

Merging First-Order Knowledge using Dilation Operators

Nikos Gorogiannis and Anthony Hunter

Department of Computer Science,
University College London,
Gower Street, WC1E 6BT, London, UK
{n.gorogiannis,a.hunter}@cs.ucl.ac.uk

Abstract The area of knowledge merging is concerned with merging conflicting information while preserving as much as possible. Most proposals in the literature work with knowledge bases expressed in propositional logic. We propose a new framework for merging knowledge bases expressed in (subsets of) first-order logic. Dilation operators (a concept originally introduced by Bloch and Lang) are employed and developed, and by combining them with the concept of comparison orderings we obtain a framework that is driven by model-based intuitions but that can be implemented in a syntax-based manner. We demonstrate specific dilation operators and comparison orderings for use in applications. We also show how postulates from the literature on knowledge merging translate into our framework and provide the conditions that dilation operators and comparison orderings must satisfy in order for the respective merging operators to satisfy the new postulates.

1 Introduction

The dynamic nature of many information sources makes the appearance of inconsistencies a natural occurrence, either due to the need to consider multiple disparate sources, or because a single source is continually updated with new information. The need to study and handle these inconsistencies is well recognised and has spawned numerous research directions in non-monotonic logics, paraconsistent logics, belief revision and knowledge merging, to name a few. Logic-based knowledge merging is concerned with combining a set of potentially conflicting sources expressed in a logical language (*knowledge bases*), producing a single consistent knowledge base that maintains as much information content of the original knowledge bases as possible.

A *merging operator* is typically a function from tuples of knowledge bases (*profiles*) to knowledge bases. In recent years, the area of knowledge merging has seen significant research results, both in proposals for specific merging operators [1,2,3,4,5] as well as in proposals for unifying frameworks, where specific merging operators can be seen as instantiations of these frameworks [6,7,8,9].

Significantly, most of these results concern merging knowledge bases expressed in propositional logic (PL). If logic-based merging is to succeed in applications, then it is clear that more expressive languages need to be employed.

Although there exist proposals for knowledge merging with richer logics (mostly consistency-based, e.g. [1,10] but see also [11] for merging on infinite propositional logics), there is still a range of unresolved issues. In addition, we believe that computational issues should remain at the core of such investigations. In this paper, we examine the problems that arise when trying to construct a framework for merging knowledge expressed in first-order logic (FOL). We present a general framework that can be used when the language is (a subset of) FOL, and which makes use of the notion of a *dilation operator*, corresponding to the notion of distance employed in the knowledge merging frameworks such as the ones of [8,9]. The concept of a *comparison ordering* is then employed to order the result of applying the dilation operator to the original profile. The merging operators of our framework are, then, defined by minimising over these comparison orderings. The resulting framework has similarities to the work by Booth [12] in that repeated weakenings of the original profile are employed with the aim of reaching consensus; in contrast to that work, our framework is not limited to propositional logic, and does not explicitly use sets of models, making it computationally more appealing for richer logics. In addition, we instantiate this framework with specific dilation operators for use in different applications.

The outline of this paper is as follows. First, we cover some of the background literature in Section 2. Then, we introduce our approach in Section 3 and cover preliminary definitions in Section 4.1. We propose concrete dilation operators in Section 4.2 and comparison orderings in Section 4.3. We define the notion of a merging operator in Section 4.4 and present worked examples of concrete merging operators in Section 4.5. The general properties of merging operators in terms of merging postulates are examined in Section 4.6. Finally, we look at the issues that remain unaddressed and suggest further research topics in Section 5.

2 Background on Propositional Knowledge Merging

Merging operators from the literature can be broadly categorised in two groups: the *model-based*, or *semantic*, operators are defined using orderings over the models of PL, expressing how close models are with respect to the knowledge bases being merged [2,3,4,5]; and the *syntax-based*, or *consistency-based*, operators, where the merging process is defined in terms of consistent unions of subsets of the original knowledge bases [1,10,13]. We believe that the semantic approach offers several advantages such as syntax-independence and generality (e.g. sometimes syntax-based operators can be expressed as semantic ones), hence our approach is geared towards a semantic framework. In fact, we draw our intuitions from an important subset of the model-based operators, those that can be defined using a notion of *distance* between models (see, e.g. [9]).

We briefly cover the core concepts behind model-based merging operators. Let \mathcal{L} be a finite PL language (i.e. a logic with a finite number of atoms), with \mathcal{M} as its set of models and the usual operations $\text{mod} : \mathcal{L} \rightarrow 2^{\mathcal{M}}$ (which returns the set of models of a formula) and $\text{form} : 2^{\mathcal{M}} \rightarrow \mathcal{L}$ (which returns a formula whose set of models is the one provided). Knowledge bases are represented by

formulae (equivalent to the conjunction of the formulae in the knowledge base). The set of all profiles (tuples of consistent formulae of \mathcal{L}) is denoted by \mathcal{E} . The concatenation of two profiles E_1, E_2 is written $E_1 \sqcup E_2$ and the length of E_1 as $|E_1|$. Two profiles $E_1 = \langle \phi_1, \dots, \phi_k \rangle, E_2$ are equivalent, written $E_1 \leftrightarrow E_2$, iff $|E_1| = |E_2|$ and there is a bijection f from E_1 to E_2 such that $\phi_i \leftrightarrow f(\phi_i)$ for all $i \leq |E_1|$. The conjunction of all formulae in a tuple E is written $\bigwedge E$. We will abbreviate $E \sqcup \{\phi\}$ with $E \sqcup \phi$, and $\phi \sqcup \dots \sqcup \phi$ (n times) with ϕ^n . The lexicographic ordering on tuples of integers is defined as follows. For two tuples of integers $A = \langle a_1, \dots, a_k \rangle, B = \langle b_1, \dots, b_k \rangle$ we write $A <_{\text{lex}} B$ iff there exists $i \leq |A|$ such that $a_i < b_i$ and for all j such that $1 \leq j < i$, $a_j = b_j$; and we write $A \leq_{\text{lex}} B$ iff $A <_{\text{lex}} B$ or $A = B$. We denote by sort^a the function that takes a tuple of integers and sorts it in ascending order and by sort^d the function that takes a tuple of integers and sorts it in descending order.

For most model-based merging operators the concept of a *distance function* is central. A distance function $d : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{N}$ evaluates how “close” two models are, and has the properties that for any $\omega_1, \omega_2 \in \mathcal{M}$, $d(\omega_1, \omega_1) = 0$ (identity of indiscernibles) and $d(\omega_1, \omega_2) = d(\omega_2, \omega_1)$ (symmetry). On top of this distance, a notion of distance between a formula $\phi \in \mathcal{L}$ and a model $\omega \in \mathcal{M}$ can be defined by setting $d(\phi, \omega) = \min \{ d(\omega_1, \omega) \mid \omega_1 \in \text{mod}(\phi) \}$. Most model-based operators use this notion of distance to construct another distance between a whole profile and a model. Then, the merging operator is defined by minimising this distance over \mathcal{M} . Some examples from the literature are given below.

