \in W \text{ and } w \neq v\}) \cup \{g(w)\}.$$

It is routine to verify that M' is a finite MDM. Indeed, $V'_{Atm_0}(w) \neq V'_{Atm_0}(v)$ for all $w, v \in W'$ such that $w \neq v$. This guarantees that M' satisfies the “functionality” constraint **C2**. Moreover, by induction on the structure of φ , it is straightforward to prove that “ $(M, v) \models \varphi$ iff $(M', v) \models \varphi$ ” for every $v \in W$. The crucial point of the proof is that $\sim'_{\square_I} = \sim_{\square_I}$ and $\sim'_{\square_F} = \sim_{\square_F}$. Thus, $(M', w_0) \models \varphi$ since $(M, w_0) \models \varphi$.

A.3 Proof of Theorem 4

Proof. The right-to-left direction is clear. We are going to prove the left-to-right direction by a filtration argument.

Let $M = (W, \sim_{\square_I}, \sim_{\square_F}, V)$ be a quasi-MDM and $w_0 \in W$ such that $(M, w_0) \models \varphi$. It is routine to verify that $(\sim_{\square_I} \cup \sim_{\square_F})^* = \sim_{\square_I} \circ \sim_{\square_F} = \sim_{\square_F} \circ \sim_{\square_I}$. Thus, we can define $M' = (W', \sim'_{\square_I}, \sim'_{\square_F}, V')$ to be the submodel of M generated from w_0 through the relation $\sim_{\square_I} \circ \sim_{\square_F}$. M' is a quasi-MDM and $(M', w_0) \models \varphi$.

Let $sf(\varphi)$ be the set of all subformulas of φ and let $sf^+(\varphi) = sf(\varphi) \cup Dec$. Moreover, for every $v \in W'$, let $\Theta(v) = \{\psi \in sf^+(\varphi) : (M', v) \models \psi\}$. For every $v, u \in W'$, we define

$$v \simeq u \text{ iff } \Theta(v) = \Theta(u).$$

Moreover, we define $[v] = \{u \in W' : v \simeq u\}$.

We construct a new model $M'' = (W'', \sim''_{\square_I}, \sim''_{\square_F}, V'')$ where:

- $W'' = \{[v] : v \in W'\};$
- $[v] \sim''_{\square_I} [u]$ iff

$$\forall \square_I \psi \in sf(\varphi), ((M', v) \models \square_I \psi \text{ iff } (M', u) \models \square_I \psi);$$
- $[v] \sim''_{\square_F} [u]$ iff

$$\forall \square_F \psi \in sf(\varphi), ((M', v) \models \square_F \psi \text{ iff } (M', u) \models \square_F \psi) \text{ and}$$

$$\forall p \in sf(\varphi) \cap Atm_0, ((M', v) \models p \text{ iff } (M', u) \models p);$$
- $V''([v]) = V'_{sf(\varphi) \cap Atm_0}(v) \cup V'_{Dec}(v).$

M'' is indeed a filtration, for it satisfies that if $v \sim_{\blacksquare} u$, then $[v] \sim_{\blacksquare} [u]$; and if $\blacksquare \psi \in sf(\varphi)$ and $(M', v) \models \blacksquare \psi$, then $(M', u) \models \psi$, for $\blacksquare \in \{\square_I, \square_F\}$. Additionally, the valuation function is defined in standard way.

To check that M'' is a finite quasi-MDM, we go through all constraints. For **C1** a crucial fact is that M' generated from w_0 through $\sim_{\square_I} \circ \sim_{\square_F}$, viz. $\forall v, u \in W', v \sim'_{\square_I} u$ and $v \sim'_{\square_F} u$.

