Towards a Semantic Construction for Belief Base Contraction: Partial Meet vs Smooth Kernel (Preliminary Report) *

Jandson S. Ribeiro

Universiy of Hagen, Germany

Abstract

We introduce novel classes of fully rational contraction operators for belief bases. These operators are founded on a plausibility relation on models, called tracks, that allow distinguishing between suitable and unsuitable models. We obtain two main representation theorems: the first one semantically characterizes the class of partial-meet operators, which are related to the rationality postulate of relevance; while the second one semantically characterizes the class of smooth kernel contraction operators, which are related to the postulates of core-retainment and relative closure. For the supplementary postulates (conjunction and intersection), we strengthen such operators by imposing the mirroring condition on the track relations. We consider logics that are both Tarskian and compact.

Keywords

Belief base, models, AGM, contraction

1. Introduction

The field of *Belief Change* [2, 3, 4] studies how an agent should rationally modify its corpus of beliefs in response to incoming pieces of information. The two most important kinds of change are: contraction, which relinquishes undesirable/obsolete information; and revision, which accommodates new information with the caveat of keeping the corpus of beliefs consistent. Each of these kinds of changes is governed by sets of rationality postulates, split into basic and supplementary rationality postulates, which prescribe adequate behaviours of change. Such rationality postulates are motivated by the principle of minimal change: in response to a piece of information, say α, an agent should remove only beliefs that either conflict with α (in the case of revision), or that contribute to entail α (in case of contraction).

Several classes of belief change operators were proposed that abide by such rationality postulates, called rational belief change operators (see [4], for a list). These classes of operators can be split into two main kinds: syntactic operators and semantic operators. Operators belonging to the first kind select sentences from the language, while operators of the second kind select models. Examples of syntactic operators are partial meet operators [2] and smooth kernel operators [5], while Grove's system of spheres [6, 3] and the faithful pre-orders of Katsuno and Mendelzon [7] are the main frameworks for constructing semantic operators. In the most fundamental case, when an agent's corpus of beliefs is represented as a logically closed set of sentences, called a theory, all these classes of operators are characterised by the rationality postulates of contraction/revision.

Theories, however, are very restrictive, as they do not distinguish between explicit and implicit beliefs. One can achieve this distinction by dropping the logical closure requirement, and simply representing an agent's corpus of beliefs as a set of sentences, called a *belief base* [4]. For bases, however, very few belief change operators are capable of satisfying the rationality postulates of belief change. The two foremost classes of syntactic operators are smooth kernel contraction and partial-meet. On theories, these two classes are equivalent, whereas on bases only partial meet remains rational for belief bases [4, 5]. On bases, smooth kernel contraction corresponds to a more permissive version of contraction. As a result, research on belief base change has focused on partial meet operators or other similar syntactic operators [4, 8]. This poses a severe limitation in advancing belief base change, as syntactic operators are highly dependent on the assumptions made about the underlying logic used to represent an agent's knowledge, as for instance, imposing that the language is closed under classical negation [9]. By devising belief change operators via models, such conditions upon the language of the logics can be easily waived.

In this work, we devise three novel classes of semantic operators for belief base contraction. Our approach consists in imposing a pre-order, called a track, upon the models of the logics. A track indicates the most plausible models, which in turn are selected to perform a contrac-

21st International Workshop on Nonmonotonic Reasoning, September 2-4, 2023, Rhodes, Greece

* A preliminary version was presented at the FCR 22 workshop [1]

jandson.ribeiro@fernuni-hagen.de (J. S. Ribeiro)
https://jandsonribeiro.github.io/home/ (J. S. Ribeiro)
0000-0003-1614-6808 (J. S. Ribeiro)

tion. We call such operators that follow this strategy tracked contraction operators. We show a representation theorem between the basic rationality postulates of belief base contraction and such a novel class of contraction operators. Equivalently, the tracked contraction operators correspond to the semantic counterpart of the partial meet operators. We then impose the mirroring condition [10] upon such tracks, and we show that tracks satisfying mirroring induce belief base contraction operators that capture the supplementary postulates of belief contraction. It is worth highlighting that, except for safe contraction [11], the study of the supplementary postulates on belief bases has been neglected. As contraction is a central operation in belief change, our result can be extended to provide semantic operators for other kinds of belief change such as revision.

We also characterize semantically the smooth kernel contraction operators for bases. For this, we explore some properties of the track relations which unveil the permissive behaviour of smooth kernel contraction on models. We then relax the tracked contraction operators in order to capture such behaviour.

Road map: Section 2 introduces some basic notations and definitions that will be used throughout this work. In Section 3, we briefly review belief contraction, including both basic and supplementary rationality postulate of contraction as well as the smooth kernel and partial meet contraction operators. For semantic operators, we review the faithful pre-orders of Katsuno and Mendelzon [7] for revision, and we translate them in terms of belief contraction. We show that such operators, though fully rational for theories, are not rational for belief bases. In Section 4, we introduce two novel classes of contraction operators and the representation theorem connecting tracks and the basic rationality postulates of contraction. In Section 5, we semantically characterize the smooth kernel contraction operators using the track relations. Finally, in Section 6 we conclude the work and discuss some future works. We sketch the proofs of the most important results. The full proofs are available in the appendix at https://jandsonribeiro.github.io/home/appendix/NMR_23_appendix.pdf

2. Notation and Technical Background

The power set of a set A is denoted by $\mathcal{P}(A)$. We treat a logic as a pair $\langle \mathcal{L}, Cn\rangle$, where $\mathcal{L}$ is a language, and $Cn : \mathcal{P}(\mathcal{L}) \to \mathcal{P}(\mathcal{L})$ is a logical consequence operator that indicates all the formulae that are entailed from a set of formulae in $\mathcal{L}$. We limit ourselves to logics whose consequence operator Cn satisfies:

monotonicity: if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$;

inclusion: $A \subseteq Cn(A)$;

idempotency: $Cn(Cn(A)) = Cn(A)$;

compactness: if $\varphi \in Cn(A)$ then there is some finite set $A' \subseteq A$ such that $\varphi \in Cn(A')$.

Consequence operators that satisfy the first three conditions above are called Tarskian. Likewise, consequence operators satisfying the compactness property will be called compact. Sometimes we say that the logic itself is Tarskian or compact. Throughout this work, unless otherwise stated, all the presented results regard logics whose consequence operators are Tarskian and satisfy compactness. A theory is a set of formulae $X \subseteq \mathcal{L}$ such that $X = Cn(X)$.

As we are interested to define semantic operators, we exploit the semantics of the logics. Given a logic $\langle \mathcal{L}, Cn\rangle$ and a set of structures $\mathcal{I}$, an interpretation or a model is an element of $\mathcal{I}$ that gives meaning to the formulae of $\mathcal{L}$; $\mathcal{I}$ is called an *interpretation domain* of that logic, whereas each subset of $\mathcal{I}$ is called an *interpretation set*. For instance, an interpretation domain for the Propositional Logic is the power set of the propositional symbols of the language. A satisfaction relation $\models \subseteq \mathcal{I} \times \mathcal{L}$ is used to indicate on which interpretations a formula is satisfied. If $M \models \alpha$, then we say that M is a model of α. If an interpretation M does not satisfy a formula α, denoted by $M \not\models \alpha$, then we say that M is a counter-model of α. The set of all models of α is given by $[\![\alpha]\!]$, while the set of all counter-models of α is given by $\overline{[\![\alpha]\!]}$.