The operator Δ_{Max} [6]. The distance from a profile to a model for this operator is simply the maximum distance of each formula-to-model distance (where $E = \langle \phi_1, \dots, \phi_k \rangle$):

$$\begin{aligned} d_{\text{Max}}(E, \omega) &= \max \{ d(\phi_1, \omega), \dots, d(\phi_k, \omega) \} \\ \Delta_{\text{Max}}(E) &= \text{form} (\min \{ \omega \in \mathcal{M} \mid d_{\text{Max}}(E, \omega) \text{ is minimal} \}) \end{aligned}$$

The operator Δ_{Σ} [6]. Here, the distance from a profile to a model is the sum of the individual formula-to-model distances:

$$\begin{aligned} d_{\Sigma}(E, \omega) &= \sum_{i=1}^k d(\phi_i, \omega) \\ \Delta_{\Sigma}(E) &= \text{form} (\min \{ \omega \in \mathcal{M} \mid d_{\Sigma}(E, \omega) \text{ is minimal} \}) \end{aligned}$$

The operator Δ^{Gmax} [6]. In this case, the tuple of formula-to-model distances is sorted in ascending order and compared lexicographically.

$$\begin{aligned} d_{\text{Gmax}}(E, \omega) &= \text{sort}^a \langle d(\phi_1, \omega), \dots, d(\phi_k, \omega) \rangle \\ \Delta^{\text{Gmax}}(E) &= \text{form} (\min \{ \omega \in \mathcal{M} \mid d_{\text{Gmax}}(E, \omega) \text{ is minimal w.r.t. } \leq_{\text{lex}} \}) \end{aligned}$$

The operator Δ^{Gmin} [5]. Not dissimilarly with Δ^{Gmax} , the tuple of formula-to-model distances is sorted in *descending* order and compared lexicographically.

$$\begin{aligned} d_{\text{Gmin}}(E, \omega) &= \text{sort}^d \langle d(\phi_1, \omega), \dots, d(\phi_k, \omega) \rangle \\ \Delta^{\text{Gmin}}(E) &= \text{form} (\min \{ \omega \in \mathcal{M} \mid d_{\text{Gmin}}(E, \omega) \text{ is minimal w.r.t. } \leq_{\text{lex}} \}) \end{aligned}$$

- | |
|---|
| <p>A1. $\Delta(E)$ is consistent.</p> <p>A2. If E is consistent, then $\Delta(E) = \bigwedge E$.</p> <p>A3. If $E_1 \leftrightarrow E_2$, then $\vdash \Delta(E_1) \leftrightarrow \Delta(E_2)$.</p> <p>A4. If $\phi \wedge \psi$ is inconsistent, then $\Delta(\phi \sqcup \psi) \not\vdash \phi$.</p> <p>A5. $\Delta(E_1) \wedge \Delta(E_2) \vdash \Delta(E_1 \sqcup E_2)$.</p> <p>A6. If $\Delta(E_1) \wedge \Delta(E_2)$ is consistent, then $\Delta(E_1 \sqcup E_2) \vdash \Delta(E_1) \wedge \Delta(E_2)$.</p> |
|---|

Figure 1. Merging postulates.

There is a set of postulates, first proposed in [6] and listed in Figure 1, that merging operators are expected to satisfy. Postulate A1 simply means that a merging operator should produce consistent results. Postulate A2 states that in the case where a profile is already consistent, the operator will not alter it. Postulate A3 expresses the condition for syntax-independence. Postulate A4 is a fairness condition, requiring that the merging operator will not give preference to one knowledge base over another. Postulate A5 requires that if the two profiles, when merged separately, agree on a compromise then merging the concatenation of the profiles should include that compromise. Postulate A6 states that if there is agreement between two merged profiles, then the merged aggregate profile should entail all the consequences the two merged profiles agree on. Note that this list has been enhanced in later publications, such as [7,8], to include integrity constraints, i.e. formulae that must be consistent with the result of the merging; for simplicity of exposition, we have used the simpler version that does not employ integrity constraints.

Bloch and Lang [14] explore how some operations from mathematical morphology can be translated and used in logic. One of the core concepts is that of a dilation operator, a function from formulae to formulae, $D : \mathcal{L} \rightarrow \mathcal{L}$.

$$D(\phi) = \text{form}(\{\omega \in \mathcal{M} \mid d(\phi, \omega) \leq 1\})$$

Intuitively, dilating a formula yields another formula, the models of which are at a distance at most one from the models of the original formula. Dilations can be iterated: $D^{n+1}(\phi) = D(D^n(\phi))$ and $D^0(\phi) = \phi$. In [14], it is shown that $\Delta_{\text{Max}}(E)$ is equivalently defined by use of dilations:

$$\Delta_{\text{Max}}(E) = D^n(\phi_1) \wedge \dots \wedge D^n(\phi_k)$$

where n is the least number such that this conjunction is consistent.

Several syntax-based merging operators have been proposed [1,10,13]. These operators generally work by producing the maximal consistent subsets of the profile (maximising either with respect to subset inclusion or using set cardinality); and then taking the disjunction of (potentially, a subset of) the maximal consistent subsets. Although these operators are useful, they fail to satisfy some of the postulates for merging operators from the literature, and have been criticised for sacrificing too much information.

3 Motivation for Dilation-based Merging

It should be clear that using such model-based operators in a FOL setting can be problematic. The number of models of even the simplest FOL formulae is usually infinite and, therefore, several results from the literature that make use of the fact that the number of models of (finite) PL is finite will fail. Also, it may be that the set of models resulting from merging a profile cannot be represented by a formula, or even worse, not even by an infinite set of formulae (even though there always exists a set of formulae whose models are a superset of a given set X , e.g. the theory of that set, $\text{Th}(X)$).¹ Although there are ways to address this issue, e.g. by using maximal consistent sets of formulae instead of sets of models, as for example in [16], such approaches are not amenable computationally. Therefore, a compromise between model-based and syntax-based approaches that aims to address these issues is desirable. We will now look at how developing the idea of dilations can help achieve these aims.

The merging operators described previously can be expressed in an equivalent form using only dilations.

- $\Delta_{\text{Max}}(E) \leftrightarrow D^n(\phi_1) \wedge \dots \wedge D^n(\phi_k)$, where n is the least number such that this conjunction is consistent [14].
- $\Delta_{\Sigma}(E) \leftrightarrow \bigvee_{c_1+\dots+c_k=n} D^{c_1}(\phi_1) \wedge \dots \wedge D^{c_k}(\phi_k)$, where n is the least number such that this disjunction is consistent [17].
- $\Delta^{\text{Gmax}}(E) \leftrightarrow \bigvee_{\langle c_1, \dots, c_k \rangle \in \text{perm}(T)} D^{c_1}(\phi_1) \wedge \dots \wedge D^{c_k}(\phi_k)$, where T is the lexicographically-least tuple of integers that is sorted in ascending order for which this disjunction is consistent [17] ($\text{perm}(T)$ is the set of all tuples that are permutations of T).
- $\Delta^{\text{Gmin}}(E) \leftrightarrow \bigvee_{\langle c_1, \dots, c_k \rangle \in \text{perm}(T)} D^{c_1}(\phi_1) \wedge \dots \wedge D^{c_k}(\phi_k)$, where T is a tuple of integers that is sorted in *descending* order and is lexicographically least, such that this disjunction is consistent [17].

Example 1. Consider the profile $E = \langle p \wedge q, \neg p \wedge \neg q \rangle$. Clearly, $\bigwedge E$ is inconsistent. Using the dilation operator defined using Dalal’s notion of distance [18], we obtain:

$$\begin{aligned} D^1(p \wedge q) &\leftrightarrow (p \wedge q) \vee (p \wedge \neg q) \vee (\neg p \wedge q) \\ D^2(p \wedge q) &\leftrightarrow \top \\ D^1(\neg p \wedge \neg q) &\leftrightarrow (\neg p \wedge \neg q) \vee (\neg p \wedge q) \vee (p \wedge \neg q) \\ D^2(\neg p \wedge \neg q) &\leftrightarrow \top \end{aligned}$$

Clearly, the following conjunctions are consistent:

$$\begin{aligned} D^0(p \wedge q) \wedge D^2(\neg p \wedge \neg q) \\ D^1(p \wedge q) \wedge D^1(\neg p \wedge \neg q) \\ D^2(p \wedge q) \wedge D^0(\neg p \wedge \neg q) \end{aligned}$$

¹ See [15] for the notion of *definability preservation* in this context.

