

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/184377/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Ozaki, Ana and Santos Ribeiro Santos, Jandson 2026. Model change for description logic concepts. Presented at: The 40th Annual AAAI Conference on Artificial Intelligence, Singapore, 20 - 27 January 2026. Published in: Koenig, Sven, Jenkins, Chad and Taylor, Matthew E. eds. Proceedings of the AAAI Conference on Artificial Intelligence. AAAI Press, pp. 19321-19328. 10.1609/aaai.v40i23.39008

Publishers page: <https://doi.org/10.1609/aaai.v40i23.39008>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Model Change for Description Logic Concepts

Ana Ozaki¹, Jandson S. Ribeiro²

¹University of Oslo, Norway

²Cardiff University, UK

ana.ozaki@uib.no, ribeiroj@cardiff.ac.uk

Abstract

We consider the problem of modifying a description logic concept in light of models represented as pointed interpretations. We call this setting *model change*, and distinguish three main kinds of changes: eviction, which consists of only removing models; reception, which incorporates models; and revision, which combines removal with incorporation of models in a single operation. We introduce a formal notion of revision and argue that it does not reduce to a simple combination of eviction and reception, contrary to intuition. We provide positive and negative results on the compatibility of eviction and reception for $\mathcal{EL}_\perp$ and $\mathcal{ALC}$ description logic concepts and on the compatibility of revision for $\mathcal{ALC}$ concepts.

1 Introduction

Keeping beliefs updated is a central problem in knowledge representation, that has been investigated in the context of different logics and applications. In the main streams of belief change research, the belief base is finite and the pieces of information expressing how to modify it are expressed as sets of formulae in the underlying logic (Hansson 1999; Gärdenfors 1988; Alchourrón, Gärdenfors, and Makinson 1985). In many scenarios, however, using sets of models to specify the observed change is more suitable than using formulae. This is well studied in the context of *learning from interpretations* (De Raedt 1997), where the goal is to find a concise formula that is consistent with models labelled as positive or negative. In the context of description logic (DL), the process of building an ontology usually goes through stages where the person creating it studies possible models of world, discarding models when they are proven false and adding new models previously not considered. The next examples illustrate model change operations for DL concepts.

Example 1. *Araci is visiting a zoo in Australia and knows little about Australian animals. She knows that a platypus is a mammal that lays eggs, so her belief on platypus is*

$$\text{Platypus} \equiv \text{Mammal} \sqcap (\exists \text{lays.Egg}).$$

She sees a platypus ‘d’ but knows nothing about their diet. So, she entertains the two following possible worlds $(\mathcal{I}_1, d)$

and $(\mathcal{I}_2, d)$ with $\{d, e\} \subseteq \Delta^{\mathcal{I}_1} = \Delta^{\mathcal{I}_2}$:

$\mathcal{I}_1 :$	$\begin{array}{l} d \in \text{Mammal}^{\mathcal{I}_1} \\ d \notin \text{Herbivore}^{\mathcal{I}_1} \\ (d, e) \in \text{lays}^{\mathcal{I}_1} \\ e \in \text{Egg}^{\mathcal{I}_1} \end{array}$	$\mathcal{I}_2 :$	$\begin{array}{l} d \in \text{Mammal}^{\mathcal{I}_2} \\ d \in \text{Herbivore}^{\mathcal{I}_2} \\ (d, e) \in \text{lays}^{\mathcal{I}_2} \\ e \in \text{Egg}^{\mathcal{I}_2} \end{array}$
-------------------	---	-------------------	--

She catches the platypus eating a small insect, which makes Araci retract the pointed interpretation $(\mathcal{I}_2, d)$. So, she changes her conceptual beliefs about platypuses to

$$\text{Platypus} \equiv \text{Mammal} \sqcap (\exists \text{lays.Egg}) \sqcap \neg \text{Herbivore}.$$

Ex. 1 illustrates the process of removing a model from the concept, producing a concept closer to the real world. This change operation is called *eviction* (Guimarães, Ozaki, and Ribeiro 2023), and should minimally modify the concept. The next example illustrates its dual: the process of adding a model, called *reception*.

Example 2. *Araci knows that kangaroos and koalas are marsupials, but knows little about Tasmanian devils. So, her beliefs about marsupials are $\text{Marsupial} \equiv \text{Koala} \sqcup \text{Kangaroo}$. In the Tasmanian devils section, she sees a devil ‘d’ and reads a sign with the information that tasmanian devils are marsupials. So, she now admits a world $(\mathcal{I}_3, d')$ where*

$$\mathcal{I}_3 : \begin{array}{l} d' \in \text{TasDevil}^{\mathcal{I}_3}, d' \in \text{Carnivore}^{\mathcal{I}_3}, d' \in \text{Marsupial}^{\mathcal{I}_3} \end{array}$$

She changes her conceptual beliefs to comply with $(\mathcal{I}_3, d')$:

$$\text{Marsupial} \equiv \text{Koala} \sqcup \text{Kangaroo} \sqcup \text{TasDevil}.$$

In both examples, Araci minimally modified her concepts. For platypus, she modified only concepts related to the diet, while for the Tasmanian devil, she modified only concepts related to marsupials and devils; no further changes related to kangaroos or koalas were carried out. Ensuring minimal change in this setting is a challenge, as it is not always possible to retract or incorporate the single input model (Guimarães, Ozaki, and Ribeiro 2023). A third and more complex kind of operation is illustrated in Ex. 3, where one must add and remove models in a single step.

Example 3. *Araci believes that koalas are marsupial mammals that are not placental. So her concept on koalas is:*

$$\text{Koala} \equiv \text{Mammal} \sqcap \text{Marsupial} \sqcap \neg \text{Placental}.$$

She sees a koala ‘ d'' ’ in the zoo, so she considers the model $(\mathcal{I}_4, d'')$ possible, but not $(\mathcal{I}_5, d'')$.

$\mathcal{I}_4 :$	$d'' \in \text{Mammal}^{\mathcal{I}_4}$	$\mathcal{I}_5 :$	$d'' \in \text{Mammal}^{\mathcal{I}_5}$
	$d'' \in \text{Marsupial}^{\mathcal{I}_4}$		$d'' \in \text{Marsupial}^{\mathcal{I}_5}$
	$d'' \notin \text{Placental}^{\mathcal{I}_4}$		$d'' \in \text{Placental}^{\mathcal{I}_5}$

She reads a sign informing that koalas are actually placental. So she is compelled to comply with $(\mathcal{I}_5, d'')$, whereas retracting $(\mathcal{I}_4, d'')$. Therefore, she changes her belief to

$$\text{Koala} \equiv \text{Mammal} \sqcap \text{Marsupial} \sqcap \text{Placental}.$$

In Ex. 3, the model $(\mathcal{I}_4, d'')$ had to be removed, while $(\mathcal{I}_5, d'')$ is added. We call this more complex operation a *revision*. After providing preliminaries (Sec. 2) and recalling eviction and reception (Sec. 3), we present our main contributions in this paper, which are:

- a theoretical study on eviction and reception for DL concepts in $\mathcal{ALC}$ and $\mathcal{EL}_\perp$ (Sec. 4); and
- the introduction of the notion of model revision (Sec. 5), with some results for DL concepts (Sec. 6).

Contrary to intuition, revision *does not* correspond to a serial combination of eviction and reception, as we show in this paper. In general, it may well be impossible to add and remove exactly and only the models of the input.