In Tarskian logics, the consequence operator can be semantically defined as: a formula $\varphi \in Cn(X)$ iff every model that satisfies all formulae in X also satisfies φ [12]. Let $\mathcal{I}$ be an interpretation domain of a logic $\langle \mathcal{L}, Cn\rangle$, and M a model in $\mathcal{I}$. The set of all formulae of $\mathcal{L}$ satisfied by M is the theory $Th(M) = \{\varphi \in \mathcal{L} \mid M \models \varphi\}$. Generalising, given a set of models A, $Th(A) = \{\varphi \mid \forall M \in A, M \models \varphi\}$ is the theory of the formulae satisfied by all models in A. Moreover, given a set $X \subseteq \mathcal{L}$, the set of models that satisfy all formulae in X is $[\![X]\!] = \{M \in \mathcal{I} \mid \forall \varphi \in X, M \models \varphi\}$. For simplicity, given a set of formulae X and a model M, we will write $M \models X$ to mean that M satisfies every formula in X.

Throughout this paper, we will provide examples to support the intuition of the proposed contraction operators. Due to its simplicity, we will use classical propositional logics to construct such examples. Observe, however, that our results are not confined to classical propositional logics. As usual, the formulae of classical propositional logics are Boolean formulae constructed from a set AP of atomic propositional symbols, via the operators of conjunction ($\wedge$), disjunction ($\vee$) and classical negation ($\neg$). The models are subsets of AP, and the satisfaction relation is defined as usual.

A pre-order on a domain $\mathcal{D}$ is binary relation $\leqslant: \mathcal{D}\times\mathcal{D}$ that satisfies transitivity and reflexivity. The minimal elements of a set $A \subseteq \mathcal{D}$ w.r.t a binary relation $\leqslant: \mathcal{D}\times\mathcal{D}$ is $\min_{\leqslant}(A) = \{a \in A \mid \text{ if } b \leqslant a \text{ then } a \leqslant b, \text{ for all } b \in A\}$. We write $a < b$ to denote that $a \leqslant b$ but $b \nleqslant a$. We also write $a \sim b$ as a shorthand for $a \leqslant b$ and $b \leqslant a$.

3. Belief Contraction

We assume that an agent's corpus of beliefs is represented as a belief base, which will be denoted by the letter $\mathcal{K}$. The term belief base has been used in the literature with two main purposes: (i) as a finite representation of an agent's beliefs [13, 14, 15], and (ii) as a more general and expressive approach that distinguishes explicit from implicit beliefs [16, 4]. We follow the latter approach, and therefore a belief base can be infinite.

Let $\mathcal{K}$ be a belief base, a contraction function for $\mathcal{K}$ is a function $\dot{-} : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ that given an unwanted piece of information α, outputs a subset of $\mathcal{K}$ which does not entail α. A contraction function is subject to the following basic rationality postulates [17, 5]:

(success): if $\alpha \notin Cn(\emptyset)$ then $\alpha \notin Cn(\mathcal{K} \dot{-} \alpha)$;

(inclusion): $\mathcal{K} \dot{-} \alpha \subseteq \mathcal{K}$;

(vacuity): if $\alpha \notin Cn(\mathcal{K})$ then $\mathcal{K} \dot{-} \alpha = \mathcal{K}$;

(uniformity): if for all $\mathcal{K}' \subseteq \mathcal{K}$ it holds that $\alpha \in Cn(\mathcal{K}')$ iff $\beta \in Cn(\mathcal{K}')$, then $\mathcal{K} \dot{-} \alpha = \mathcal{K} \dot{-} \beta$;

(core-retainment): if $\beta \in \mathcal{K} \setminus (\mathcal{K} \dot{-} \alpha)$ then there is a $\mathcal{K}' \subseteq K$ s.t $\alpha \notin Cn(\mathcal{K}')$ but $\alpha \in Cn(\mathcal{K}' \cup \{\beta\})$;

(**relative closure**): $\mathcal{K} \cap Cn(\mathcal{K} \dot{-} \alpha) \subseteq \mathcal{K} \dot{-} \alpha$;

(relevance): if $\beta \in \mathcal{K} \setminus (\mathcal{K} \dot{-} \alpha)$ then there is some $\mathcal{K}'$ such that $\mathcal{K} \dot{-} \alpha \subseteq \mathcal{K}' \subseteq \mathcal{K}$, $\alpha \notin Cn(\mathcal{K}')$ but $\alpha \in Cn(\mathcal{K}' \cup \{\beta\})$.

For a discussion on the rationale of these postulates, see [4]. The postulate of uniformity guarantees that contraction is not syntax sensitive: if two formulae, say α and β, are entailed exactly by the same subsets of $\mathcal{K}$ (we say α and β are $\mathcal{K}$-uniform), then α and β must present the same contraction result. We call the set of rationality postulates listed above the basic rationality postulates of contraction. A contraction function that satisfies all the basic rationality postulates above will be dubbed a *rational contraction function*. It is worth highlighting that relevance implies core-retainment. Moreover, in Tarskian logics, relevance also implies relative closure [4].

There are other two postulates, called supplementary postulates [2, 18, 4]:

(***intersection***) $\mathcal{K} \dot{-} \varphi \cap \mathcal{K} \dot{-} \psi \subseteq \mathcal{K} \dot{-} \varphi \wedge \psi$

(***conjunction***) If $\varphi \notin Cn(\mathcal{K} \dot{-} \varphi \wedge \psi)$ then $\mathcal{K} \dot{-} (\varphi \wedge \psi) \subseteq \mathcal{K} \dot{-} \varphi$.

It is important to stress that the study of the supplementary postulates has been confined to theories, and very little is known about their behaviours on belief bases. Rational contraction operators that satisfy the supplementary postulates will be dubbed *fully rational*.

3.1. Partial Meet and Smooth Kernel Contractions

Several rational contraction operators were proposed in the literature. The two most influential ones are partial meet (Definition 4), and Smooth Kernel (Definition 9). Partial meet makes use of remainders.

Definition 1. *Given a belief base $\mathcal{K}$ and formula α, an α-remainder of $\mathcal{K}$ is a set $X \subseteq \mathcal{K}$ such that: $\alpha \notin Cn(X)$, and if $X \subset Y \subseteq \mathcal{K}$, then $\alpha \in Cn(Y)$. The set of all α-remainders of $\mathcal{K}$ is denoted by $\mathcal{K} \perp \alpha$.*

Each member of $\mathcal{K} \perp \alpha$ is called a remainder, and it is a maximal subset of $\mathcal{K}$ that does not entail α. A partial meet operator works by selecting remainders and intersecting them. As a remainder set might have many remainders, a choice must be made about which ones are the best to perform the contraction. This choice is done via an extra-logical mechanism called a *selection function*:

Definition 2. *A selection function γ picks some remainder from $\mathcal{K} \perp \alpha$ such that,*

(i) $\gamma(\mathcal{K} \perp \alpha) \neq \emptyset$; and

(ii) $\gamma(\mathcal{K} \perp \alpha) \subseteq \mathcal{K} \perp \alpha$, if $\mathcal{K} \perp \alpha \neq \emptyset$; and

(iii) $\gamma(\mathcal{K} \perp \alpha) = \{\mathcal{K}\}$, if $\mathcal{K} \perp \alpha = \emptyset$.