Therefore, the following hold:

$$\begin{aligned}
\Delta_{\text{Max}}(E) &\leftrightarrow D^1(p \wedge q) \wedge D^1(\neg p \wedge \neg q) \leftrightarrow (p \wedge \neg q) \vee (\neg p \wedge q) \\
\Delta_{\Sigma}(E) &\leftrightarrow ((D^0(p \wedge q) \wedge D^2(\neg p \wedge \neg q)) \vee \\
&\quad (D^1(p \wedge q) \wedge D^1(\neg p \wedge \neg q)) \vee \\
&\quad (D^2(p \wedge q) \wedge D^0(\neg p \wedge \neg q))) \leftrightarrow \top \\
\Delta^{\text{Gmax}}(E) &\leftrightarrow D^1(p \wedge q) \wedge D^1(\neg p \wedge \neg q) \leftrightarrow (p \wedge \neg q) \vee (\neg p \wedge q) \\
\Delta^{\text{Gmin}}(E) &\leftrightarrow (D^0(p \wedge q) \wedge D^2(\neg p \wedge \neg q)) \vee (D^2(p \wedge q) \wedge D^0(\neg p \wedge \neg q)) \leftrightarrow \\
&\quad (p \wedge q) \vee (\neg p \wedge \neg q)
\end{aligned}$$

The advantage of the above expressions is that, assuming a syntactic formulation for dilation is provided, they are entirely syntax-based and are, therefore, appealing in the context of a FOL merging framework without compromising the model-based properties of the original operators. In addition, a number of observations are in order.

- A dilation operator encapsulates a concept of distance in the sense that if a model is in $\text{mod}(D^n(\phi))$ and not in $\text{mod}(D^{n-1}(\phi))$, then it follows that the distance of the model to ϕ is n . In other words, given a dilation operator we can recover the corresponding formula-to-model distance. Also, a dilation operator can be seen as a *weakening* operator in that $D(\phi)$ represents a weakened version of ϕ , in the sense that $\phi \vdash D(\phi)$. This is especially appealing for knowledge merging, as the process of merging can be understood as a succession of compromises on the original points of view, until consensus is reached.
- In some sense, the merging operators presented above work by manipulating conjunctions of dilations of each of the formulae in a profile ($D^{c_1}(\phi_1) \wedge \dots \wedge D^{c_k}(\phi_k)$). This is important because it exemplifies that models are not the smallest unit necessary for building merging operators, and therefore, that in our move from PL to FOL we need not use models explicitly.
- All of these merging operators can be expressed as a disjunction of conjunctions of dilations of profile formulae (for Δ_{Max} it suffices to use the fact that $\phi \vdash D(\phi)$ to see this). The set of conjunctions is determined by an ordering which is used for minimisation. This ordering effectively works on tuples of numbers. In other words, for a given profile, a tuple of dilations $\langle D^{c_1}(\phi_1), \dots, D^{c_k}(\phi_k) \rangle$ corresponds to the *distance tuple* $\langle c_1, \dots, c_k \rangle$. In order to produce the merged knowledge base we only need to know how to find the minimal distance tuples among all those that correspond to consistent conjunctions of dilations of profile formulae. Therefore, orderings of distance tuples (which we call *comparison orderings*) encapsulate the nature of the merging operator independently of the dilation operator (or distance concept) used.

Whilst it is not impossible to use the model-based operators presented earlier in a FOL setting, there are shortcomings in doing so. An approach demonstrated in several papers [2,19,4] assumes a function-free FOL language and modifies

Dalal’s distance to one that counts the number of tuples that belong to the symmetric difference of the interpretations of a predicate. The models of the knowledge bases are, then, produced for *the minimum domain size* that is admitted by the number of constants in the language. The merging operator works by minimising over these sets of models, producing a subset of models at that domain size. This approach, we feel, restricts the expressive power of FOL too much as it is geared primarily towards databases, and so leads to problems when used more generally as illustrated by the following example. Suppose that there is one predicate, P , and that the profile we want to merge is $\langle \exists x P(x), \exists x \neg P(x) \rangle$. Since there are no constants, the merge will happen at domain size 1, where these formulae are obviously inconsistent, yielding a tautology, an unsatisfactory result. We can avoid this pitfall by skolemising, but at the potential cost of introducing functions, which the notion of distance proposed is unable to handle. Also, after such a merge, one is left with a set of models of a particular domain size, and therefore one that entails many more facts than were implicit in the original profile. In addition, we feel that using such a version of Dalal’s distance introduces a level of granularity that can be too fine-grained. Our framework, which is based on dilations, is a way around these problems.

An additional consideration is that the most frequently used notion of distance d between models in PL is the Dalal (or Hamming) distance [18,20]. In FOL it is much harder to come up with a single, all-encompassing notion of distance (indeed, even in PL, the Dalal distance suffers from well-known issues [21], e.g. some propositional letters may be more relevant to the merging process than others). To address this problem our framework admits several dilation operators, corresponding to different notions of distance.

Furthermore, PL distances like Dalal’s enjoy a certain property of *commensurateness* in the sense that the Dalal dilation of a formula changes it by a degree that is in a sense equal to the degree of change that results from the application of Dalal’s dilation to another formula (this condition corresponds to the symmetry condition for the distance function employed). This is especially hard to maintain in a FOL setting, as we will demonstrate in Section 4.2. Our approach, first, relaxes this constraint for dilation operators by letting them be relations as opposed to functions; and, second, changes the type of merging operators from a function that returns a single knowledge base, to a function that returns a set of alternative knowledge bases. The user, then, would assess the resulting set of candidate knowledge bases and act depending on the application. In other words, in order to get around the problem of lacking commensurateness, we increase the freedom of choice over what the merged profile should be.

4 Framework for Dilation-based Merging

4.1 Preliminaries

We will work with subsets of FOL, although the proposed framework should be usable with other classical logics such as modal logics. So, let $L = \langle \mathcal{L}, \vdash \rangle$ be a sound and complete logic, where \mathcal{L} is the set of formulae of the logic, and

$\vdash \subseteq \mathcal{L} \times \mathcal{L}$ is the (transitive and reflexive) entailment relation. Two formulae are equivalent, $\phi \equiv \psi$, iff $\phi \vdash \psi$ and $\psi \vdash \phi$. We assume that \mathcal{L} has the connectives \wedge , \vee and \neg and that they behave classically. A formula $\phi \in \mathcal{L}$ is consistent iff $\phi \not\vdash \perp$.

We use the notation for profiles introduced previously with the following additions. The notation for concatenation of tuples of formulae, \sqcup , repeated concatenation, \cdot^n , and the length of a tuple of formulae, $|\cdot|$, will also be used for tuples of numbers. A profile is a tuple of consistent formulae of \mathcal{L} and denoted by E . The set of all profiles of finite length is denoted by \mathcal{E} . Two profiles $E_1 = \langle \phi_1, \dots, \phi_k \rangle, E_2$ are equivalent, written $E_1 \equiv E_2$, iff $|E_1| = |E_2|$ and there is a bijection f from E_1 to E_2 such that $\phi_i \equiv f(\phi_i)$ for all $i \leq |E_1|$. We also use the \equiv notation for sets of formulae with the corresponding meaning.

If R is a binary relation over a set S and $x \in S$ then we denote the image of x through R as $R(x)$, i.e. $R(x) = \{y \in S \mid \langle x, y \rangle \in R\}$. We extend this notation over sets of elements, i.e. if $X \subseteq S$ then $R(X) = \{y \in S \mid \exists x \in X, \langle x, y \rangle \in R\}$. We define $R^n(x)$ inductively as $R^1(x) = R(x)$ and $R^{n+1}(x) = R^n(R(x))$ and use the natural extension of this notation over sets, e.g. $R^n(X)$.

If $A = \langle n_1, \dots, n_k \rangle$ is a tuple of natural numbers, then $\sum A = \sum_{i=1}^k n_i$ and $\max A = \max \{n_i \mid 1 \leq i \leq k\}$.

4.2 Dilation Operators

Similarly to the literature on propositional knowledge merging, we will list some postulates that the concepts we define may satisfy. However, we do not make these postulates a part of the definitions as we want to highlight the dependencies of the results as well as to allow more freedom as to which postulates a particular instantiation of the definitions happens to satisfy. Note that a further modification to the notion of a dilation is that we set it to depend on the profile in question. We do this to allow a dilation operator to perform a guided weakening to the formulae involved; given the difficulty of ensuring commensuratensness for dilation, this may allow for more effective definitions of concrete dilation operators.