Related Work Our contribution is closest to the work by (Guimarães, Ozaki, and Ribeiro 2023), with the main differences being that they neither consider revision nor DL concepts. Although their results on eviction and reception of DL ontologies inspired some of our proofs, the settings are considerably different, with the DL concept case that we present here being arguably more natural to the process of modelling concepts, which is fundamental for building ontologies. In this line, there has been work on computing least common subsumers, which generalize a set of concepts with minimal change (Baader, Sertkaya, and Turhan 2007). Recent work considered learnability of DL concept expressions with pointed interpretations labelled as positive and negative (ten Cate, Koudijs, and Ozaki 2024). One can see the positive and negative pointed interpretations in their work as the sets of models to be added and removed in our setting. The difference from their work to ours is that here we focus on the existence of operators following minimal change rationality postulates, while they focus on the existence of sets of labelled models that characterize a DL concept, that is, that can be used to distinguish a particular concept from all the others in a DL language for concepts. Ontology learning from interpretations has been investigated by (Klarman and Britz 2015). Other works investigated learnability of DL concepts in a data retrieval setting (Funk, Jung, and Lutz 2021; Funk et al. 2019), using inductive logic programming (Lehmann 2009; Fanizzi, d’Amato, and Esposito 2008; Lehmann and Haase 2009; Lehmann and Hitzler 2010), and with counterfactuals (Iannone, Palmisano, and Fanizzi 2007). We also point out works relating learning, epistemic logic, and belief revision (Baltag, Gierasimczuk, and Smets 2019; Baltag et al. 2019; Ozaki and Troquard 2019; Schwind et al. 2025).

2 Preliminaries

The power set of a set A is denoted by $\wp(A)$, while the set of all finite subsets of A is denoted by $\wp_f(A)$. Given a pre-order $\leq \subseteq \mathcal{D} \times \mathcal{D}$ on a domain $\mathcal{D}$, and a set $A \subseteq \mathcal{D}$, the set of all maximal and minimal elements of A w.r.t. $\leq$ are respectively $\max_{\leq}(A) = \{x \in A \mid \text{for all } y \in A \text{ if } x \leq y, \text{ then } y \leq x\}$, and $\min_{\leq}(A) = \{x \in A \mid \text{for all } y \in A \text{ if } y \leq x, \text{ then } x \leq y\}$. We write $\wp^*(A)$ to denote the non-empty subsets of A . Following Aiguier et al. (2018); Delgrande, Peppas, and Woltran (2018), and Guimarães, Ozaki, and Ribeiro (2023), we use satisfaction systems to define logics. A *satisfaction system* is a triple $\Lambda = (\mathcal{L}, \mathfrak{M}, \models)$, where $\mathcal{L}$ is a non-empty (possibly countably infinite) language, $\mathfrak{M}$ is a set of models, and $\models \subseteq \mathfrak{M} \times \wp(\mathcal{L})$ is a relation, called the *satisfaction relation*, which relates models to subsets of the language. We use the infix notation $M \models B$ as a shorthand for $(M, B) \in \models$ and say that M *satisfies* B . Every subset of $\mathcal{L}$ is called a *base* (which can be finite or infinite), denoted B . We denote by $\text{mod}_\Lambda(B)$ the set $\{M \in \mathfrak{M} \mid M \models B\}$. We write $\text{mod}(B)$ when the satisfaction system is clear from the context. Satisfaction systems facilitate the generalisation of some results that do not depend on certain properties of the consequence relation of the logic. A set of models $\mathbb{M} \subseteq \mathfrak{M}$ within Λ is *finitely representable* iff there is $B \in \wp_f(\mathcal{L})$ such that $\text{mod}(B) = \mathbb{M}$. Let $\text{FR}(\Lambda)$ denote all *finitely representable sets of models in Λ* , i.e.,

$$\text{FR}(\Lambda) = \{\mathbb{M} \subseteq \mathfrak{M} \mid \exists B \in \wp_f(\mathcal{L}) : \text{mod}(B) = \mathbb{M}\}.$$

Given a set $\mathbb{M}$ of models, the greatest finitely-representable subsets of $\mathbb{M}$ and the least finitely-representable supersets of $\mathbb{M}$ are given respectively by

$$\text{MaxFRSubs}(\mathbb{M}, \Lambda) = \max_{\subseteq}(\{ \mathbb{M}' \in \text{FR}(\Lambda) \mid \mathbb{M}' \subseteq \mathbb{M} \}),$$

$$\text{MinFRSups}(\mathbb{M}, \Lambda) = \min_{\subseteq}(\{ \mathbb{M}' \in \text{FR}(\Lambda) \mid \mathbb{M} \subseteq \mathbb{M}' \}).$$

We say that a set of formulae $B \subseteq \mathcal{L}$ is *finitely representable* iff there is $B' \in \wp_f(\mathcal{L})$ with $\text{mod}(B) = \text{mod}(B')$. We write $\times$ for the Cartesian product of two sets. Also, we denote the logical closure of a base in a satisfaction system Λ by Cn_Λ , omitting the subscript when clear from the context.

DL Concepts Let N_C and N_R be countable and pairwise disjoint sets of concept names and role names, respectively. In this work, N_C and N_R can be finite or infinite. We explicitly indicate when these sets, called *signature*, are finite (otherwise, they are assumed to be infinite). $\mathcal{EL}$ concepts are built according to the rule: $C, D ::= \top \mid A \mid (C \sqcap D) \mid (\exists r.C)$, where $A \in N_C$ and $r \in N_R$. $\mathcal{EL}_\perp$ concepts extend $\mathcal{EL}$ by allowing $\perp$ (interpreted as the empty set). $\mathcal{ALC}$ concepts extend $\mathcal{EL}$ concepts with the rule $\neg C$ (recall that $C \sqcap \neg C$ is equivalent to $\perp$, so $\mathcal{ALC}$ extends $\mathcal{EL}_\perp$). We may write $\exists r^n. \top$, with $n \in \mathbb{N}$, as a shorthand for the nesting of n existential quantifiers (that is, $\exists r^{n+1}. \top = \exists r. (\exists r^n. \top)$) and $\exists r^0. \top = \top$.

Semantics The semantics of $\mathcal{EL}$, $\mathcal{EL}_\perp$, and $\mathcal{ALC}$ concepts is defined using pointed interpretations (Baader et al. 2017; Agi et al. 2003). An interpretation $\mathcal{I}$ is a pair $(\Delta^\mathcal{I}, \cdot^\mathcal{I})$, where $\Delta^\mathcal{I}$ is a non-empty set, called the *domain*, and $\cdot^\mathcal{I}$ is a function that maps every $A \in N_C$ to a subset of $\Delta^\mathcal{I}$ and every $r \in N_R$ to a subset of $\Delta^\mathcal{I} \times \Delta^\mathcal{I}$. It is *finite* if $\Delta^\mathcal{I}$ is finite. A *pointed*

interpretation is a pair $(\mathcal{I}, d)$ where $\mathcal{I} = (\cdot^{\mathcal{I}}, \Delta^{\mathcal{I}})$ is an interpretation and $d \in \Delta^{\mathcal{I}}$. A pointed interpretation $(\mathcal{I}, d)$ satisfies a concept C iff $d \in C^{\mathcal{I}}$. We define tree-shaped pointed interpretations using the notion of unfolding (Dummett and Lemmon 1959) (see also (Konev et al. 2016)). We say that a concept C entails a concept D if for all pointed interpretations $(\mathcal{I}, d)$ (over the signature), $d \in C^{\mathcal{I}}$ implies $d \in D^{\mathcal{I}}$. Two concepts C, D are *equivalent*, written $C \equiv D$, iff C entails D and D entails C . Given a fixed but arbitrary $r \in \mathbb{N}_R$, we define $M^n = (\mathbb{N}, \cdot^{M^n})$ where $r^{M^n} = \{(i, i+1) \mid i \in \mathbb{N}, 0 \leq i < n\}$ and similarly $M^\infty = (\mathbb{N}, \cdot^{M^\infty})$ where $r^{M^\infty} = \{(i, i+1) \mid i \in \mathbb{N}\}$.