A selection function works as an extra-logical mechanism that realises the agent's epistemic preferences. In the original work of [2], the authors propose to represent an agent's preferences as a binary relation $\leqslant$ on all remainders. Precisely, a pair $A \leqslant B$ means that the remainder A is at least as preferable as B. The agent picks the most preferable α-remainders w.r.t $\leqslant$.

Definition 3. *A selection function γ is relational iff there exists some binary relation $\leqslant$ on all remainders such that $\gamma(\mathcal{K} \perp \alpha) = \min_{\leqslant}(\mathcal{K} \perp \alpha)$, for all $\mathcal{K} \perp \alpha \neq \emptyset$. If $\leqslant$ is transitive then γ is called transitive relational.*

Remainder sets and selection functions are used to define a contraction operator called *partial meet contraction*:

Definition 4. *Given a belief base $\mathcal{K}$, and a selection function γ, the operation $\dot{-}_\gamma$ defined as $\mathcal{K} \dot{-}_\gamma \alpha = \bigcap \gamma(\mathcal{K} \perp \alpha)$ is a partial meet contraction function.*

Theorem 5. *[19] A contraction operator is rational iff it is a partial meet contraction operator.*

For theories, the transitive relational partial meet operators are characterised by all the rationality postulates of contraction.

Theorem 6. *[2] On theories, a contraction operator is fully rational iff it is a transitive relational partial meet contraction operator.*

As Hansson [18] shows, the transitive relational partial meet operators are not strong enough to satisfy the two supplementary postulates on belief bases. Hansson proposes to strengthen the transitive relations with a property called maximising. However, a representation theorem is not obtained.

Another influential class of rational contraction operations is the class of smooth kernel contraction operations, which are defined on kernels and incision functions:

Definition 7. *An α-kernel of a belief base $\mathcal{K}$ is a set X such that (1) $X \subseteq \mathcal{K}$; (2) $\alpha \in Cn(X)$; and (3) if $X' \subset X$ then $\alpha \notin Cn(X')$.*

An α-kernel of a belief base $\mathcal{K}$ is a minimal subset of $\mathcal{K}$ that does entail α. The set of all α-kernels of a belief base $\mathcal{K}$ is denoted by $\mathcal{K} \perp\!\!\!\perp \alpha$. Formulae that do not appear in any α-kernel are not responsible for entailing the formula α to be contracted, and therefore they should be kept intact. In contrast, only formulae that appear in the kernels should be picked for removal. This choice of removal is realised by an incision function:

Definition 8. *Let $\mathcal{C}(\mathcal{K}) = \{\mathcal{K} \perp\!\!\!\perp \alpha \mid \alpha \in \mathcal{L}\}$ be the set of all kernel sets of $\mathcal{K}$. An incision function on a belief base $\mathcal{K}$ is a function $\sigma : \mathcal{C}(\mathcal{K}) \to \mathcal{P}(\mathcal{L})$ such that*

(1) $\sigma(\mathcal{K} \perp\!\!\!\perp \alpha) \subseteq \bigcup \mathcal{K} \perp\!\!\!\perp \alpha$;

(2) if $X \in \mathcal{K} \perp\!\!\!\perp \alpha$ and $X \neq \emptyset$, then $X \cap \sigma(\mathcal{K} \perp\!\!\!\perp \alpha) \neq \emptyset$.

Intuitively, in order to contract a formula α, an agent chooses at least one formula from each α-kernel, and only formulae from such kernels. An incision function works as an extra-logical device that realises an agent's epistemic preferences, and it chooses the least preferable formulae in each α-kernel to be removed. A contraction operation can be constructed by removing the formulae picked by an incision function. Contraction operations that follow this recipe are called kernel contractions:

Definition 9. *[5] Given a belief base $\mathcal{K}$ and an incision function σ for $\mathcal{K}$, the kernel contraction function $\dot{-}_\sigma$ is defined as: $\mathcal{K} \dot{-}_\sigma \alpha = \mathcal{K} \setminus \sigma(\mathcal{K} \perp\!\!\!\perp \alpha)$.*

Kernel contractions functions, however, are not strong enough to satisfy relevance and relative closure. To capture relative closure, Hansson [5] has proposed the *smoothness* property for incision functions:

smoothness: if $\mathcal{K}' \subseteq \mathcal{K}$, $\varphi \in Cn(\mathcal{K}')$ and $\varphi \in \sigma(\mathcal{K} \perp\!\!\!\perp \alpha)$ then $\mathcal{K}' \cap \sigma(\mathcal{K} \perp\!\!\!\perp \alpha) \neq \emptyset$.

Incision functions that satisfy smoothness are called *smooth incision function* and the respective kernel contractions are called *smooth kernel contraction operations*. Intuitively, smoothness states that any removed formula cannot be entailed by the remaining formulae.

The *smooth kernel contraction operations* are characterised by the first six rationality postulates:

Theorem 10. *[5, 20] A contraction function satisfies* success, inclusion, vacuity, uniformity, core-retainment, *and* relative closure *iff it is a smooth kernel contraction function.*

3.2. Semantic Contraction Operators

We start by explaining how belief contraction works on models when the agent's corpora of beliefs are represented as theories. After that, we show why such strategies do not work for belief bases.

In terms of models, in order to contract a formula α from a theory $\mathcal{K}$, it suffices to obtain a theory that is a subset of $\mathcal{K}$ (due to the inclusion postulate) and it is satisfied by some counter-models of α. This can be formalised by taking a function $\mu : \mathcal{L} \to \mathcal{P}(\mathcal{I})$ that picks, for every non-tautological formula α, some counter-models of α. For tautological formulae α, we make $\mu(\alpha) = \emptyset$, as tautologies have no counter-models. When $\alpha \notin Cn(\mathcal{K})$, there is nothing to be removed, and $\mathcal{K}$ should be kept untouched, according to *vacuity*. Therefore, in this case, we make $\mu(\alpha) = \overline{[\![\alpha]\!]} \cap [\![\mathcal{K}]\!]$, that is, the most plausible counter-models of α are those ones that satisfy $\mathcal{K}$. Moreover, if two formulae α and β are logically equivalent, then $\mu(\alpha) = \mu(\beta)$. This guarantees that the choice function is not syntax sensitive. We say that μ is a model choice function.