Definition 1. A dilation operator is a relation $D_E \subseteq \mathcal{L} \times \mathcal{L}$, where $E \in \mathcal{E}$.

A dilation operator D may satisfy the following postulates.

- D1. If $\phi \equiv \psi$ then $D_E(\phi) \equiv D_E(\psi)$ (syntax independence).
- D2. If $\chi \in D_E(\phi)$ then $\phi \vdash \chi$ (weakening).
- D3. Let $E = \langle \phi, \psi \rangle$. For all $a \in \mathbb{N}$ and $\phi' \in D_E^a(\phi)$ such that $\phi' \wedge \psi$ is consistent, there exists $b \leq a$ and $\psi' \in D_E^b(\psi)$ such that $\phi \wedge \psi' \not\vdash \perp$ (commensuratensness).

Postulate D1 is a standard syntax-independence postulate. Postulate D2 is important because it demands that the dilation operator is, in effect, a subset of the entailment relation. Postulate D3 corresponds to the symmetry condition on distances in the sense that if a certain number of dilations of ϕ is enough

to make it consistent with ψ then it should be the case that an equal or lower number of dilations of ψ will make it consistent with ϕ .

We list, below, several concrete dilation operators. We begin with a trivial dilation operator that corresponds to the *drastic distance* from the PL literature.

Definition 2. Let \mathcal{L} be a FOL fragment. Then for all $\phi \in \mathcal{L}$,

$$D_E^{\text{drastic}}(\phi) = \{\top\}.$$

This operator can be used to reconstruct some of the consistency-based operators from the literature and we will show specific examples of that in Section 4.5. It is easy to see that postulates *D1*, *D2* and *D3* trivially hold.

Another simple dilation operator is obtained by taking the disjunction of the formula being dilated with some other formula from the profile.

Definition 3. Let \mathcal{L} be a FOL fragment which is closed under disjunction. Then, if $E = \langle \psi_1, \dots, \psi_k \rangle$,

$$D_E^{\text{disj}}(\phi) = \{\phi \vee \psi_i \mid 1 \leq i \leq k \text{ and } \psi_i \not\vdash \phi\}$$

Again, it is trivial to see that *D1*, *D2* and *D3* are satisfied. As an example, if $E = \langle a \rightarrow b, a, \neg b, c, d \rangle$ then,

$$\begin{aligned} D_E^{\text{disj}}(a \rightarrow b) &= \{(a \rightarrow b) \vee a, (a \rightarrow b) \vee \neg b, (a \rightarrow b) \vee c, (a \rightarrow b) \vee d\} \\ &\equiv \{\top, (a \rightarrow b) \vee c, (a \rightarrow b) \vee d\} \end{aligned}$$

$$D_E^{\text{disj}}(D_E^{\text{disj}}(a \rightarrow b)) \equiv \{\top, (a \rightarrow b) \vee c, (a \rightarrow b) \vee d, (a \rightarrow b) \vee c \vee d\}$$

Finally, the following definition concerns a dilation operator that weakens formulae by changing universal quantifiers to existential. Its effect on the input formulae is rather profound, and as such it may be suitable only for specific applications. We present a more illuminating worked example in Section 4.4.

Definition 4. Let \mathcal{L} be a (prenex) FOL fragment and ψ a formula of the form $\psi = Q_1 x_1 \dots Q_n x_n \phi$ where Q_i is a quantifier and ϕ is quantifier-free.

$$D_E^Q(\psi) = \left\{ Q'_1 x_1 \dots Q'_n x_n \phi \left| \begin{array}{l} \text{there exists } j \leq n \text{ such that} \\ Q_j = \forall \text{ and } Q'_j = \exists \text{ and} \\ \text{for all } i \leq n \text{ such that } i \neq j, Q'_i = Q_i \end{array} \right. \right\}$$

As an example, for some profile E ,

$$\begin{aligned} D_E^Q(\forall x \forall y P(x, y)) &= \{\forall x \exists y P(x, y), \exists x \forall y P(x, y)\} \\ D_E^Q(D_E^Q(\forall x \forall y P(x, y))) &= \{\exists x \exists y P(x, y)\}. \end{aligned}$$

Postulate *D2* is trivially satisfied. Syntax independence (postulate *D1*) fails, since $\forall x \forall y P(x, y) \equiv \forall y \forall x P(x, y)$ but, $D_E^Q(\forall y \forall x P(x, y)) \not\equiv D_E^Q(\forall x \forall y P(x, y))$ as can be seen below.

$$D_E^Q(\forall y \forall x P(x, y)) = \{\forall y \exists x P(x, y), \exists y \forall x P(x, y)\}$$

Commensurateness (D3) also fails. Consider the formulae $\phi = \forall x \forall y P(x, y)$ and $\psi = \exists x \forall y \neg P(x, y)$. Clearly, $\phi \wedge \psi$ is inconsistent and there exists a formula in $D_E^Q(\phi)$, namely $\chi = \exists x \forall y P(x, y)$ such that $\chi \wedge \psi$ is consistent. However, the only formula in $D_E^Q(\exists x \forall y \neg P(x, y))$ is $\exists x \exists y \neg P(x, y)$ which is not consistent with ϕ .

If we restrict the logical language to propositional logic, then we can define a dilation operator that exactly captures Dalal's dilation. Supposing that ϕ is written in Disjunctive Normal Form (DNF), i.e. $\phi = \bigvee_{i=1}^l \bigwedge_{j=1}^{m_i} L_{i,j}$ with $L_{i,j}$ being literals, then it can be shown (e.g. [14]) that

$$D_E^{\text{Dalal}}(\psi) = \left\{ \bigvee_{i=1}^l \bigvee_{j=1}^{m_i} \bigwedge_{\substack{k=1 \\ k \neq j}}^{m_i} L_{i,k} \right\}.$$

This dilation operator allows us to subsume the PL merging operators from the literature within our framework.

4.3 Comparison Orderings

The notion of a comparison ordering is meant to capture the nature of a merging operator independently of the particular dilation operator used in its definition. Since dilating a formula constitutes a discrete step, the purpose of a comparison ordering is to evaluate how much compromise is involved in making a profile consistent by applying the dilation operator to the profile formulae a number of times. In other words, the comparison ordering is meant to order distance tuples in how far they are from the original profile. This concept is partly related to the orderings over models used in frameworks such as the one in [8] (syncretic assignments) but also to the aggregation functions of [9].

Definition 5. A comparison ordering is a collection of total preorders over tuples of natural numbers, one for each length k , $\sqsubseteq \subseteq \mathbb{N}^k \times \mathbb{N}^k$.

Comparison orderings may satisfy the following postulates.

- C1. For any $\langle n_1, \dots, n_k \rangle \in \mathbb{N}^k$, $\langle 0, \dots, 0 \rangle \sqsubseteq \langle n_1, \dots, n_k \rangle$. Furthermore, if there exists $i \leq k$ such that $n_i > 0$, then $\langle 0, \dots, 0 \rangle \sqsubset \langle n_1, \dots, n_k \rangle$
- C2. If A_1, A_2 and B_1, B_2 are pairs of tuples of equal respective lengths, then $A_1 \sqsubseteq A_2$ and $B_1 \sqsubseteq B_2$ implies $A_1 \sqcup B_1 \sqsubseteq A_2 \sqcup B_2$.
- C3. If A_1, A_2 and B_1, B_2 are pairs of tuples of equal respective lengths, then $A_1 \sqsubset A_2$ and $B_1 \sqsubseteq B_2$ implies $A_1 \sqcup B_1 \sqsubset A_2 \sqcup B_2$.
- C4. For all $a, b \in \mathbb{N}$, if $a \leq b$, then $\langle a, 0 \rangle \sqsubseteq \langle 0, b \rangle$.

Postulate C1 corresponds to the conditions of *non-decreasingness* and *minimality* from [9] and to the postulate A2 in a way that will become obvious after we define what a merging operator is, in Definition 11. Postulates C2 and C3 correspond to the fifth and sixth conditions for syncretic assignments respectively [8].