DL concepts in Satisfaction Systems In a satisfaction system, we use the term ‘base’ for a subset of formulas in a logic language. We treat a set of concepts and a concept formed by the conjunction of the elements of the set interchangeably. So we may refer to a base as a concept or as a finite set of concepts (in the latter, we mean the concept formed by the conjunction). The notion of a ‘model’ in a satisfaction system corresponds to the notion of a pointed interpretation.

3 Model Reception and Eviction

This section addresses the problem of modifying a finite base in light of a set of models, as illustrated in Ex. 1 and Ex. 2. We call such kinds of operators *model change operators*. Guimarães, Ozaki, and Ribeiro (2023) distinguished two main primitive kinds of model change operators:

eviction: remove a set $\mathbb{M}$ of models from a base $\mathcal{B}$, that is, turn $\mathcal{B}$ into a base $\mathcal{B}'$ whose models are not in $\mathbb{M}$;

reception: incorporate all models from the set $\mathbb{M}$ into a base $\mathcal{B}$, that is, turn $\mathcal{B}$ into a base $\mathcal{B}'$ such that all models in $\mathbb{M}$ satisfy $\mathcal{B}'$.

In some scenarios, not all sets of models need to be taken into account. For instance, some DLs have the finite model property or the tree-shaped property. So, it makes sense to also consider model change operators that only take into account such classes of models. Formally, a class of models on a satisfaction system $\Lambda = (\mathcal{L}, \mathfrak{M}, \models)$ is a set $\mathcal{C} \subseteq \mathcal{P}(\mathfrak{M})$ with sets of interpretations. Model change operators are, therefore, defined on a given class of models.

Definition 4. A model change operator in a class $\mathcal{C}$ of models is a function $\circ : \mathcal{P}_f(\mathcal{L}) \times \mathcal{C} \rightarrow \mathcal{P}_f(\mathcal{L})$, mapping each finite base $\mathcal{B}$ into a finite base $\mathcal{B}'$ in light of a set of models.

For conciseness, we may omit the reference to the class $\mathcal{C}$ of models if $\mathcal{C}$ is clear from the context. The main challenge of model change operators is to guarantee the finiteness of the new base, which presents two main hurdles:

1. some sets of models cannot be uniquely added/removed to/from some bases, as Ex. 5 below illustrates.
2. the simple addition/removal of a set of models might not be finitely representable. This occurs because the language of logic is not expressive enough to distinguish all the models in a set from those which are not in the set. This issue is illustrated at Ex. 6.

Example 5. Consider the $\mathcal{EL}$ concept $\exists r^3.\top$, with $N_C = \emptyset$, $N_R = \{r\}$, and the pointed models $(\mathcal{I}_1, d_1)$ and $(\mathcal{I}_2, d_1)$ with

domains $\Delta^{\mathcal{I}_1} = \{d_1, d_2\}$ and $\Delta^{\mathcal{I}_2} = \{d_1, d_2, d_3\}$, where

$$\mathcal{I}_1 : r^{\mathcal{I}_1} = \{(d_1, d_2)\} \quad \mathcal{I}_2 : r^{\mathcal{I}_2} = \{(d_1, d_2), (d_2, d_3)\}$$

Neither $(\mathcal{I}_1, d_1)$ nor $(\mathcal{I}_2, d_1)$ are models of $\exists r^3.\top$. Suppose that we want to add $(\mathcal{I}_1, d_1)$ to the models of $\exists r^3.\top$. In $\mathcal{EL}_\perp$, however, every concept satisfied by $(\mathcal{I}_1, d_1)$ and $\text{mod}(\exists r^3.\top)$ is also satisfied by $(\mathcal{I}_2, d_1)$. There is no concept separating $(\mathcal{I}_2, d_1)$ from $\text{mod}(\exists r^3.\top) \cup \{(\mathcal{I}_1, d_1)\}$.

Example 6. Let $\mathcal{B} = \{\exists r.\top\}$ be an $\mathcal{EL}$ concept and let the signature be $N_C = \{A\}$, $N_R = \{r\}$. We want to evict the pointed model $(\mathcal{I}, d)$ with $\Delta^{\mathcal{I}} = \{d\}$, $r^{\mathcal{I}} = \{(d, d)\}$, and $A^{\mathcal{I}} = \emptyset$. We have that $(\mathcal{I}, d)$ satisfies $\mathcal{B}$. The removal of $(\mathcal{I}, d)$ from $\text{mod}(\mathcal{B})$ yields the base $\mathcal{B}' = \{\exists r.A, \exists r^2.A, \dots, \exists r^n.A, \dots\}$, which is not finitely representable.

On both cases 1 and 2, as illustrated respectively on Ex. 5 and Ex. 6, extra models must be added/removed to achieve a finite base. Such addition/removal should be minimised, so only models that do contribute to reaching finiteness are considered. In this case, a “closest” finite base is produced. Such minimality criteria are properly addressed in the form of rationality postulates. The appropriate class of all eviction/reception operators abiding by such postulates are identified. We show that in several classes of models, the minimality principles cannot be guaranteed, which is due to the non-existence of a “closest” finite base, known as the compatibility problem (Guimarães, Ozaki, and Ribeiro 2023). We briefly review reception and eviction, respectively, on Sec. 3.1 and Sec. 3.2.

3.1 Reception

We denote belief change operators related to reception as rcp. Guimarães, Ozaki, and Ribeiro (2023) proposed the following rationality postulates to govern reception

(success) $\mathbb{M} \subseteq \text{mod}(\text{rcp}(\mathcal{B}, \mathbb{M}))$

(persistence) $\text{mod}(\mathcal{B}) \subseteq \text{mod}(\text{rcp}(\mathcal{B}, \mathbb{M}))$

(finite temperance) $\mathbb{M}' \notin \text{FR}(\Lambda)$, if $\text{mod}(\mathcal{B}) \cup \mathbb{M} \subseteq \mathbb{M}'$ and $\mathbb{M}' \subset \text{mod}(\text{rcp}(\mathcal{B}, \mathbb{M}))$

(uniformity) $\text{mod}(\text{rcp}(\mathcal{B}, \mathbb{M})) = \text{mod}(\text{rcp}(\mathcal{B}', \mathbb{M}'))$, if $\text{MinFRSup}(\text{mod}(\mathcal{B}) \cup \mathbb{M}, \Lambda) = \text{MinFRSup}(\text{mod}(\mathcal{B}') \cup \mathbb{M}', \Lambda)$.

Success ensures that each model from $\mathbb{M}$ must be incorporated. The purpose of reception is to accommodate new models. *Persistence* ensures that no model is removed in the process. *Finite temperance* ensures that the addition of extra models is minimized, adding only models that contribute to reaching a finite base. *Uniformity* ensures that reception is neither syntax sensitive nor sensitive to model structure. For instance, in Ex. 5, the pointed interpretation $(\mathcal{I}_1, d_1)$ cannot be separated from $(\mathcal{I}_2, d_1)$ in the presence of the models of $\mathcal{B}$, so reception, on $\mathcal{B}$, of the sets $\{(\mathcal{I}_1, d_1)\}$ and $\{(\mathcal{I}_1, d_1), (\mathcal{I}_2, d_1)\}$ must coincide.

A *reception operator* on a class $\mathcal{C}$ of models is a model change operator $\text{rcp} : \mathcal{P}_f(\mathcal{L}) \times \mathcal{C} \rightarrow \mathcal{P}_f(\mathcal{L})$ that satisfies *success*. A reception operator satisfying all rationality postulates is called *rational*. When there is no finite base for $\text{mod}(\mathcal{B}) \cup \mathbb{M}$, some further models must be added in favor of finiteness.

Such extra removal must be minimized as finite-temperance demands. This corresponds to picking a least finitely representable superset of $\text{mod}(\mathcal{B}) \cup \mathbb{M}$, that is, picking a set from $\text{MinFRSups}(\text{mod}(\mathcal{B}) \cup \mathbb{M}, \Lambda)$. However, for some classes of models, such a least finitely representable superset does not exist, that is, $\text{MinFRSups}(\text{mod}(\mathcal{B}) \cup \mathbb{M}, \Lambda)$ is empty. In such cases, therefore, finite-temperance cannot be satisfied, which implies in the inexistence of rational reception operators. Classes of models in which such least finitely representable supersets exist are called reception-compatible.