Definition 11. *The contraction function induced by a model choice function μ is the operator*

$$\mathcal{K} \dot{-}_\mu \alpha = \{\varphi \in \mathcal{K} \mid \mu(\alpha) \subseteq [\![\varphi]\!]\}.$$

Indeed, the basic rationality postulates of contraction characterise such a class of semantic contraction operators for theories, as long as the underlying logic is Tarskian, compact and the language is closed under classical negation and disjunction[12, 10][1]. Examples of such logics include classical propositional logic, first-order logic, a number of temporal logics, and several normal-modal logics such as the systems K, T and S_4.

Theorem 12. *A contraction function $\dot{-}$ on a theory $\mathcal{K}$ is rational iff it is induced by some model choice function μ.*

[1] The language of the logic contains the classical boolean operators of negation and disjunction and they are interpreted as usual.

For full rationality, there are two main classes of belief operators: the revision operators based on faithful pre-orders of Katsuno and Mendelzon (KM, for short) [7] and the revision operators based on Grove's spheres[6]. Although both classes of operators were originally framed for belief revision, they can be easily translated to contraction. In the following, we present a translation of KM operators based on faithful pre-orders in terms of contraction. Caridroit et al. [21] have shown a similar translation, for classical propositional logic, where a theory is represented as a single formula. The translation we present below works directly on bases (sets of formulae instead of a single formula).

Definition 13. *[7][2] Given a belief base $\mathcal{K}$, a pre-order $\leqslant_{\mathcal{K}}$ is faithful w.r.t $\mathcal{K}$ iff it satisfies the two following conditions:*

(1) if $M, M' \in [\![\mathcal{K}]\!]$ then $M \not<_{\mathcal{K}} M'$;

(2) if $M \in [\![\mathcal{K}]\!]$ and $M' \notin [\![\mathcal{K}]\!]$ then $M <_{\mathcal{K}} M'$.

Definition 14. *Given a faithful pre-order $\leqslant_{\mathcal{K}}$ on a belief base $\mathcal{K}$, the faithful contraction operator founded on $\leqslant_{\mathcal{K}}$ is the operation $\dot{-}_{\leqslant_{\mathcal{K}}}$ such that $[\![\mathcal{K} \dot{-}_{\leqslant_{\mathcal{K}}} \alpha]\!] = [\![\mathcal{K}]\!] \cup \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$. If $\leqslant_{\mathcal{K}}$ is total then $\dot{-}_{\leqslant_{\mathcal{K}}}$ is a total faithful contraction operator.*

A faithful pre-order works as an epistemic preference relation on models. In order to contract a formula α, the agent chooses exactly the most plausible counter-models of α. In the current presentation, KM operators are suitable only for theories, because, for belief bases, there is no guarantee that $\mathcal{K} \dot{-}_{\leqslant_{\mathcal{K}}} \alpha$ outputs a subset of $\mathcal{K}$, as the inclusion postulate demands. Towards this end, in order to satisfy the inclusion postulate we need only to rewrite faithful contraction in the spirit of Definition 11: get the greatest subset of $\mathcal{K}$ satisfied by the minimal counter-models of the formula α to be contracted. Indeed, within classical propositional logics, the KM operations is a special kind of contraction induced by a model choice function as per Definition 11. In classical propositional logics, for theories, the faithful contraction operators on total pre-orders are fully rational:

Theorem 15. *[7, 21] In classical propositional logics, a contraction operator on a theory $\mathcal{K}$ is fully rational iff it is a total faithful contraction operator.*

Observe that the representation theorems above (Theorem 12 and Theorem 15) are established only for theories. Indeed, as Example 1 below illustrates, both representation theorems break down for bases, which is due to violation of the relevance postulate.

[2]Originally, KM defines an assignment that maps each formula to a pre-order, and defines such an assignment to be faithful. This assignment has only the purpose to provide general contraction operators. As here we focus on local contraction, we opt to remove this complication and operate directly on the pre-orders.

Example 1. *Consider the belief base $\mathcal{K} = \{p, q, p \vee q, \neg q \vee p\}$, expressed in classical propositional logics, with $AP = \{p, q\}$. We want to contract the formula $p \wedge q$. There are only three rational contraction results:*

$$A_1 = \{p, p \vee q, \neg q \vee p\}, \quad A_2 = \{q, p \vee q\},$$
$$A_3 = \{p \vee q\}.$$

Not every model choice function, however, induces a rational contraction operator. To see this, note that we have only four models

$$M_1 = \{p, q\}, M_2 = \{p\}, M_3 = \{q\}, \text{ and } M_4 = \emptyset.$$

Observe that $\overline{[\![p \wedge q]\!]} = \{M_2, M_3, M_4\}$. Let $\leqslant_{\mathcal{K}}$ be the following faithful pre-order on $\mathcal{K}$:

$$M_1 \leqslant_{\mathcal{K}} M_4 \leqslant_{\mathcal{K}} M_3 \leqslant_{\mathcal{K}} M_2.$$

Let σ be a model choice function such that $\sigma(p \wedge q) = \min_{\leqslant_{\mathcal{K}}}(\overline{[\![p \wedge q]\!]}) = \{M_4\}$. The only formula of $\mathcal{K}$ that M_4 satisfies is $\neg q \vee p$. Thus, $\mathcal{K} \dot{-}_{\sigma} p \wedge q = \{\neg q \vee p\}$. However, this does not correspond to any of the three possible rational contraction results listed above.

4. Tracks and Mirrors: Belief Base Contraction on Models

In this section, we provide a novel class of semantic contraction operators for belief bases.

In terms of models, contracting a formula α from a theory $\mathcal{K}$ consists in picking some counter-models of α and maintaining the formulae in $\mathcal{K}$ satisfied by all such picked counter-models. While this strategy yield rational contractions for theories (Theorem 12), it fails for belief bases as Example 1 illustrates. This occurs because some counter-models of α might satisfy less formulae than allowed by the relevance postulate. For instance, looking back at Example 1, according to relevance the formula $p \vee q$ must be kept. Observe that this formula appears in all the three possible rational contraction results. The counter-model M_4, however, does not satisfy $p \vee q$, which makes it unsuitable for performing a rational contraction, as picking it would remove $p \vee q$. The main hurdle is to properly distinguish between suitable and unsuitable models. To solve this problem, we establish a plausibility relation $\leqslant$ on the models. Intuitively, a pair $M \leqslant M'$ means that the model M is at least as plausible as M'. Towards this end, in order to contract a formula α, only the most plausible counter-models of α w.r.t $\leqslant$ should be chosen, that is, only models within $\min_{\leqslant}(\overline{[\![\alpha]\!]})$. The question at hand is which properties a pre-order on models should satisfy in order to be an adequate plausibility relation that distinguishes between suitable and unsuitable models.

Here, we propose such plausibility relations be defined upon the notion of information preservation. Intuitively, the more information from $\mathcal{K}$ a model preserves the more plausible it is. The set of all formulae from $\mathcal{K}$ satisfied by a model M is given by the set $Pres(M \mid \mathcal{K}) = \{\varphi \in \mathcal{K} \mid M \models \varphi\}$. Generalising, given a set X of models, $Pres(X \mid \mathcal{K}) = \{\varphi \in \mathcal{K} \mid M \models \varphi, \text{ for all } M \in X\}$. Definition 16 below formalises a class of pre-orders based on this notion, which we call tracks.