Note that $C3$ entails $C2$. These two postulates allow us to infer facts about the relation between two concatenations of tuples when the corresponding relations are known, but there is no way to infer facts about concatenations in different permutations. Postulate $C4$ fills in this gap and is related to commensurateness.

The PL merging operators listed in Section 2 can be translated into our framework as follows.

Definition 6. *The comparison ordering \sqsubseteq_{Max} is defined as follows.*

$$\langle a_1, \dots, a_k \rangle \sqsubseteq_{\text{Max}} \langle b_1, \dots, b_k \rangle \text{ iff } \max_{1 \leq i \leq k} a_i \leq \max_{1 \leq i \leq k} b_i$$

It is easy to see that the \sqsubseteq_{Max} satisfies $C1$. It also satisfies $C2$ since if $\max A_1 \leq \max A_2$ and $\max B_1 \leq \max B_2$ then it will be the case that $\max A_1 \sqcup A_2 \leq \max B_1 \sqcup B_2$. Note, however, that \sqsubseteq_{Max} does not satisfy $C3$, because if $\max A_1 < \max A_2$ but $\max B_1 = \max B_2$ we cannot conclude that $\max A_1 \sqcup A_2 < \max B_1 \sqcup B_2$. Finally, it is easy to see that \sqsubseteq_{Max} satisfies $C4$.

Definition 7. *The comparison ordering \sqsubseteq_{Σ} is defined as follows.*

$$\langle a_1, \dots, a_k \rangle \sqsubseteq_{\Sigma} \langle b_1, \dots, b_k \rangle \text{ iff } \sum \langle a_1, \dots, a_k \rangle \leq \sum \langle b_1, \dots, b_k \rangle$$

The ordering \sqsubseteq_{Σ} clearly satisfies $C1$. It is also easy to see that it satisfies $C3$ (and therefore $C2$ as well): if $\sum A_1 < \sum A_2$ and $\sum B_1 \leq \sum B_2$ then $\sum A_1 \sqcup B_1 = \sum A_1 + \sum B_1 < \sum A_2 + \sum B_1 \leq \sum A_2 + \sum B_2 = \sum A_2 \sqcup B_2$. It is clear that \sqsubseteq_{Σ} satisfies $C4$ as well.

Definition 8. *The comparison ordering $\sqsubseteq_{\text{Gmax}}$ is defined as follows.*

$$\langle a_1, \dots, a_k \rangle \sqsubseteq_{\text{Gmax}} \langle b_1, \dots, b_k \rangle \text{ iff } \text{sort}^d \langle a_1, \dots, a_k \rangle \leq_{\text{lex}} \text{sort}^d \langle b_1, \dots, b_k \rangle$$

It is clear that Δ^{Gmax} satisfies $C1$. That it satisfies $C3$ follows from a result presented in [6], namely that if A_1, A_2 and B_1, B_2 are tuples of numbers of equal lengths, and $\text{sort}^d A_1 <_{\text{lex}} \text{sort}^d A_2$ and $\text{sort}^d B_1 \leq_{\text{lex}} \text{sort}^d B_2$ then $\text{sort}^d A_1 \sqcup B_1 <_{\text{lex}} \text{sort}^d A_2 \sqcup B_2$. It is also easy to check that $\sqsubseteq_{\text{Gmax}}$ satisfies $C4$.

Definition 9. *The comparison ordering $\sqsubseteq_{\text{Gmin}}$ is defined as follows.*

$$\langle a_1, \dots, a_k \rangle \sqsubseteq_{\text{Gmin}} \langle b_1, \dots, b_k \rangle \text{ iff } \text{sort}^a \langle a_1, \dots, a_k \rangle \leq_{\text{lex}} \text{sort}^a \langle b_1, \dots, b_k \rangle$$

Again, it is obvious that $\sqsubseteq_{\text{Gmin}}$ satisfies $C1$. In similar manner to that of the result on $\sqsubseteq_{\text{Gmax}}$ we can show that $\sqsubseteq_{\text{Gmin}}$ satisfies $C3$ as well. Also, it is clear that $C4$ holds of $\sqsubseteq_{\text{Gmin}}$ as well.

4.4 Merging Operators

As noted in Section 3, the smallest unit that we can distinguish using a profile and a dilation operator is the conjunction of dilations of the profile formulae. This set of candidate tuples of dilations is captured by the next definition. Note that we discard tuples that yield inconsistent conjunctions as they represent a set of compromises that has not reached consensus, and therefore cannot form part of the merging.

Definition 10. The function $\mathcal{C}_D : \mathcal{E} \rightarrow 2^{\mathcal{E}}$ generates the set of all consistent combinations of dilations of a profile, using the dilation operator D_E . Let $E = \langle \phi_1, \dots, \phi_k \rangle$.

$$\langle \psi_1, \dots, \psi_k \rangle \in \mathcal{C}_D(E) \text{ iff } \begin{cases} \psi_1 \wedge \dots \wedge \psi_k \text{ is consistent and} \\ \text{for all } i \leq k \text{ there exists } n_i, \text{ such that } \psi_i \in D_E^{n_i}(\phi_i) \end{cases}$$

In effect, \mathcal{C}_D generates the set of all possible tuples that may play a role in the merging process.

For a profile $E = \langle \phi_1, \dots, \phi_k \rangle$ of length k , the function $dt : \mathcal{C}_D(E) \rightarrow \mathbb{N}^k$ (*distance tuple*) returns the tuple of the (minimum) numbers of dilations of each profile formula:

$$dt(\langle \psi_1, \dots, \psi_k \rangle) = \langle n_1, \dots, n_k \rangle \text{ iff } \begin{cases} \text{for all } i \leq k, n_i \text{ is the least number} \\ \text{such that } \psi_i \in D_E^{n_i}(\phi_i) \end{cases}$$

We can now give the definition of a merging operator, using the notions of a dilation operator and comparison ordering.

Definition 11. A merging operator is a function $\Delta_{D, \sqsubseteq} : \mathcal{E} \rightarrow 2^{\mathcal{L}}$ defined in terms of a dilation operator D and a comparison ordering \sqsubseteq in the following manner.

$$\bigwedge A \in \Delta_{D, \sqsubseteq}(E) \text{ iff } \begin{cases} A \in \mathcal{C}_D(E) \text{ and} \\ dt(A) \sqsubseteq dt(B) \text{ for all } B \in \mathcal{C}_D(E) \end{cases}$$

Note that a merging operator according to this definition is a function that returns a set of formulae instead of a single formula. This is because the dilation operators are now relations and this opens up the question of how to treat the resulting conjunctions. In PL, merging operators form the union of the corresponding sets of models. This is not appropriate here, since the union may discard too much information. As an example suppose that for any $\psi \in D(\phi)$, $\psi \not\equiv \top$, but $\bigvee D(\phi) \equiv \top$. To address this, instead of taking the disjunction of the formulae in the result of the merging, the user is free to choose which operation to apply according to what makes sense in the context of the application in question. For example, the original behaviour of PL merging operators can be recovered by taking the disjunction of the members of this set.

In the following, we will abuse notation slightly and say that a tuple $A \in \mathcal{C}_D(E)$ is \sqsubseteq -minimal when $dt(A)$ is actually \sqsubseteq -minimal. We will also omit \sqsubseteq and D as subscripts of a merging operator when they are clear from the context.

4.5 Examples of Dilation-based Merging

In this section we demonstrate some concrete merging operators by combining specific dilation operators with the comparison orderings presented in Section 4.3. We begin with the drastic dilation and show how some consistency-based merging operators can be subsumed by our framework.