To see that fact, by construction of M' we have $w_0 \sim'_{\square_I} u$ and $w_0 \sim'_{\square_F} v$. This means $w_0 \sim'_{\square_I} v_1 \sim'_{\square_F} v$ and $w_0 \sim'_{\square_I} u_1 \sim'_{\square_F} u$ for some $u_1, v_1 \in W' \subseteq W$. Then we have $v_1 \sim'_{\square_I} u$ by the Euclidean of \sim'_{\square_I} . Then, since $v_1, u \in W$ and by **C1** of M , we have $v_1 \sim_{\square_F} v_2 \sim_{\square_I} u$ for some $v_2 \in W$. Since $w_0 \sim_{\square_I} v_1 \sim_{\square_F} v_2$, we are sure that $v_2 \in W'$, which gives us $v_1 \sim'_{\square_F} v_2 \sim'_{\square_I} u$. So now we have $v \sim'_{\square_F} v_1 \sim'_{\square_F} v_2 \sim'_{\square_I} u$, and by Euclidean of \sim'_{\square_F} , we have $v \sim'_{\square_F} u$. The case of $u \sim'_{\square_I} v$ is proven in the same way.

C3 holds because of the definition of \sim''_{\square_F} . **C4, C5** hold, since V'' not only considers $sf(\varphi) \cap Atm_0$ but also Dec .

It is routine to verify that $M'' = (W'', \sim''_{\square_I}, \sim''_{\square_F}, V'')$ is a filtration of M' and is a finite quasi-MDM. Therefore, $(M'', [w_0]) \models \varphi$.

A.4 Proof of Theorem 6

Proof. Suppose $|Atm_0|$ is finite. Then, the class **MCM** is bounded by some integer k . So, in order to determine whether a formula φ is satisfiable for the class **MCM**, it is sufficient to verify whether φ is satisfied in one of these MCMs. This verification takes a polynomial time in the size of φ since it is a repeated model checking and model checking in the product modal logic $S5^2$ is polynomial.

A.5 Proof of Theorem 7

Proof. The satisfiability problem of formulas in \mathcal{L} relative to the class **MCM** is clearly NP-hard since there exists a polynomial-time reduction of SAT to it. As for membership, WPLC is a proper normal extension of the product modal logic S5². It is proved in [4] (see also [5]) that all proper normal extensions of the product modal S5² are in NP. WPLC is moreover sound and complete relative to the class **MCM** (Corollary 2) which implies that checking whether a formula is consistent in the logic amounts to checking whether the formula is satisfiable relative to the class **MCM**. So, we can conclude that satisfiability checking of \mathcal{L} -formulas in class **MCM** is in NP.

A.6 Proof of the validity of the sixth equivalence in Definition 8

$$\begin{aligned}
(\Gamma, s, f) \models [\varphi]\Box_F\psi &\iff \text{if } (\Gamma, s, f) \models \Box_I\varphi \text{ then } (\Gamma^\varphi, s, f) \models \Box_F\psi; \\
&\iff \text{if } (\Gamma, s, f) \models \Box_I\varphi \text{ then } \forall f' \in \Phi^\varphi, (\Gamma^\varphi, s, f') \models \psi; \\
&\iff \text{if } (\Gamma, s, f) \models \Box_I\varphi \text{ then } \forall f' \in \Phi, \\
&\quad (\text{if } \forall s' \in S, (\Gamma, s', f') \models \psi \text{ then } (\Gamma^\varphi, s, f') \models \psi); \\
&\iff \text{if } (\Gamma, s, f) \models \Box_I\varphi \text{ then } \forall f' \in \Phi, \\
&\quad (\text{if } (\Gamma, s, f') \models \Box_I\psi \text{ then } (\Gamma^\varphi, s, f') \models \psi); \\
&\iff \text{if } (\Gamma, s, f) \models \Box_I\varphi \text{ then } \forall f' \in \Phi, (\Gamma, s, f') \models [\varphi]\psi; \\
&\iff (\Gamma, s, f) \models \Box_I\varphi \rightarrow \Box_F[\varphi]\psi.
\end{aligned}$$