Definition 7. A class $\mathcal{C}$ of models on a satisfaction system Λ , is reception-compatible iff for all $\mathbb{M} \in \mathcal{C}$ and base $\mathcal{B} \in \wp_f(\mathcal{L})$, $\text{MinFRSups}(\text{mod}(\mathcal{B}) \cup \mathbb{M}, \Lambda) \neq \emptyset$.

Several classes of all models are not reception-compatible, as we show in Sec. 4. For instance, in the DL $\mathcal{EL}_\perp$, the class of all its models is not reception compatible, as Ex. 8 illustrates.

Example 8. Consider the concept $\perp$ in $\mathcal{EL}_\perp$ and a signature with $\mathcal{N}_R = \{r\}$. Suppose $\mathbb{M}^+ = \{(M^\infty, 0)\}$ (see def. in Sec. 2). We would like to include an interpretation pointed at an infinite chain but we are not asking the finite chains M^n to be included. There is no $\mathcal{EL}_\perp$ concept that best represents adding this set to $\perp$. Given a concept $\exists r^n. \top$ one can always create another concept $\exists r^{n+1}. \top$ that includes $\mathbb{M}^+$ but has strictly less models: $\text{mod}(\exists r. \top) \supset \text{mod}(\exists r^2. \top) \supset \dots \supset \text{mod}(\exists r^n. \top) \supset \dots$.

Ex. 9 illustrates a reception operation.

Example 9. (continued from Ex. 5). Recall we want the reception of $\exists r^3. \top$ with $(\mathcal{I}_1, d_1)$. Both $\exists r^2. \top$ and $\exists r. \top$ are more general than $\exists r^3. \top$. From these, the closest that contains $(\mathcal{I}_1, d_1)$ is $\exists r. \top$. So, $\text{rcp}(\{\exists r^3. \top\}, \{(\mathcal{I}_1, d_1)\}) = \{\exists r. \top\}$.

On reception-compatible classes of models, Guimarães, Ozaki, and Ribeiro (2023) have proposed the family of maxi-choice reception operators, which chooses one finitely representable base from MinFRSups . Maxi-choice reception operators coincide with all rational reception operators.

3.2 Eviction

Eviction operators, denoted by evc , are governed by the following postulates:

(success) $\mathbb{M} \cap \text{mod}(\text{evc}(\mathcal{B}, \mathbb{M})) = \emptyset$.

(inclusion) $\text{mod}(\text{evc}(\mathcal{B}, \mathbb{M})) \subseteq \text{mod}(\mathcal{B})$.

(finite retainment) $\mathbb{M}' \notin \text{FR}(\Lambda)$, if $\text{mod}(\text{evc}(\mathcal{B}, \mathbb{M})) \subset \mathbb{M}'$ and $\mathbb{M}' \subseteq \text{mod}(\mathcal{B}) \setminus \mathbb{M}$.

(uniformity) $\text{mod}(\text{evc}(\mathcal{B}, \mathbb{M})) = \text{mod}(\text{evc}(\mathcal{B}', \mathbb{M}'))$, if $\text{MaxFRSups}(\text{mod}(\mathcal{B}) \setminus \mathbb{M}, \Lambda) = \text{MaxFRSups}(\text{mod}(\mathcal{B}') \setminus \mathbb{M}', \Lambda)$

Success ensures that each model from $\mathbb{M}$ must be relinquished. As the purpose of eviction is to remove models, *inclusion* ensures that no models are added. *Uniformity*, as for reception, guarantees that eviction is neither syntax sensitive nor sensitive to model structure. *Finite retainment* ensures that the removal of extra models is minimized, retracting only models that contribute to reaching a finite base. An *eviction operator* on a class $\mathcal{C}$ of models is a model change operator $\text{evc} : \wp_f(\mathcal{L}) \times \mathcal{C} \rightarrow \wp_f(\mathcal{L})$ that satisfies *success*.

An eviction operator satisfying all rationality postulates is called *rational*. In the best scenario, to evict a set of models

$\mathbb{M}$ from a finite base $\mathcal{B}$, one would simply identify a base for the set $\text{mod}(\mathcal{B}) \setminus \mathbb{M}$. However, as not all sets of models are finitely representable, some further models must be minimally removed in favour of finiteness, as finite-retainment demands. This corresponds to picking a greatest finitely representable subset from $\text{mod}(\mathcal{B}) \setminus \mathbb{M}$, that is, picking a set from $\text{MaxFRSups}(\text{mod}(\mathcal{B}) \setminus \mathbb{M}, \Lambda)$. This set, in general, is not a singleton, and a choice must be made among the most plausible candidates. The choice, as for reception, is realised by a choice function. Similarly to reception, depending on the underlying class $\mathcal{C}$ of models, $\text{MaxFRSups}(\text{mod}(\mathcal{B}) \setminus \mathbb{M}, \Lambda)$ might be empty. In such a case, finite retainment cannot be satisfied, which means that in such classes rational eviction operators do not exist. Classes of models in which the greatest finitely representable subsets exist are called eviction-compatible.

Definition 10. A class $\mathcal{C}$ of models, on a satisfaction system Λ , is eviction-compatible iff for all $\mathbb{M} \in \mathcal{C}$ and base $\mathcal{B} \in \wp_f(\mathcal{L})$, $\text{MaxFRSups}(\text{mod}(\mathcal{B}) \setminus \mathbb{M}, \Lambda) \neq \emptyset$.

On eviction-compatible classes of models, Guimarães, Ozaki, and Ribeiro (2023) have proposed the family of maxi-choice eviction operators, which chooses one finitely representable base from FR. Maxi-choice reception operators are characterised by all rationality postulates of reception.

4 Eviction and Reception: DL Concepts

Here we investigate eviction and reception on DL concepts, focusing on the prototypical DLs $\mathcal{ALC}$ and $\mathcal{EL}$. In the following, we denote by $\Lambda(\mathcal{EL}_{\perp \text{concepts}})$ and $\Lambda(\mathcal{ALC}_{\text{concepts}})$ the satisfaction systems for $\mathcal{EL}_\perp$ and $\mathcal{ALC}$ concepts with pointed interpretations as models. Table 1 summarises our results.

Sat. System	Eviction	Reception
$\Lambda(\mathcal{EL}_{\perp \text{con.}})$	yes (Thm. 11)	no (Thm. 13)
$\Lambda(\mathcal{EL}_{\perp \text{con.}})^\dagger$	yes (Thm. 11)	yes (Thm. 14)
$\Lambda(\mathcal{ALC}_{\text{con.}})$	no (Thm. 12)	no (Thm. 13)
$\Lambda(\mathcal{ALC}_{\text{con.}})^\ddagger$	yes (Thm. 16)	yes (Thm. 16)

Table 1: Eviction and reception-compatibility for DL concepts. $\dagger$ is for the case pointed interpretations can only be tree-shaped and $\ddagger$ is for the case they can only be tree-shaped, over finite signatures, and sets of models can only be finite.

Theorem 11. $\Lambda(\mathcal{EL}_{\perp \text{concepts}})$ is eviction-compatible.

Theorem 11 does not hold for $\mathcal{EL}$ (without $\perp$) as any language that cannot express inconsistencies is not eviction-compatible (Guimarães, Ozaki, and Ribeiro 2023). We also do not have eviction-compatibility for $\mathcal{ALC}$.

Theorem 12. $\Lambda(\mathcal{ALC}_{\text{concepts}})$ is not eviction-compatible.

The next theorem establishes that reception-compatibility neither holds for $\mathcal{EL}_\perp$ nor for $\mathcal{ALC}$ concepts.