Example 2. Suppose $E = \langle a \rightarrow b, a, \neg b, c \rangle$. It is easy to see that $\bigwedge E$ is inconsistent and that any two formulae of $\{a \rightarrow b, a, \neg b\}$ are consistent with c . Hence,

$$\mathcal{C}_{D_E^{\text{drastic}}}(E) = \left\{ \begin{array}{l} \langle \top, a, \neg b, c \rangle, \\ \langle a \rightarrow b, \top, \neg b, c \rangle, \\ \langle a \rightarrow b, a, \top, c \rangle \\ \vdots \end{array} \right\}.$$

Using this dilation operator and the \sqsubseteq_{Max} comparison ordering we obtain the following result.

$$\Delta_{\sqsubseteq_{\text{Max}}}(E) = \left\{ \bigwedge A \mid A \in \mathcal{C}_{D_E^{\text{drastic}}}(E) \right\}$$

Obviously, this operator is very coarse in the comparisons it makes and that is a direct result of the deficiencies of \sqsubseteq_{Max} . PL merging operators work by taking the disjunction of the returned set (e.g. $\bigvee \psi$ for all $\psi \in \Delta_{\sqsubseteq_{\text{Max}}}(E)$), which in our case would yield a tautology.

Using \sqsubseteq_{Σ} , $\sqsubseteq_{\text{Gmax}}$ or $\sqsubseteq_{\text{Gmin}}$ we obtain the same intuitive result corresponding to the maximal consistent subsets of the profile E :

$$\Delta_{\sqsubseteq_{\Sigma}}(E) = \Delta_{\sqsubseteq_{\text{Gmax}}}(E) = \Delta_{\sqsubseteq_{\text{Gmin}}}(E) = \left\{ \begin{array}{l} a \wedge \neg b \wedge c, \\ a \rightarrow b \wedge \neg b \wedge c, \\ a \rightarrow b \wedge a \wedge c \end{array} \right\}.$$

The disjunction of the members of this set is equivalent to $(b \rightarrow a) \wedge c$.

It is easy to see that, when using the drastic dilation operator, the three comparison orderings \sqsubseteq_{Σ} , $\sqsubseteq_{\text{Gmax}}$ or $\sqsubseteq_{\text{Gmin}}$ yield the same merging operator, which selects the cardinality-maximal consistent subsets of the original profile. This is, effectively, one of the consistency-based merging operators found in [1,10,13].

Example 3. Suppose $E = \langle \forall x \forall y P(x, y) \leftrightarrow Q(y, x), \forall x P(x, x) \leftrightarrow \neg Q(x, x) \rangle$. The first formula effectively forces the relation Q to be the transpose of the relation P . The second formula prescribes that, if we view P and Q as incidence matrices, that they disagree on their diagonals. Clearly, E is inconsistent, since the first formula entails the negation of the second. Using the dilation operator D_E^Q , we obtain:

$$\begin{aligned} D_E^Q(\forall x \forall y P(x, y) \leftrightarrow Q(y, x)) &= \left\{ \begin{array}{l} \forall x \exists y P(x, y) \leftrightarrow Q(y, x), \\ \exists x \forall y P(x, y) \leftrightarrow Q(y, x) \end{array} \right\} \\ D_E^Q(D_E^Q(\forall x \forall y P(x, y) \leftrightarrow Q(y, x))) &= \{ \exists x \exists y P(x, y) \leftrightarrow Q(y, x) \} \\ D_E^Q(\forall x P(x, x) \leftrightarrow \neg Q(x, x)) &= \{ \exists x P(x, x) \leftrightarrow \neg Q(x, x) \}. \end{aligned}$$

Choosing any comparison ordering out of \sqsubseteq_{Σ} , $\sqsubseteq_{\text{Gmax}}$, $\sqsubseteq_{\text{Gmin}}$ will yield the same merged set, i.e.

$$\Delta_{\Sigma}(E) = \Delta^{\text{Gmax}}(E) = \Delta^{\text{Gmin}}(E) = \left\{ \bigwedge \left\{ \begin{array}{l} \forall x \exists y P(x, y) \leftrightarrow Q(y, x), \\ \forall x P(x, x) \leftrightarrow \neg Q(x, x) \end{array} \right\} \right\}$$

The result entails that for each pair of elements in P there is a transposed pair in Q and that the two relations disagree on their diagonals.

It is clear that a merging operator based on D_E^Q will have a very strong effect on the input profile and, therefore, may be used in limited contexts.

Next, we look at a dilation operator that weakens universal formulae by adding exceptions using the constants of the language.

Definition 12. Assume \mathcal{L} is a (prenex) universal FOL fragment with a finite number of constants c_1, \dots, c_n . Let $\psi \in \mathcal{L}$ be of the form $\forall x_1 \dots \forall x_k \phi$.

$$\begin{aligned} exc_{i,j}(\phi) &= \forall x_1 \dots \forall x_k (x_i \neq c_j \rightarrow \phi) \\ D_E^{exc}(\phi) &= \{ exc_{i,j}(\phi) \mid 1 \leq i \leq k, 1 \leq j \leq n \text{ and } exc_{i,j}(\phi) \not\vdash \perp \} \end{aligned}$$

It should be easy to see that postulates $D1$ and $D2$ clearly hold.

Example 4. Let us assume a language with three constants, a, b, c and the following profile $E = \langle \forall y P(c, y), \forall y P(a, y), \forall x \neg P(x, b) \rangle$. The dilation operator applied to the profile formula $\forall y P(c, y)$ gives:

$$D_E^{exc}(\forall y P(c, y)) = \left\{ \begin{array}{l} \forall y (y \neq a \rightarrow P(c, y)), \\ \forall y (y \neq b \rightarrow P(c, y)), \\ \forall y (y \neq c \rightarrow P(c, y)) \end{array} \right\}.$$

Similar results are obtained by dilating the rest of the profile formulae. By using $\sqsubseteq_{\Sigma}, \sqsubseteq_{G_{\max}}, \sqsubseteq_{G_{\min}}$ we obtain the following results:

$$\begin{aligned} \Delta_{\sqsubseteq_{G_{\max}}}(E) &= \left\{ \bigwedge \left\{ \begin{array}{l} \forall y (y \neq b \rightarrow P(c, y)), \\ \forall y (y \neq b \rightarrow P(a, y)), \\ \forall x \neg P(x, b) \end{array} \right\} \right\} \\ \Delta_{\sqsubseteq_{G_{\min}}}(E) &\equiv \left\{ \bigwedge \left\{ \begin{array}{l} \forall y P(c, y), \\ \forall y P(a, y), \\ \forall x (x \neq a \wedge x \neq c) \rightarrow \neg P(x, b) \end{array} \right\} \right\} \end{aligned}$$

$$\Delta_{\sqsubseteq_{\Sigma}}(E) \equiv \Delta_{\sqsubseteq_{G_{\max}}}(E) \cup \Delta_{\sqsubseteq_{G_{\min}}}(E)$$

The operator $\Delta_{\sqsubseteq_{\Sigma}}$ is selecting dilated profiles based on the total number of dilation operations applied (here the minimum for a consistent result is 2), $\Delta_{\sqsubseteq_{G_{\max}}}$ sorts the distance tuples in descending order and therefore selects the dilated profiles that have (any permutation of) $\langle 1, 1, 0 \rangle$ as their distance tuple. Finally, $\Delta_{\sqsubseteq_{G_{\min}}}$ sorts in ascending order and, thus, selects the dilated profiles that correspond to the distance tuple $\langle 2, 0, 0 \rangle$ and its permutations.

We note that if the language \mathcal{L} does not allow function symbols and the input profile does not contain existential quantifiers, then this merging operator can

be equivalently defined using the notion of Dalal’s distance on the symmetrical set differences of the interpretation of the predicates, as explained in Section 3 in relation to the papers [2,19,4]. However, if we allow function symbols in the language and/or existential quantifiers in the input profile, then our dilation-based merging operator becomes distinct from the approach described in the papers cited above.

4.6 Properties of Merging Operators

Obviously, the list of postulates presented in Figure 1 is not appropriate for operators defined through Definition 11. To address this, we present a list in Figure 2 that we believe captures the intention of the original postulates and fits with our framework. We comment on the modified postulates below. Note that most postulates are essentially quantified versions of the originals, over the set returned by the merging operator.