Definition 16. *A track of a belief base $\mathcal{K}$ is a pre-order $\leqslant_{\mathcal{K}} \subseteq \mathcal{I} \times \mathcal{I}$ such that*

(1) If $Pres(M \mid \mathcal{K}) = Pres(M' \mid \mathcal{K})$ then $M' \leqslant_{\mathcal{K}} M$ and $M \leqslant_{\mathcal{K}} M'$; and

(2) If $Pres(M \mid \mathcal{K}) \subset Pres(M' \mid \mathcal{K})$ then $M' <_{\mathcal{K}} M$.

In short, a track relation imposes models that strictly preserve more information to be strictly more plausible (condition 2), while models that preserve the same set of information are equally plausible (condition 1). Thus, in every track for a belief a base $\mathcal{K}$, the models of $\mathcal{K}$ are the most plausible ones, and they are also all equally plausible.

A least track of a knowledge base $\mathcal{K}$ is a least relation satisfying all conditions of Definition 16. It is easy to see that every belief base has a unique least track. We denote the least track of a belief base $\mathcal{K}$ as $\leqslant^{-}_{\mathcal{K}}$.

Proposition 17. *If $\mathcal{K}$ is a consistent belief base and $\leqslant_{\mathcal{K}}$ is a track of $\mathcal{K}$ then $\min_{\leqslant_{\mathcal{K}}}(\mathcal{I}) = [\![\mathcal{K}]\!]$.*

Example 2 (continued from Example 1). *The beliefs in $\mathcal{K} = \{p, q, p \vee q, p \vee \neg q\}$ preserved by each of the four models are:*

$$\begin{aligned}
&Pres(M_1 \mid \mathcal{K}) = \mathcal{K}\\
&Pres(M_2 \mid \mathcal{K}) = \{p, p \vee q, \neg q \vee p\}\\
&Pres(M_3 \mid \mathcal{K}) = \{q, p \vee q\}\\
&Pres(M_4 \mid \mathcal{K}) = \{\neg q \vee p\}.
\end{aligned}$$

Fig. 1 (on the right) illustrates the set inclusion relation between the preservation sets of each model, while Fig. 1 (on the left) depicts the least track relation of $\mathcal{K}$. As M_1 is the only model of $\mathcal{K}$, it is strictly more plausible than all other models. Models M_2 and M_3 are incomparable, since they preserve different beliefs in $\mathcal{K}$. For the same reason, M_4 and M_3 are incomparable. However, M_2 is strictly more plausible than M_4, as M_2 preserves strictly more information than M_4. At this point, we can see that a track can distinguish between suitable and unsuitable models. According to this track, both models M_2 and M_3 are the most plausible counter-models of $p \wedge q$. If we choose either M_2 or M_3 then we get a rational contraction: either $A_1 = \{p, p \vee q, \neg q \vee p\}$, or $A_2 = \{q, p \vee q\}$. By picking both models we get the last rational contraction $A_3 = \{p \vee q\}$. The only non-rational contractions are those involving the model M_4 which is not among the most plausible ones (the suitable ones). Also, observe that other tracks exist: for instance, augmenting the illustrated track by making M_2 and M_3 comparable or even M_3 and M_4 comparable. However, for any of the possible tracks, M_4 is never among the suitable ones, as it must be strictly less plausible than M_2, due to condition 2 of the track's definition. This suggests that tracks can be used as an adequate class of plausibility relations to distinguish between suitable and unsuitable models.

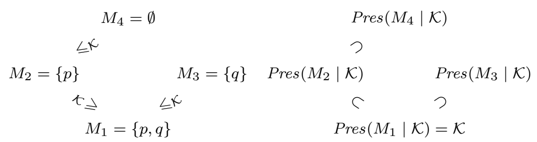

Figure 1: The least track relation $\leqslant_{\mathcal{K}}$(on the left), and the set inclusion relation on the preservation set of the models (on the right).

As tracks establish an adequate notion of plausibility between models, the most plausible ones to contract a formula α are the minimal counter-models of α. In classical propositional logics, such minimal models always exist, as there is only a finite number of models. However, for more expressive logics, such as First Order Logics and several Description Logics [22], there are formulae with an infinite number of (counter-)models. In the presence of an infinite amount of models, some tracks arrange the models through infinite chains. In general, these infinite chains prevent identifying the most plausible counter-models for some formulae. Thus, we need to constrain ourselves to tracks that do not present such bad behaviour, that is, tracks that are *founded*:

Definition 18. *A relation $\leqslant \subseteq \mathcal{I} \times \mathcal{I}$ is founded iff $\min_{\leqslant}(\overline{[\![\alpha]\!]}) \neq \emptyset$ for every non-tautological formula α.*

Relying on founded tracks guarantees that for every non-tautological formula α, there is at least one counter-model to be picked to perform such a contraction. In fact, as long as the underlying Tarskian logic satisfies compactness, every belief base presents at least one founded track: its least track.

Theorem 19. *If a logic $\langle \mathcal{L}, Cn \rangle$ is Tarskian and compact then for every belief base $\mathcal{K} \subseteq \mathcal{L}$, the least track is founded.*

We can then define a function that selects among the most plausible models:

Definition 20. *Let $\leqslant_{\mathcal{K}}$ be a founded track. A tracking selection function on $\leqslant_{\mathcal{K}}$ is a function $\delta_{\leqslant_{\mathcal{K}}} : \mathcal{L} \to \mathcal{P}(\mathcal{I})$ such that*

1. $\delta_{\leqslant_{\mathcal{K}}}(\alpha) \subseteq \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$;

2. $\delta_{\leqslant_{\mathcal{K}}}(\alpha) \neq \emptyset$, *if* α *is not a tautology;*
3. *if* α *and* β *are* $\mathcal{K}$*-uniform,* $M \in \delta_{\leqslant_{\mathcal{K}}}(\alpha)$, $M \sim_{\mathcal{K}} M'$ *and* $M' \in \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\beta]\!]})$ *then* $M' \in \delta_{\leqslant_{\mathcal{K}}}(\beta)$.

A tracking selection function works similarly to the model choice function for theories. The main difference is that model choice functions can choose any counter-models of a formula α, while tracking selection functions choose only among the most plausible (w.r.t a track relation) counter-models of α. Condition 3 is related to the postulate of uniformity, and guarantees that a tracking selection function is not syntax sensitive. Precisely, it states that if two models M and M' are respectively counter-models of α and β and they are equally preferable, then picking M to contract α implies picking M' to contract β. Example 3 illustrate a tracking selection function and the role of this condition. When it is clear from context, we drop the subscript $_{\leqslant_{\mathcal{K}}}$ and write δ.