Theorem 13. $\Lambda(\mathcal{EL}_{\perp \text{concepts}})$ and $\Lambda(\mathcal{ALC}_{\text{concepts}})$ are not reception-compatible. This holds even if we restrict to the class of (possibly infinite) sets of (possibly infinite) tree-shaped pointed interpretations or if we restrict to the class of sets of pointed interpretations over a unique finite signature.

We now concentrate on finding a restricted class of concepts where reception-compatibility holds.

Theorem 14. $\Lambda(\mathcal{EL}_{\perp \text{concepts}})$ is reception-compatible in the class of (possibly infinite) sets of finite tree-shaped pointed interpretations (over a possibly infinite signature).

Thm. 12 and Thm. 13 compel us to investigate more restricted classes of models to obtain a positive result for $\mathcal{ALC}$. There are two natural ways of restricting this class: restricting to finite sets of finite tree-shaped interpretations or to finite tree-shaped interpretations with finite signature. The next theorem establishes that these restrictions alone are not sufficient $\mathcal{ALC}$ reception-compatibility.

Theorem 15. $\Lambda(\mathcal{ALC}_{\text{concepts}})$ is neither reception-compatible in the class of finite sets of finite tree-shaped pointed interpretations over a (possibly infinite) signature; nor in the class of (possibly infinite) sets of finite tree-shaped pointed interpretations over a (unique) finite signature.

Theorem 16. $\Lambda(\mathcal{ALC}_{\text{concepts}})$ is reception-compatible and eviction-compatible in the class of finite sets of finite tree-shaped pointed interpretations over any finite signature.

5 Model Revision

In this section, we introduce a new kind of model change operation, which we call *model revision*. Model revision incorporates a set of models while also guaranteeing that another set of models is removed. For instance, in Ex. 3, $(\mathcal{I}_4, d'')$ had to be removed, while $(\mathcal{I}_5, d'')$ had to be added. Revision cannot be defined by assembling reception and eviction, as reception can add models required to be evicted and vice versa. For revision, we consider change operators defined on a class $\mathcal{C} \subseteq \wp(\mathfrak{M}) \times \wp(\mathfrak{M})$ of pairs of models. A class of pairs of models is called a *binary class*.

Example 17. Let $\mathcal{B} = \{\exists r.\top\}$ be an $\mathcal{EL}_{\perp}$ concept on the signature $\mathcal{N}_{\mathcal{C}} = \{A\}$ and $\mathcal{N}_{\mathcal{R}} = \{r\}$. Let $(\mathcal{I}_1, d_1)$ and $(\mathcal{I}_2, d_1)$ be pointed models with $\Delta^{\mathcal{I}_1} = \{d_1, d_2\}$, $\Delta^{\mathcal{I}_2} = \{d_1, d_2, d_3\}$ and

$$\begin{aligned} \mathcal{I}_1 : & \boxed{A^{\mathcal{I}_1} = \{d_2\}, r^{\mathcal{I}_1} = \{(d_1, d_2)\}} \\ \mathcal{I}_2 : & \boxed{A^{\mathcal{I}_2} = \{d_1\}, r^{\mathcal{I}_2} = \{(d_1, d_2), (d_2, d_3)\}} \end{aligned}$$

We want to revise $\mathcal{B}$ with $(\{(\mathcal{I}_1, d_1)\}, \{(\mathcal{I}_2, d_1)\})$, that is, receive $(\mathcal{I}_1, d_1)$ and evict $(\mathcal{I}_2, d_1)$. Combining rational eviction with rational reception, in any order, is not strong enough to achieve revision. A rational eviction of $\mathcal{B}$ with $(\mathcal{I}_2, d_1)$ is $\mathcal{B}' = \{\exists r^3.\top\}$. However, incorporating $(\mathcal{I}_1, d_1)$ to it yields the base $\{\exists r.\top\}$ again, which contains $(\mathcal{I}_2, d_1)$. On the other hand, reception of $\mathcal{B}$ with $(\mathcal{I}_1, d_1)$ does not change $\mathcal{B}$, as $(\mathcal{I}_1, d_1)$ is a model of $\exists r.\top$. Eviction of $\mathcal{B}$ with $(\mathcal{I}_2, d_1)$ gives $\exists r^3.\top$, which does not contain $(\mathcal{I}_1, d_1)$.

Definition 18. A revision operator, on a binary class $\mathcal{C}$, is function $\text{rev} : \wp_{\text{f}}(\mathcal{L}) \times \mathcal{C} \rightarrow \wp_{\text{f}}(\mathcal{L})$ which satisfies the postulate

success $\mathbb{M}^- \cap \text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-)) = \emptyset$ and $\mathbb{M}^+ \subseteq \text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-))$.

For clarity, we denote $\mathbb{M}^-$ as the set of models to be removed, while $\mathbb{M}^+$ denotes the set of models to be added.

Success guarantees that all models in $\mathbb{M}^-$ are removed while all models in $\mathbb{M}^+$ are added. Clearly, one cannot demand to add and remove the same model. Success cannot be satisfied for $(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-)$, if $\mathbb{M}^+$ and $\mathbb{M}^-$ are not disjoint. Unfortunately, the possibility of success goes beyond identifying whether or not $\mathbb{M}^+$ and $\mathbb{M}^-$ are disjoint. It depends on the logic's underlying satisfaction system. For example, for $\mathcal{ALC}$ concepts, models are closed under bisimulation (Goranko and Otto 2007; Blackburn, van Benthem, and Wolter 2007), which means that some distinct but bisimilar pointed interpretations $(\mathcal{I}_1, d_1)$ and $(\mathcal{I}_2, d_2)$ satisfy precisely the same formulae. Thus, a revision that demands incorporation of $(\mathcal{I}_1, d_1)$ and removal of $(\mathcal{I}_2, d_2)$ cannot occur. For success, revision must at least be defined on classes of pairs of models in which incorporating $\mathbb{M}^+$ does not conflict with eliminating $\mathbb{M}^-$. We call such classes *revision-realisable*.

Definition 19. A binary class of models $\mathcal{C}$ is revision-realisable iff for all $(\mathbb{M}^+, \mathbb{M}^-) \in \mathcal{C}$, there is a finitely representable set $\mathbb{M}$ of models such that $\mathbb{M}^+ \subseteq \mathbb{M}$ and $\mathbb{M}^- \cap \mathbb{M} = \emptyset$.

For conciseness, unless otherwise explicitly stated, we assume that all binary classes of models are revision-realisable. Success alone is not enough to bring rationality to revision. We introduce other rationality postulates to capture the minimal change principle for revision. We start with

vacuous-expansion: if $\mathbb{M}^+ \subseteq \text{mod}(\mathcal{B})$ then $\text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-)) \subseteq \text{mod}(\mathcal{B})$,

vacuous-removal: if $\mathbb{M}^- \cap \text{mod}(\mathcal{B}) = \emptyset$ then $\text{mod}(\mathcal{B}) \subseteq \text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-))$.

If all models of $\mathbb{M}^+$ are models of $\mathcal{B}$, one should only evict $\mathbb{M}^-$ (*vacuous expansion*). On the other hand, if none of the models in $\mathbb{M}^-$ satisfy $\mathcal{B}$ (*vacuous removal*), all we need to do is to add models. If all models in $\mathbb{M}^+$ satisfy $\mathcal{B}$, and none of $\mathbb{M}^-$ violate $\mathcal{B}$, the base $\mathcal{B}$ should be left untouched. We call this *lethargy*.

lethargy: if $\mathbb{M}^+ \subseteq \text{mod}(\mathcal{B})$ and $\mathbb{M}^- \cap \text{mod}(\mathcal{B}) = \emptyset$, then $\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-) = \text{mod}(\mathcal{B})$.

Proposition 20. If a revision operator satisfies *vacuous-expansion* and *vacuous-removal* then it satisfies *lethargy*.