- B1.* For any $\phi \in \Delta(E)$, ϕ is consistent.
- B2.* If $\bigwedge E$ is consistent, then $\Delta(E) \equiv \{\bigwedge E\}$.
- B3.* If $E_1 \equiv E_2$, then $\Delta(E_1) \equiv \Delta(E_2)$.
- B4.* If $\phi \wedge \psi$ is inconsistent, then there exists $\chi \in \Delta(\phi \sqcup \psi)$ such that $\chi \not\vdash \phi$.
- B5.* If $\phi_1 \in \Delta(E_1)$, $\phi_2 \in \Delta(E_2)$ and $\phi_1 \wedge \phi_2 \not\vdash \perp$, then $\phi_1 \wedge \phi_2 \in \Delta(E_1 \sqcup E_2)$.
- B6.* If there exist $\phi_1 \in \Delta(E_1)$, $\phi_2 \in \Delta(E_2)$ such that $\phi_1 \wedge \phi_2$ is consistent, then for any $\chi \in \Delta(E_1 \sqcup E_2)$ there are $\psi_1 \in \Delta(E_1)$ and $\psi_2 \in \Delta(E_2)$ such that $\chi = \psi_1 \wedge \psi_2$.

Figure 2. Merging postulates for our framework.

Postulate *B1* is a straightforward translation of *A1* requiring that any member of $\Delta(E)$ is consistent. Similarly, postulates *B2*, *B3* correspond directly to *A2* and *A3*. Postulate *B4* corresponds to *A4* but requires some explanation; since the PL merging operators can be expressed as disjunctions of conjunctions of dilations, in order to satisfy *A4* it suffices to have a conjunction of dilations that does not entail ϕ and this is why *B4* is existentially quantified. In order to translate *A5* and *A6* we interpret the conjunction of two merged profiles as an operation that produces all consistent conjunctions from the two sets $\Delta(E_1)$ and $\Delta(E_2)$. In other words, we translate $\Delta(E_1) \wedge \Delta(E_2)$ into the set $\{\phi_1 \wedge \phi_2 \mid \phi_1 \in \Delta(E_1), \phi_2 \in \Delta(E_2), \phi_1 \wedge \phi_2 \not\vdash \perp\}$, over which the postulates *B5* and *B6* are quantified.

Using Definitions 1, 10, 5 and 11 we can now explore the conditions under which a merging operator will satisfy the postulates listed in Figure 2.

Proposition 1. *Any operator Δ defined through Definition 11 in terms of a dilation operator D and a comparison ordering \sqsubseteq will satisfy the following postulates, under the listed conditions.*

- If D satisfies $D2$ then Δ satisfies $B1$.
- If \sqsubseteq satisfies $C1$ then Δ satisfies $B2$.
- If D satisfies $D1$ then Δ satisfies $B3$.
- If D satisfies $D3$ and \sqsubseteq satisfies $C1$, $C2$ and $C4$, then Δ satisfies $B4$.
- If \sqsubseteq satisfies $C2$ then Δ satisfies $B5$.
- If \sqsubseteq satisfies $C3$ then Δ satisfies $B6$.

The proof for Proposition 1 is included in the appendix.

So, returning to Example 2, we can see that the merging operators $\Delta_{\sqsubseteq_{\Sigma}}(E)$, $\Delta_{\sqsubseteq_{G_{\max}}}(E)$ and $\Delta_{\sqsubseteq_{G_{\min}}}(E)$, when based on the drastic dilation operator, will satisfy all the above postulates. On the other hand, in Example 4, the operators $\Delta_{\sqsubseteq_{\Sigma}}(E)$, $\Delta_{\sqsubseteq_{G_{\max}}}(E)$ and $\Delta_{\sqsubseteq_{G_{\min}}}(E)$, when based on D_E^{exc} , will satisfy all the postulates apart from $B4$ since it is an open question as to whether D_E^{exc} satisfies $D3$.

5 Conclusions

We have outlined a dilation-based framework for knowledge merging that is appropriate for more expressive logics than propositional logic. This framework is flexible, admitting different notions of compromise in the form of dilation operators, as well as allowing different ways of minimising the extent of compromise in the form of comparison orderings. In addition, it can subsume important proposals from the literature, such as some kinds of consistency-based merging, as well as the PL merging operators based on Dalal’s distance.

We believe that the framework and its specific instantiations presented, avoid the problems described in Section 3. Artefacts that result from using variants of Dalal’s distance on the interpretations of predicates, such as the exceedingly fine-grained nature of the resulting knowledge base or the inability to allow function symbols in the language, are avoided. In addition, merging operators that return a multiplicity of weakened versions of a profile are allowed, enabling the user of the merging operator to make choices that may depend on extra-logical information, or information that is not available to the merging operator. Furthermore, the user can compute the disjunction of the resulting set of profiles, thus recovering the behaviour of many of the merging operators in the literature.

There are many possibilities for additional dilation operators. One that we will explore in the future is the following. The intuition behind it is to try and capture some of the characteristics of the Dalal dilation, but in a FOL setting.

Definition 13. Let \mathcal{L} be a (prenex) universal FOL fragment and $\psi \in \mathcal{L}$ a formula of the form $Q_1x_1 \cdots Q_nx_n\phi$ where Q_i is a quantifier and ϕ is quantifier-free and in DNF form. This means that $\phi = \bigvee_{i=1}^l \bigwedge_{j=1}^{m_i} L_{i,j}$ with $L_{i,j}$ being possibly negated atomic formulae.

$$D_E^{\text{Dalal}}(\psi) = \left\{ Q_1x_1 \cdots Q_nx_n \bigvee_{i=1}^l \bigvee_{j=1}^{m_i} \bigwedge_{\substack{k=1 \\ k \neq j}}^{m_i} L_{i,k} \right\}$$

Several questions remain open. For example, are there conditions that can capture the essence behind arbitration and majority operators, as defined in the literature (see, e.g. [6])? Also, the properties we have outlined are sufficient to force a merging operator to satisfy the postulates listed in Figure 2, but can this result can be strengthened towards a characterisation theorem?

In addition, although we believe that the format of our definitions lends itself nicely to implementation, decidability is obviously a deeper question, especially with regards to the exact logical language chosen. For example, our definition of dilation allows for an infinity of weakened versions of a formula to be returned, something that would obviously complicate computation, even though all the examples of dilation operators we presented do not suffer from this issue. Also related to this is the possibility of infinite chains of weakenings that never produce consensus in the original profile (this is interrelated to the existence of infinitely-descending chains of models in frameworks that employ distances or orderings). While this is not a definitional problem of the framework $(\mathcal{C}_D(E)$ and $\Delta(E)$ will both simply be empty), it would be interesting to designate the conditions under which this does not happen.

Other open issues include whether we can define a distance function between models based on a syntactic construction such as a dilation operator, similar to the work done by Lehmann et al on belief revision [15]. Finally, even though the use of integrity constraints is very appealing for purposes like expressing background knowledge, for simplicity of exposition, we have omitted their treatment. Extending the framework to allow integrity constraints to be expressed and handled appropriately would be another possible research direction.