Example 3. *Let* $\mathcal{K} = \{p \vee q, p \leftrightarrow q\}$ *be a knowledge base. Observe that the formulae* p *and* q *are* $\mathcal{K}$*-uniform. There are only three possible results to contract either* p *or* q *that satisfy relevance, which are*

$$A_1 = \{p \vee q\}, \quad A_2 = \{p \leftrightarrow q\} \quad \text{and } A_3 = \emptyset.$$

Recall that $\leqslant^-_{\mathcal{K}}$ *denotes the least track of* $\mathcal{K}$*. Assume we want the solution* A_1 *for contracting either the formulae* p *or* q*. Thus, a track selection function* $\delta_{\leqslant^-_{\mathcal{K}}}$ *can pick only counter-models that satisfy* A_1*, when contracting such formulae. We have only four models:*

$$M_1 = \{p, q\} \qquad M_2 = \{p\}$$
$$M_3 = \{q\} \qquad M_4 = \emptyset.$$

Fig. 2 illustrates the least track $\leqslant^-_{\mathcal{K}}$ *on the base* $\mathcal{K}$*. For clarity, in Fig. 2, we depict within rectangles the formulae from* $\mathcal{K}$ *that are satisfied by each model. The counter-models of* p *are* M_3 *and* M_4*, and the only one satisfying* A_1 *is* M_3*. So, we make* $\delta_{\leqslant^-_{\mathcal{K}}}(p) = \{M_3\}$*. As* p *and* q *are* $\mathcal{K}$*-uniform, their contraction must coincide. Ideally, we would make* $\delta_{\leqslant_{\mathcal{K}}}(p) = \delta_{\leqslant_{\mathcal{K}}}(q)$*. However, this is not possible, as* M_3 *is not a counter-model of* q*. In fact, the only counter-models of* q *are* M_2 *and* M_4*. Observe that* M_2 *is the only counter-model of* q *that satisfy* A_1*. Therefore, the track selection function must choose* M_2*, that is,* $\delta_{\leqslant^-_{\mathcal{K}}}(q) = \{M_2\}$*. Not surprisingly,* M_2 *and* M_3 *are equally preferable modulo* $\leqslant^-_{\mathcal{K}}$*, and according to Condition 3 from the definition of track selection function* M_2 *must be picked for contracting* q*, since* M_1 *was chosen to contract* p*. This condition, as this example illustrates, ensures uniformity.*

Following the same strategy as for theories, a contraction on a belief base is performed by keeping the formulae from the current belief base that are satisfied by all the counter-models selected by a tracking selection function.

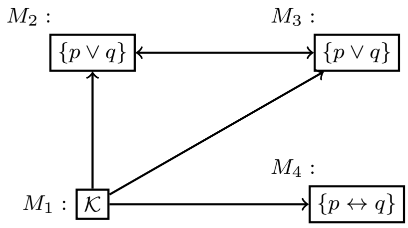

Figure 2: The least track on the base of Example 3.

Definition 21. *Let* δ *be a tracking selection function. The tracked contraction founded on* δ *is defined as*

$$\mathcal{K} \dot{-}_{\delta} \alpha = \{\varphi \in \mathcal{K} \mid \delta(\alpha) \subseteq [\![\varphi]\!]\}.$$

Example 4 (continued from Example 2)**.** *Let* $\leqslant^-_{\mathcal{K}}$ *be the least track of the belief base* $\mathcal{K} = \{p, q, p \vee q, \neg q \vee p\}$*. Observe that* $\min_{\leqslant^-_{\mathcal{K}}}(p \wedge q) = \{M_2, M_3\}$*. Then, we can choose any combination of* M_2 *and* M_3 *to contract* $p \wedge q$*. Let* δ_1, δ_2 *and* δ_3 *be tracked selection functions founded on* $\leqslant^-_{\mathcal{K}}$ *such that* $\delta_1(p \wedge q) = \{M_2\}, \delta_2(p \wedge q) = \{M_3\}$ *and* $\delta_3(p \wedge q) = \{M_2, M_3\}$*. They induce the following tracked contraction operators:* $\mathcal{K} \dot{-}_{\delta_1} \neg q \vee p = \{p, p \vee q, \neg q \vee p\}$*,* $\mathcal{K} \dot{-}_{\delta_2} \neg q \vee p = \{q, p \vee q\}$*, and* $\mathcal{K} \dot{-}_{\delta_3} \neg q \vee p = \{p \vee q\}$*. As one can easily check, each one of them is a rational contraction operator.*

Theorem 22. *Every tracked contraction function is rational.*

Proof sketch. Postulates of *success, inclusion, vacuity* and *uniformity* are easy to prove. We focus on relevance. Let $\beta \in \mathcal{K} \setminus (\mathcal{K} \dot{-}_{\delta_{\leqslant}} \alpha)$. Thus, there is some model $M \in \delta_{\leqslant_{\mathcal{K}}}(\alpha)$ such that $M \not\models \beta$. As $M \in \delta_{\leqslant_{\mathcal{K}}}(\alpha)$, we have that $M \models \mathcal{K} \dot{-}_{\delta_{\leqslant}} \alpha$ and $M \in \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$. Thus,

$$\mathcal{K} \dot{-}_{\delta_{\leqslant}} \alpha \subseteq \mathit{Pres}(M \mid \mathcal{K}) \subseteq \mathcal{K}.$$

Let us suppose for contradiction that $\alpha \notin Cn(\mathit{Pres}(M \mid \mathcal{K}) \cup \{\beta\})$. Thus, there is some model $M' \in \overline{[\![\alpha]\!]}$ such that

$$M' \models \mathit{Pres}(M \mid \mathcal{K}) \cup \{\beta\},$$

which implies $\mathit{Pres}(M \mid \mathcal{K}) \cup \{\beta\} \subseteq \mathit{Pres}(M' \mid \mathcal{K})$. As $M \not\models \beta$, we have that $\beta \notin \mathit{Pres}(M \mid \mathcal{K})$. Therefore, $\mathit{Pres}(M \mid \mathcal{K}) \subset \mathit{Pres}(M \mid \mathcal{K}) \cup \{\beta\}$. This means that $\mathit{Pres}(M \mid \mathcal{K}) \subset \mathit{Pres}(M' \mid \mathcal{K})$ which implies that $M' <_{\mathcal{K}} M$. Therefore, $M \notin \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$, which is a contradiction. ☐

Theorem 23. *Every rational base contraction function is a tracked contraction function.*

Since a track establishes a plausibility relation between models, it is natural to expect that a track also works as an epistemic preference relation. Therefore, instead of simply picking some of the most plausible models w.r.t a track, it would be rational to pick all such most plausible models. We will call contraction operators that follow this strategy full tracked contraction:

Definition 24. *Let $\leqslant_\mathcal{K}$ be a founded tracking of a belief base $\mathcal{K}$. The full tracked selection of $\leqslant_\mathcal{K}$ is the function $\nu_{\leqslant_\mathcal{K}}$ such that $\nu_{\leqslant_\mathcal{K}}(\alpha) = \min_{\leqslant_\mathcal{K}}(\overline{[\![\alpha]\!]})$. Tracked contraction operators founded on full tracking selection functions are full tracked contraction operators.*

Full tracked contraction operators do satisfy *intersection*, due to the transitivity of tracks.