These three postulates capture the most fundamental features of revision. Yet, such postulates are not enough, as they allow for drastic removal and addition of models. For example, if some of the models in $\mathbb{M}^+$ does not satisfy $\mathcal{B}$ and some model in $\mathbb{M}^-$ violates $\mathcal{B}$ then **vacuous-expansion** and **vacuous-removal** allow the removal of all models of $\mathcal{B}$. However, ideally, changes should be minimised.

In the case that the trivial removal of $\mathbb{M}^-$ and the trivial addition of $\mathbb{M}^+$ reach a finitely representable set $\mathbb{M}$, the revision should correspond to $\mathbb{M}$ as finiteness is trivially obtained. However, if finiteness is not reached in this way, then one must remove and add some extra interpretations in favour of finiteness. Such additions and removals must be minimised. Such a minimal change principle is conceptualised in the form of the **circumspection** postulate.

circumspection: if $X^+, X^- \in \mathcal{C}$ are disjoint sets and conditions (1) to (3) below are jointly satisfied, then condition (4) is satisfied:

1. $X^- \subseteq (\text{mod}(\mathcal{B}) \setminus \text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-)))$, and $\mathbb{M}^- \cap \text{mod}(\mathcal{B}) \subseteq X^-$
2. $\mathbb{M}^+ \setminus \text{mod}(\mathcal{B}) \subseteq X^+ \subseteq (\text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-)) \setminus \text{mod}(\mathcal{B}))$ and
3. $((\text{mod}(\mathcal{B}) \setminus X^-) \cup X^+) \neq \text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-))$
4. $((\text{mod}(\mathcal{B}) \setminus X^-) \cup X^+) \notin \text{FR}(\Lambda)$.

In **circumspection** above, the set X^- (condition 1) denotes the extra interpretations to be removed, while X^+ (condition 2) denotes the extra interpretations to be added during revision. These extra additions and removals can only occur in favour of finiteness, and all of them must be necessary to achieve finiteness. Therefore, every smaller combination of removals or additions to form a revision candidate (condition 3) does not reach finiteness (condition 4). The postulate **circumspection** captures **lethargy**.

Proposition 21. *If a revision operator satisfies **circumspection** then it satisfies **lethargy**.*

Defining operators capable of satisfying principles of minimal change has been proved a challenge in the field of *belief change*. In some logics, operators satisfying minimal change principles cannot even be defined (Flouris 2006; Ribeiro, Nayak, and Wassermann 2018; Guimarães, Ozaki, and Ribeiro 2023). This occurs due to the strong semantics and properties of the logics. Guimarães, Ozaki, and Ribeiro (2023) have shown that in some logics satisfaction systems, eviction and reception do not exist. For revision, this would not be different. We shall direct the effort of defining revision operators to classes of models that are compatible with such rationality postulates.

Definition 22. *A binary class $\mathcal{C}$ is revision-compatible iff there is a revision operator on $\mathcal{C}$ satisfying **success**, **vacuous-expansion**, **vacuous-removal** and **circumspection**.*

Given that we are working on a binary class of models that is revision-compatible, we would like to know how to construct a revision operator satisfying all the rationality postulates presented so far. We will frame the precise class of operators that satisfy such postulates. For this, we need to define some auxiliary tools. Let

$$\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-) = \{Y \in \text{FR}(\Lambda) \mid \mathbb{M}^+ \subseteq Y \text{ and } \mathbb{M}^- \cap Y = \emptyset\},$$

be the set which contains exactly all finite bases satisfied by all models in a given set $\mathbb{M}^+$ but violated by all models in $\mathbb{M}^-$. One can regard $\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-)$ as the set of all potential candidates to revise a base with the pair $(\mathbb{M}^+, \mathbb{M}^-)$.

Not all sets in $\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-)$, however, are suitable to revise a given base $\mathcal{B}$, as some of them might add or remove more than allowed by **circumspection**. If we could “measure” the changes incurred on $\text{mod}(\mathcal{B})$ to achieve a finite representable set of models, then we can just select the finite representable sets with the minimal incurred changes. The symmetric difference between two sets A and B provide exactly the changes necessary to turn one set into another. To turn a set A into a set B , we need only to add the elements of B that are not in A , and remove the elements in A that are not in B . We can, therefore, use the symmetric difference to “measure” changes between sets of interpretations, and

choose those that minimise the changes. On each set of interpretations $\mathbb{M}$, we define the relation $\leq_{\mathbb{M}} \subseteq \wp(\mathfrak{M}) \times \wp(\mathfrak{M})$, such that $\mathbb{M}_1 \leq_{\mathbb{M}} \mathbb{M}_2$ iff $(\mathbb{M} \oplus \mathbb{M}_1) \subseteq (\mathbb{M} \oplus \mathbb{M}_2)$.

Intuitively, $\mathbb{M}_1 \leq_{\mathbb{M}} \mathbb{M}_2$ means that turning $\mathbb{M}$ into $\mathbb{M}_2$ incurs in at least as much change as turning $\mathbb{M}$ into $\mathbb{M}_1$. This means that turning $\mathbb{M}$ into $\mathbb{M}_1$ is equally cheap or cheaper than turning $\mathbb{M}$ to $\mathbb{M}_2$. We can use this closeness relation to revise a base $\mathcal{B}$ by a pair $(\mathbb{M}^+, \mathbb{M}^-)$. As the revision must minimise the changes, we choose, from $\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-)$, one of the closest options to $\text{mod}(\mathcal{B})$ that is, one from $\min_{\leq_{\text{mod}(\mathcal{B})}}(\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-))$. To define a revision function using this strategy, we need the condition that for every base $\mathcal{B}$ and pair $(\mathbb{M}^+, \mathbb{M}^-)$, there is at least one choice on $\min_{\leq}(\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-))$, that is, the symmetric difference indeed minimises the distance from the base to the revision candidates. This condition follows from revision-compatibility.

Theorem 23. *If $\mathcal{C}$ is revision-compatible, then for all finite base $\mathcal{B}$ and $(\mathbb{M}^+, \mathbb{M}^-) \in \mathcal{C}$, $\min_{\leq_{\text{mod}(\mathcal{B})}}(\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-)) \neq \emptyset$.*

We get revision operators from symmetric difference.

Definition 24. *A naïve relational revision operator, on a binary class of models $\mathcal{C}$, is a map $\text{rev}^\oplus : \wp_f(\mathcal{L}) \times \mathcal{C} \rightarrow \wp_f(\mathcal{L})$ s.t. $\text{mod}(\text{rev}_{\text{sel}}^\oplus(\mathbb{M}^+, \mathbb{M}^-)) \in \min_{\leq_{\text{mod}(\mathcal{B})}}(\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-))$.*

The naïve operator minimises the distance between the models of a base $\mathcal{B}'$ to all bases satisfied by $\mathbb{M}^+$ and violated by all models in $\mathbb{M}^-$. Indeed, the naïve revision operators is strongly connected with **circumspection**.

Theorem 25. *A revision operator satisfies **circumspection** iff it is a naïve relational revision operator.*

Although naïve revision operators capture **circumspection**, they are too weak to satisfy **vacuous-removal** and **vacuous-expansion**, as Ex. 26 illustrates.