Theorem 25. *Every full tracked contraction satisfies intersection.*

Although tracks capture *intersection*, they are not strong enough to capture *conjunction*. Observe that tracks form a special case of faithful pre-orders (Definition 14). It would be natural then to simply impose totality upon the tracks in the hope of capturing *conjunction*. Totality, however, has been criticised in the literature for being too demanding, as an agent might be indifferent or ignorant on how to grade some of its beliefs [23]. Moreover, works such as [10, 24] have observed that totality is not strong enough to capture *conjunction*, even for theories, in more expressive logics. As a solution, Ribeiro et al. [10] has introduced mirroring:

> ***mirroring:*** if $A \not\leqslant B$ and $B \not\leqslant A$ but $C \leqslant A$ then $C \leqslant B$.

Mirroring is similar to the modular relations introduced at [23] which was based on the modular partial orders of Lehmann and Magidor [25]. Though modular relations are defined as partial orders, we do not impose such restrictions. According to mirroring, if two models are incomparable then they should agree upon their preferences. We will show here that by employing mirroring upon tracks, *conjunction* is also captured for belief bases.

Theorem 26. *If a founded track $\leqslant_\mathcal{K}$ satisfies mirroring than its full tracked contraction operator satisfies conjunction.*

5. Smooth Kernel Contraction: A Semantic Perspective

In this section, we characterise semantically the class of smooth kernel contraction operations. In terms of rationality postulates, we are capturing core-retainment and relative closure. While semantic operators satisfying relevance, as shown in the previous section, select only countermodels of the formula α being contracted; some operations satisfying core-retainment do incorporate models of α. This exhibits the permissive and drastic behaviour of smooth kernel contraction for bases. Example 5 illustrates this behaviour.

Example 5. *Let $\mathcal{K} = \{p, p \to q, p \vee q, r\}$, and suppose that we want to contract 'q'. There are only four possible solutions satisfying both core-retainment and relative closure:*

$$A_1 = \{p, p \vee q, r\} \qquad A_2 = \{p \to q, r\}$$
$$A_3 = \{p \vee q, r\} \qquad A_4 = \{r\}.$$

Solutions A_1, A_2 and A_4 satisfy relevance, while A_3 does not satisfy relevance but core-retainment. The base A_3 can only be obtained selecting the models $\{p, r\}$ and $\{q, r\}$. Observe that the latter model satisfies q. Therefore, in order to capture core-retainment, it is necessary to relax the selection functions to choose both models and counter-models of the formulae to be contacted.

As Example 5 illustrates, we need to allow selection functions to pick not only counter-models but also models of the formulae being contracted. However, even for core-retainment, not all models can be chosen. For instance, although $M' = \{q\}$ is a model of 'q', M' violates all the four rational solution for contracting q in Example 5, as M' violates r. On one hand, we need to relax the selection functions to pick models of the formulae being contracted. On the other hand, we need to constrain the selection function so we do not choose unsuitable models. The tracks still capture enough information to allow distinguishing between such suitable and unsuitable models. We slightly modify the definition of the tracking selection function to capture this permissive behaviour:

Definition 27. *A permissive selection function on a founded track $\leqslant_\mathcal{K}$ is a map $\lambda_{\leqslant_\mathcal{K}} : \mathcal{L} \to \mathcal{P}(\mathcal{I})$ such that*

(1) $\lambda_{\leqslant_\mathcal{K}}(\alpha) = \emptyset$, if α is a tautology;

(2) $\lambda_{\leqslant_\mathcal{K}}(\alpha) \cap \overline{[\![\alpha]\!]} \neq \emptyset$, if α is not a tautology;

(3) $\lambda_{\leqslant_\mathcal{K}}(\alpha) = \lambda_{\leqslant_\mathcal{K}}(\beta)$, if α and β are $\mathcal{K}$-uniform;

*(4) **permissiveness:** if $M \in \lambda_{\leqslant_\mathcal{K}}(\alpha)$, then*
$$Pres(\min_{\leqslant_\mathcal{K}}(\overline{[\![\alpha]\!]}) \mid \mathcal{K}) \subseteq Pres(M \mid \mathcal{K}).$$

As tautologies cannot be contracted, Condition 1 enforces that no model will be picked for tautologies. Condition 2 relaxes the selection mechanism to choose both models and counter-models, while enforcing that at least one counter-model will be chosen, so the contraction is successful. Condition 3 is related to the *uniformity* postulate, and states that $\mathcal{K}$-uniform formulae present the same choice. Since models are allowed to be picked, the last condition, permissiveness, dictates how permissive

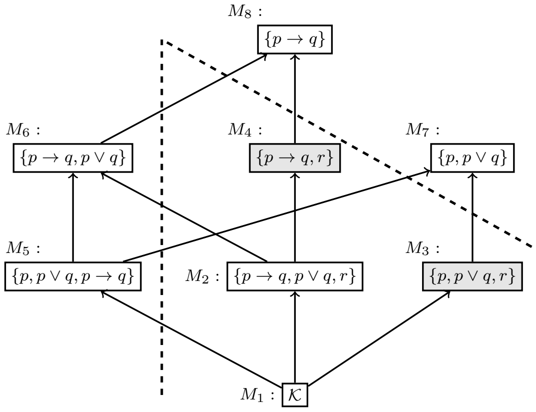

Figure 3: The least track on the base of Example 6. The relation is transitive, but to avoid visual pollution we omit edges obtained by transitivity.

the selection mechanism can be. While contracting a formula α, instead of picking only the best models w.r.t the track relation, permissiveness allows any (counter)model M to be chosen, as long as M preserves as much information as the best counter-models of α. For clarity, we will omit the subscript $\leqslant_{\mathcal{K}}$ and simply write λ.

Example 6. *(continued from Example 5). We have eight models in total:*

$$\begin{aligned} M_1 &= \{p, q, r\} & M_5 &= \{p, q\} \\ M_2 &= \{q, r\} & M_6 &= \{q\} \\ M_3 &= \{p, r\} & M_7 &= \{p\} \\ M_4 &= \{r\} & M_8 &= \emptyset. \end{aligned}$$

Fig. 3 illustrates the least track $\leqslant^-_{\mathcal{K}}$ for the knowledge base $\mathcal{K}$. For clarity, in Fig. 3, we depict within rectangles the formulae from $\mathcal{K}$ that are satisfied by each model. Observe that the counter-models of q are $\{M_3, M_4, M_7, M_8\}$, and $\min_{\leqslant^-_{\mathcal{K}}}(\overline{[\![q]\!]}) = \{M_3, M4\}$ which are coloured in gray. A selection function that picks only M_3 or M_4 yields respectively the solutions A_1 and A_2, while picking both M_3 and M_4 yields the solution A_4. The solution A_3, which satisfies core-retainment but fails relevance, can only be obtained by choosing the model M_2. Observe that M_2 preserves as much as M_3 and M_4 combined, that is,

$$Pres(M_3 \mid \mathcal{K}) \cap Pres(M_4 \mid \mathcal{K}) \subseteq Pres(M_2 \mid \mathcal{K}).$$

Therefore, according to permissiveness, a selection function can choose any of the models in $\{M_2, M_3, M_4\}$. Notice that M_2 is a model of 'q', while M_3 and M_4 are counter-models of 'q'. The models that preserve as much information as M_3 and M_4 combined are depicted within the dashed lines.