Example 26. *Let $\mathcal{B} = (B \sqcap C)$ be an $\mathcal{EL}$ concept and consider the interpretations $\mathcal{I}_i = (\Delta^{\mathcal{I}_i}, \mathcal{I}_i)$, with $i \in \{1, 2, 3, 4\}$, where $\Delta^{\mathcal{I}_i} = \{d\}$, and each $\mathcal{I}_i$ is as follows.*

$A^{\mathcal{I}_1} = \{d\}$	$A^{\mathcal{I}_2} = \emptyset$	$A^{\mathcal{I}_3} = \emptyset$	$A^{\mathcal{I}_4} = \emptyset$
$B^{\mathcal{I}_1} = \emptyset$	$B^{\mathcal{I}_2} = \emptyset$	$B^{\mathcal{I}_3} = \emptyset$	$B^{\mathcal{I}_4} = \{d\}$
$C^{\mathcal{I}_1} = \{d\}$	$C^{\mathcal{I}_2} = \emptyset$	$C^{\mathcal{I}_3} = \{d\}$	$C^{\mathcal{I}_4} = \{d\}$

By definition, $(\mathcal{I}_4, d)$ is a model of $\mathcal{B} = (B \sqcap C)$. Assume we want to revise $\mathcal{B}$ with $(\{(\mathcal{I}_1, d)\}, \{(\mathcal{I}_2, d)\})$, that is, accommodate $(\mathcal{I}_1, d)$ while relinquishing $(\mathcal{I}_2, d)$. Since $d \notin (B \sqcap C)^{\mathcal{I}_2}$, according to **vacuous-removal**, we should only add models to $\mathcal{B}$. Adding only $(\mathcal{I}_1, d)$ is not possible, because $(\mathcal{I}_4, d)$ is a model of $\mathcal{B}$ and every $\mathcal{EL}$ concept satisfied by both $(\mathcal{I}_4, d)$ and $(\mathcal{I}_1, d)$ is also satisfied by $(\mathcal{I}_3, d)$. So, every revision satisfying **vacuous-removal** contains the set $\{(\mathcal{I}_1, d), (\mathcal{I}_3, d), (\mathcal{I}_4, d)\}$. This corresponds to revising $\mathcal{B} = (B \sqcap C)$ to $\mathcal{C}$. If we want to avoid adding $(\mathcal{I}_3, d)$, we could remove $(\mathcal{I}_4, d)$, violating **vacuous-removal**, and staying only with $(\mathcal{I}_1, d)$. So, a naïve revision operator can output $(A \sqcap C)$ as a solution for revising $(B \sqcap C)$.

In Ex. 26, **vacuous-removal** is violated, as the operator removes further models from $\mathcal{B}$ when $\mathbb{M}^-$ does not violate $\mathcal{B}$. Also, when $\mathbb{M}^+$ satisfies $\mathcal{B}$, the naïve operator allows adding further models. We strengthen the naïve operator.

Definition 27. A symmetric-differential revision function on a binary class $\mathcal{C}$ of models, is a function $\text{rev} : \wp_f(\mathcal{L}) \times \mathcal{C} \rightarrow \wp_f(\mathcal{L})$, such that for all $(\mathbb{M}^+, \mathbb{M}^-) \in \mathcal{C}$, $\text{mod}(\text{rev}(\mathcal{B}, \mathbb{M}^+, \mathbb{M}^-)) = \mathbb{M}$ where,

(i) if $\mathbb{M}^+ \subseteq \text{mod}(\mathcal{B})$, then

$$\mathbb{M} \in \min_{\leq \text{mod}(\mathcal{B})} (\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^- \cup (\mathfrak{M} \setminus \text{mod}(\mathcal{B}))))$$

(ii) if $\mathbb{M}^+ \not\subseteq \text{mod}(\mathcal{B})$, but $\mathbb{M}^- \cap \text{mod}(\mathcal{B}) = \emptyset$, then

$$\mathbb{M} \in \min_{\leq \text{mod}(\mathcal{B})} (\chi_\Lambda(\mathbb{M}^+ \cup \text{mod}(\mathcal{B}), \mathbb{M}^-))$$

(iii) otherwise, $\mathbb{M} \in \min_{\leq \text{mod}(\mathcal{B})} (\chi_\Lambda(\mathbb{M}^+, \mathbb{M}^-))$.

For case (i), when $\mathbb{M}^+$ satisfies the base, according to *vacuous-expansion*, no interpretation can be incorporated. This corresponds to enforcing all counter models of $\mathcal{B}$ to be removed jointly with $\mathbb{M}^-$. For case (ii), analogous to case (i), whenever all models of $\mathbb{M}^-$ violate $\mathcal{B}$, as per *vacuous-removal*, no model of $\mathcal{B}$ should be removed. This corresponds to enforcing all models of $\text{mod}(\mathcal{B})$ to be incorporated jointly with $\mathbb{M}^+$. As for case (iii), the cases (i) and (ii) give enough protection to *vacuous-removal* and *vacuous-expansion*. Hence, in any other case, we can just perform a naïve revision. The symmetric-differential revision operators are characterised by the rationality postulates of revision.

Theorem 28. A revision operator rev satisfies **success**, **vacuous-expansion**, **vacuous-removal** and **circumspection** iff it is a symmetric differential revision operator.

Revision-compatibility is tightly connected to both reception-compatibility and eviction-compatibility. In the case that $\mathbb{M}^-$ is empty, revising a base with $(\mathbb{M}^+, \mathbb{M}^-)$ intuitively corresponds to performing a reception, as **vacuous-removal** would forbid removal of interpretations. Analogously, if $\mathbb{M}^+$ is empty, then revising with $(\mathbb{M}^+, \mathbb{M}^-)$, with a non-empty $\mathbb{M}^-$, would correspond to evicting $\mathbb{M}^-$, as **vacuous-expansion** would forbid adding interpretations. From this perspective, we can trace an important connection between revision-compatibility with eviction-compatibility and reception-compatibility. To establish this connection, the underlying class of models must cover the cases that we can perform revision of the kind $(\mathbb{M}^+, \emptyset)$ and $(\emptyset, \mathbb{M}^-)$, that is, cover the possibility of solely adding or removing interpretations. We call such classes decomposable classes of models.

Definition 29. A binary class of models $\mathcal{C}$ is decomposable iff for all $(\mathbb{M}^+, \mathbb{M}^-) \in \mathcal{C}$, $(\mathbb{M}^+, \emptyset) \in \mathcal{C}$ and $(\emptyset, \mathbb{M}^-) \in \mathcal{C}$.

Given a binary class of models $\mathcal{C}$, let $\mathcal{C}^+ = \{\mathbb{M}^+ \in \mathfrak{M} \mid (\mathbb{M}^+, \mathbb{M}^-) \in \mathcal{C}\}$, and $\mathcal{C}^- = \{\mathbb{M}^- \in \mathfrak{M} \mid (\mathbb{M}^+, \mathbb{M}^-) \in \mathcal{C}\}$. The set $\mathcal{C}^+$ is the greatest subclass of $\mathcal{C}$ with reception candidates, whereas $\mathcal{C}^-$ is the largest subclass with eviction candidates. Every rational revision operator induces an eviction and a reception operator. Consequently, compatibility of revision implies compatibility with both reception and eviction.

Theorem 30. Let $\mathcal{C}$ be a decomposable revision-compatible binary class of models. The following hold for $\mathcal{C}$.

1. $\mathcal{C}^+$ is reception-compatible and $\mathcal{C}^-$ is eviction-compatible.
2. If evc is a eviction operator on $\mathcal{C}^-$, then there is a revision operator rev on $\mathcal{C}$ satisfying all rationality postulates such that $\text{evc}(\mathcal{B}, \mathbb{M}^-) = \text{rev}(\mathcal{B}, \emptyset, \mathbb{M}^-)$.

3. If rcp is a reception operator on $\mathcal{C}^+$, then there is a revision operator rev on $\mathcal{C}$ satisfying all rationality postulates such that $\text{rcp}(\mathcal{B}, \mathbb{M}^+) = \text{rev}(\mathcal{B}, \mathbb{M}^+, \emptyset)$.

From Thm. 30, eviction and reception are special cases of revision, whereas revision can only be performed in classes of models compatible with both reception and eviction. This connection between revision with eviction and reception allows to translate (in)compatibility results from eviction and reception to revision, as we see in Sec. 6.

6 Revision: DL Concepts

Here, we briefly consider the revision of DL concepts. It follows from Thm. 30 and the results in Table 1 that neither $\Lambda(\mathcal{EL}_{\perp \text{concepts}})$ nor $\Lambda(\mathcal{ALC}_{\text{concepts}})$ are, in general, revision-compatible. We establish that these satisfaction systems are also not revision-compatible when we restrict to finite tree-shaped pointed interpretations.

Theorem 31. $\Lambda(\mathcal{EL}_{\perp \text{concepts}})$ and $\Lambda(\mathcal{ALC}_{\text{concepts}})$ are not revision-compatible in the binary class of finite tree-shaped pointed interpretations. This also holds if we restrict to finite sets of models and if we restrict to a finite signature.

So we restrict the binary class of models that we consider even further. We consider the binary class of finite tree-shaped pointed interpretations for finite sets of models *union their closure under bisimulation* over a finite signature. We argue that this class is revision-compatible. We say that two sets of pointed interpretations are *bisimulation disjoint* iff there is no pointed interpretation in one of the sets that is bisimilar to a pointed interpretation in the other set.

Theorem 32. $\Lambda(\mathcal{ALC}_{\text{concepts}})$ is revision-compatible in the binary class of sets of pointed interpretations which are the closure under bisimulation of finite sets of finite tree-shaped pointed interpretations over a (unique) finite signature.

Regarding the case of $\mathcal{EL}_{\perp}$, we do not have the same expressivity we have in $\mathcal{ALC}$ for restricting the models that are satisfied by a concept. We leave it as an open question. From the positive results for eviction and reception for certain classes of models for $\mathcal{EL}_{\perp}$ and $\mathcal{ALC}$ in Table 1, we obtain the existence of the revision operators in Thm. 30.

7 Conclusion

We investigated eviction and reception for DL concepts, establishing various results considering different classes of models. We find classes of models where we obtain eviction and reception compatibility for both $\mathcal{ALC}$ and $\mathcal{EL}_{\perp}$ concepts. It turns out that the class of models where we obtain positive results for $\mathcal{ALC}$ is much more restricted than the class for $\mathcal{EL}_{\perp}$. We also introduce the notion of model revision and relate various postulates with the revision operation. Revision cannot be seen as a mere combination of eviction and reception, which is evidenced by negative results for DL concepts. As future work, it would be interesting to investigate model eviction, reception, and revision of DL concepts in a more practical setting, expand our work to more expressive DLs.

Acknowledgments

Ozaki is supported by the Research Council of Norway, project (316022, 322480). This work was also supported by the Research Council of Norway, Integreat - Norwegian Centre for knowledge-driven machine learning (332645).

References

- Agi, K.; Wolter, F.; Zakharyashev, M.; and Gabbay, D. 2003. *Many-dimensional modal logics : theory and applications*. Amsterdam Boston: Elsevier North Holland. ISBN 0444508260.
- Aiguier, M.; Atif, J.; Bloch, I.; and Hudelot, C. 2018. Belief revision, minimal change and relaxation: A general framework based on satisfaction systems, and applications to description logics. *Artificial Intelligence*, 256: 160–180.
- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 50(2): 510–530.
- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press. ISBN 978-0521695428.
- Baader, F.; Sertkaya, B.; and Turhan, A. 2007. Computing the least common subsumer w.r.t. a background terminology. *J. Appl. Log.*, 5(3): 392–420.
- Baltag, A.; Gierasimczuk, N.; Özgün, A.; Sandoval, A. L. V.; and Smets, S. 2019. A dynamic logic for learning theory. *J. Log. Algebraic Methods Program.*, 109.
- Baltag, A.; Gierasimczuk, N.; and Smets, S. 2019. Truth-Tracking by Belief Revision. *Stud Logica*, 107(5): 917–947.
- Blackburn, P.; van Benthem, J. F. A. K.; and Wolter, F., eds. 2007. *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*. North-Holland. ISBN 978-0-444-51690-9.
- De Raedt, L. 1997. Logical settings for concept-learning. *Artificial Intelligence*, 95(1): 187–201.
- Delgrande, J. P.; Peppas, P.; and Woltran, S. 2018. General Belief Revision. *J. ACM*, 65(5): 29:1–29:34.
- Dummett, M. A. E.; and Lemmon, E. J. 1959. Modal Logics Between S 4 and S 5. *Mathematical Logic Quarterly*, 5(14-24): 250–264.
- Fanizzi, N.; d’Amato, C.; and Esposito, F. 2008. DL-FOIL Concept Learning in Description Logics. In Zelezný, F.; and Lavrac, N., eds., *ILP*, volume 5194 of *Lecture Notes in Computer Science*, 107–121. Springer.
- Flouris, G. 2006. On belief change in ontology evolution. *AI Communications*, 19(4): 395–397.
- Funk, M.; Jung, J. C.; and Lutz, C. 2021. Actively Learning Concepts and Conjunctive Queries under ELr-Ontologies. In Zhou, Z., ed., *IJCAI*, 1887–1893. ijcai.org.
- Funk, M.; Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2019. Learning Description Logic Concepts: When can Positive and Negative Examples be Separated? In Kraus, S., ed., *IJCAI*, 1682–1688. ijcai.org.
- Gärdenfors, P. 1988. *Knowledge in flux: Modeling the dynamics of epistemic states*. The MIT press.
- Goranko, V.; and Otto, M. 2007. Model theory of modal logic. In Blackburn, P.; van Benthem, J. F. A. K.; and Wolter, F., eds., *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*, 249–329. North-Holland.
- Guimarães, R.; Ozaki, A.; and Ribeiro, J. S. 2023. Finite Based Contraction and Expansion via Models. In *AAAI*, 6389–6397.
- Hansson, S. O. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Applied Logic Series. Kluwer Academic Publishers.
- Iannone, L.; Palmisano, I.; and Fanizzi, N. 2007. An algorithm based on counterfactuals for concept learning in the Semantic Web. *Appl. Intell.*, 26(2): 139–159.
- Klarman, S.; and Britz, K. 2015. Ontology Learning from Interpretations in Lightweight Description Logics. In Inoue, K.; Ohwada, H.; and Yamamoto, A., eds., *ILP*, volume 9575 of *Lecture Notes in Computer Science*, 76–90. Springer.
- Konev, B.; Lutz, C.; Wolter, F.; and Zakharyashev, M. 2016. Conservative Rewritability of Description Logic TBoxes. In Kambhampati, S., ed., *IJCAI*, 1153–1159. IJCAI Press.
- Lehmann, J. 2009. DL-Learner: Learning Concepts in Description Logics. *J. Mach. Learn. Res.*, 10: 2639–2642.
- Lehmann, J.; and Haase, C. 2009. Ideal Downward Refinement in the *EL* Description Logic. In Raedt, L. D., ed., *ILP*, volume 5989 of *Lecture Notes in Computer Science*, 73–87. Springer.
- Lehmann, J.; and Hitzler, P. 2010. Concept learning in description logics using refinement operators. *Mach. Learn.*, 78(1-2): 203–250.
- Ozaki, A.; and Troquard, N. 2019. Learning Ontologies with Epistemic Reasoning: The *E*/*L* Case. In Calimeri, F.; Leone, N.; and Manna, M., eds., *JELIA*, volume 11468 of *Lecture Notes in Computer Science*, 418–433. Springer.
- Ribeiro, J. S.; Nayak, A.; and Wassermann, R. 2018. Towards Belief Contraction without Compactness. In *KR 2018*, 287–296. AAAI Press.
- Schwind, N.; Inoue, K.; Konieczny, S.; and Marquis, P. 2025. Iterated Belief Change as Learning. In *IJCAI*, 4669–4677. ijcai.org.
- ten Cate, B.; Koudijs, R.; and Ozaki, A. 2024. On the Power and Limitations of Examples for Description Logic Concepts. In *IJCAI*, 3567–3575. ijcai.org.