The contraction function is defined analogously to the tracked contractions:

Definition 28. *Let λ be a permissive selection function on a track $\leqslant_{\mathcal{K}}$. The permissive contraction founded on λ is defined as $\mathcal{K} \dot{-}_\lambda \alpha = \{\varphi \in \mathcal{K} \mid \lambda(\alpha) \subseteq [\![\varphi]\!]\}$.*

The permissive contraction operators are as rational as smooth kernel contraction operators:

Theorem 29. *Every permissive contraction function satisfies success, inclusion, vacuity, uniformity, core-retainment and relative closure.*

Theorem 30. *If $\dot{-}$ satisfies success, inclusion, vacuity, uniformity and core-retainment and relative closure, then $\dot{-}$ is a permissive contraction.*

Our representation result follows from Theorem 29 and Theorem 30 which jointly state that the most basic rationality postulates characterize the class of permissive contraction operations. This result jointly with Theorem 10 implies that smooth kernel contraction operations and permissive contraction correspond to the same class of operators: being the latter the semantic counterpart of the former.

6. Conclusion and Future Works

While both syntactic and semantic operators are well known for belief theory contraction (and other forms of belief change), only syntactic operators are known to be rational on belief bases. In this work, we have introduced new classes of semantic contraction operators for belief bases: *tracked contraction operators*, *full tracked contraction operators*, and *permissive tracked contraction operators*. These operators rely on plausibility relations between models, called tracks.

In order to contract a formula α, the (full) tracked contraction operators select among the most plausible counter-models of α w.r.t a track relation (the most reliable ones). The permissive tracked contraction relaxes the selection mechanisms, allowing to pick models instead of only counter-models, as long as some innocuous requirements are satisfied. We have established two important representation theorems: the first one connects tracked contraction operations with relevance and the other basic rationality postulates, while the second one connects the permissive tracked contraction operators with core-retainment and the most basic rationality postulates. Equivalently, the tracked contraction operations semantically characterize the partial meet operators, while the permissive tracked contraction operators characterize semantically the smooth kernel contraction operators. A track unveils an agent's epistemic preferences: the most plausible models coincide with the most reliable ones, and therefore the agent should pick all such models. Tracked contractions following this strategy are called full tracked contractions. We have shown that tracks that satisfy the mirroring condition yield full tracked contraction satisfying the two supplementary postulates.

As future work, we shall investigate if mirroring suffices to establish a representation theorem between fully tracked contractions and the supplementary postulates. This connection with the supplementary postulates is important, because the study of such postulates has been restricted to belief change operators on theories. Particularly, the connection between contraction operators and the supplementary postulates has been established via epistemic preferences relations such as Epistemic Entrenchment [3] and Hierarchies (for safe contraction) [4]. Although all such epistemic preferences work well for theories, their connection with such rationality postulates easily disappears for bases. The only known exception is safe contraction, which still connects with the supplementary postulates only when a base $\mathcal{K}$ is finite and it is as expressive as its theory $Cn(\mathcal{K})$: for every formula $\alpha \in Cn(\mathcal{K})$ there is a formula in $\mathcal{K}$ logically equivalent to α.

We shall extend our results for more expressive logics by dispensing with compactness and widening our results to Tarskian logics. Although we have focused on contraction, our results can be easily translated to revision: instead of selecting counter-models, one needs only to select models of the formulae α to be revised.

Acknowledgments

This research is supported by the German Research Association (DFG), project number 465447331.

References

[1] J. S. Ribeiro, Towards a semantic construction for belief base contraction (A preliminary report), in: FCR@KI, volume 3242 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 4–15.

[2] C. E. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: partial meet contraction and revision functions, The Journal of Symbolic Logic 50 (1985) 510–530.

[3] P. Gärdenfors, Knowledge in flux: Modeling the dynamics of epistemic states., The MIT press, 1988.

[4] S. O. Hansson, A textbook of belief dynamics - theory change and database updating, volume 11 of *Applied logic series*, Kluwer, 1999.

[5] S. O. Hansson, Kernel contraction, J. Symb. Log. 59 (1994) 845–859.

[6] A. Grove, Two modellings for theory change, J. Philos. Log. 17 (1988) 157–170.

[7] H. Katsuno, A. O. Mendelzon, Propositional knowledge base revision and minimal change, Artificial Intelligence 52 (1991) 263–294.

[8] J. S. Ribeiro, M. Thimm, Consolidation via tacit culpability measures: Between explicit and implicit degrees of culpability, in: KR 2021, 2021, pp. 529–538.

[9] M. M. Ribeiro, R. Wassermann, G. Flouris, G. Antoniou, Minimal change: Relevance and recovery revisited, Artif. Intell. 201 (2013) 59–80.

[10] J. S. Ribeiro, A. Nayak, R. Wassermann, Towards belief contraction without compactness, in: KR 2018, AAAI Press, 2018, pp. 287–296.

[11] C. E. Alchourrón, D. Makinson, On the logic of theory change: Safe contraction, Studia Logica 44 (1985) 405–422.

[12] J. S. R. Santos, Belief change without compactness, Ph.D. thesis, University of São Paulo, 2020.

[13] B. Nebel, Reasoning and Revision in Hybrid Representation Systems, volume 422 of *Lecture Notes in Computer Science*, Springer, 1990.

[14] S. E. Dixon, Belief revision: A computational approach, Ph.D. thesis, University of Sydney, 1994.

[15] M. Dalal, Investigations into a theory of knowledge base revision, in: Proceedings of the 7th National Conference on Artificial Intelligence, AAAI Press / The MIT Press, 1988, pp. 475–479.

[16] A. Fuhrmann, Theory contraction through base contraction, J. Philos. Log. 20 (1991) 175–203.

[17] S. O. Hansson, Belief contraction without recovery, Stud Logica 50 (1991) 251–260.

[18] S. O. Hansson, Reversing the levi identity, J. Philos. Log. 22 (1993) 637–669.

[19] S. O. Hansson, R. Wassermann, Local change, Stud Logica 70 (2002) 49–76.

[20] J. S. Ribeiro, Kernel contraction and the order of relevance, in: KR 2022, 2022, pp. 299–308.

[21] T. Caridroit, S. Konieczny, P. Marquis, Contraction in propositional logic, Int. J. Approx. Reason. 80 (2017) 428–442.

[22] F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, Cambridge University Press, 2017.

[23] T. A. Meyer, W. A. Labuschagne, J. Heidema, Refined epistemic entrenchment, J. Log. Lang. Inf. 9 (2000) 237–259.

[24] J. S. Ribeiro, A. Nayak, R. Wassermann, Belief update without compactness in non-finitary languages, in: IJCAI 2019, ijcai.org, 2019, pp. 1858–1864.

[25] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artif. Intell. 55 (1992) 1–60.