References

1. Baral, C., Kraus, S., Minker, J.: Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering* **3**(2) (1991) 208–220
2. Revesz, P.Z.: On the semantics of arbitration. *Journal of Algebra and Computation* **7**(2) (1997) 133–160
3. Liberatore, P., Schaerf, M.: Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering* **10**(1) (1998) 76–90
4. Lin, J., Mendelzon, A.O.: Merging databases under constraints. *International Journal of Cooperative Information Systems* **7**(1) (1998) 55–76
5. Everaere, P., Konieczny, S., Marquis, P.: Quota and Gmin merging operators. In: Nineteenth International Joint Conference on Artificial Intelligence, IJCAI'05. (2005) 424–429
6. Konieczny, S., Pérez, R.P.: On the logic of merging. In: Sixth International Conference on Principles of Knowledge Representation and Reasoning, KR'98. (1998) 488–498
7. Konieczny, S., Pérez, R.P.: Merging with integrity constraints. In: Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, ECSQARU'95. Volume 1638 of Lecture Notes in Computer Science., Springer-Verlag (1999) 233–244
8. Konieczny, S., Pérez, R.P.: Merging Information Under Constraints: A Logical Framework. *Journal of Logic and Computation* **12**(5) (2002) 773–808

9. Konieczny, S., Lang, J., Marquis, P.: DA^2 merging operators. *Artificial Intelligence* **157**(1-2) (2004) 49–79
10. Baral, C., Kraus, S., Minker, J., Subrahmanian, V.S.: Combining knowledge bases consisting of first-order theories. *Computational Intelligence* **8** (1992) 45–71
11. Chacón, J.L., Pérez, R.P.: Merging operators: Beyond the finite case. *Information Fusion* **7**(1) (2006) 41–60
12. Booth, R.: Social contraction and belief negotiation. In: *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning, KR2002*, Morgan Kaufmann (2002) 375–384
13. Konieczny, S.: On the difference between merging knowledge bases and combining them. In Cohn, A.G., Giunchiglia, F., Selman, B., eds.: *Seventh International Conference on Principles of Knowledge Representation and Reasoning, KR2000*, San Francisco, Morgan Kaufmann (2000) 135–144
14. Bloch, I., Lang, J.: Towards Mathematical Morpho-Logics. In: *Technologies for Constructing Intelligent Systems. Volume 2*. Springer-Verlag (2002) 367–380
15. Lehmann, D., Magidor, M., Schlechta, K.: Distance semantics for belief revision. *Journal of Symbolic Logic* **66**(1) (2001) 295–317
16. Levesque, H.J., Lakemeyer, G.: *The Logic of Knowledge Bases*. MIT Press, Cambridge, MA, USA (2000)
17. Gorogiannis, N., Hunter, A.: Implementing semantic merging operators using binary decision diagrams. Technical report, University College London (2007)
18. Dalal, M.: Investigations into a theory of knowledge base revision. In: *Proceedings of the Seventh National Conference on Artificial Intelligence, AAAI’88. Volume 2*. (1988) 475–479
19. Grahne, G., Mendelzon, A.O., Revesz, P.Z.: Knowledgebase transformations. *Journal of Computer and System Sciences* **54**(1) (1997) 98–112
20. Hamming, R.W.: Error detecting and error correcting codes. *Bell System Technical Journal* **29** (1950) 147–160
21. Lafage, C., Lang, J.: Propositional distances and compact preference representation. *European Journal of Operational Research* **160**(3) (2005) 741–761

A Proof of Proposition 1

- If D satisfies $D2$ then Δ satisfies $B1$.
Let $E = \langle \phi_1, \dots, \phi_k \rangle$. From $D2$ and the fact that for all i , ϕ_i is consistent we can deduce that for all i, j , any member of $D_E^j(\phi_i)$ is consistent. This fact in conjunction with the restriction that $\mathcal{C}_D(E)$ consists of consistent tuples of dilations concludes the proof.
- If \sqsubseteq satisfies $C1$ then Δ satisfies $B2$.
If $\bigwedge E$ is consistent, then it will be the case that $E \in \mathcal{C}_D(E)$, and $dt(E) = \langle 0, \dots, 0 \rangle$. From the first half of $C1$ we conclude that for all tuples $E' \in \mathcal{C}_D(E)$, it will be the case that $dt(E) \sqsubseteq dt(E')$, and as such through Definition 11 we conclude that $\bigwedge E \in \Delta(E)$. From the second half of $C1$ we obtain that for any tuple $E' \in \mathcal{C}_D(E)$ such that $E \neq E'$ it will be the case that $dt(E) \sqsubset dt(E')$. This concludes the proof that $\Delta(E) \equiv \{\bigwedge E\}$.
- If D satisfies $D1$ then Δ satisfies $B3$.
Follows from $D1$ and Definition 11.

- If D satisfies $D3$ and \sqsubseteq satisfies $C4$ and $C2$, then Δ satisfies $B4$.
 We assume that $\chi \in \Delta(\phi \sqcup \psi)$. If $\chi \not\vdash \phi$ then we are done. So we assume that for all $\chi \in \Delta(\phi \sqcup \psi)$ it is the case that $\chi \vdash \phi$. By Definition 11, it follows that there exist $a_1, b_1 \in \mathbb{N}$ and $\phi_1 \in D_E^{a_1}(\phi)$, $\psi_1 \in D_E^{b_1}(\psi)$ such that $\langle \phi_1, \psi_1 \rangle$ is minimal in $\mathcal{C}_D(E)$. From $\chi \vdash \phi$ we get that $\phi_1 \wedge \psi_1 \vdash \phi$, therefore $\phi_1 \wedge \psi_1 \vdash \phi \wedge \psi_1$. By Definition 10, $\phi_1 \wedge \psi_1$ is consistent and, therefore, so is $\phi \wedge \psi_1$. By applying $D3$, we obtain $a_2 \in \mathbb{N}$ and $\phi_2 \in D_E^{a_2}(\phi)$ such that $\phi_2 \wedge \psi_1$ is consistent and $a_2 \leq b_1$. Using $C4$, we obtain $\langle a_2, 0 \rangle \sqsubseteq \langle 0, b_1 \rangle$. However, from $C1$ and $C2$ we know that $\langle 0, b_1 \rangle \sqsubseteq \langle a_1, b_1 \rangle$ and by transitivity we conclude that $\langle a_2, 0 \rangle \sqsubseteq \langle a_1, b_1 \rangle$ which means that, since $\phi_2 \wedge \psi_1$ is consistent, $\langle \phi_2, \psi_1 \rangle$ is minimal in $\mathcal{C}_D(\phi \sqcup \psi)$, or equivalently, that $\phi_2 \wedge \psi_1 \in \Delta(\phi \sqcup \psi)$. But then it must be the case that $\phi_2 \wedge \psi_1 \vdash \phi$, which gives $\phi_2 \wedge \psi_1 \vdash \phi \wedge \psi_1$ which is a contradiction, since by assumption $\phi \wedge \psi$ is inconsistent.
- If \sqsubseteq satisfies $C2$ then Δ satisfies $B5$.
 The assumptions mean that there exist $A_1 \in \mathcal{C}_D(E_1)$ and $A_2 \in \mathcal{C}_D(E_2)$ such that $\phi_1 = \bigwedge A_1$, $\phi_2 = \bigwedge A_2$, and that $\bigwedge A_1 \wedge \bigwedge A_2$ is consistent. Moreover, A_1 and A_2 are \sqsubseteq -minimal, i.e. for all $B_1 \in \mathcal{C}_D(E_1)$, $dt(A_1) \sqsubseteq dt(B_1)$ and for all $B_2 \in \mathcal{C}_D(E_2)$, $dt(A_2) \sqsubseteq dt(B_2)$. Since $\bigwedge A_1 \wedge \bigwedge A_2$ is consistent, it follows that $A_1 \sqcup A_2 \in \mathcal{C}_D(E_1 \sqcup E_2)$. In addition, from $C2$ and the minimalities of A_1 and A_2 as above, it follows that for all $C \in \mathcal{C}_D(E_1 \sqcup E_2)$, $dt(A_1 \sqcup A_2) \sqsubseteq dt(C)$ and thus, $\bigwedge A_1 \sqcup \bigwedge A_2 \in \Delta(E_1 \sqcup E_2)$.
- If \sqsubseteq satisfies $C3$ then Δ satisfies $B6$.
 Assume that $\chi \in \Delta(E_1 \sqcup E_2)$, which is to say that there is a \sqsubseteq -minimal tuple $A \in \mathcal{C}_D(E_1 \sqcup E_2)$ and $\chi = \bigwedge A$. Since χ is consistent, there exist tuples $B_1 \in \mathcal{C}_D(E_1)$, $B_2 \in \mathcal{C}_D(E_2)$ such that $A = B_1 \sqcup B_2$. Also, we know that there are \sqsubseteq -minimal tuples in $\mathcal{C}_D(E_1)$ and $\mathcal{C}_D(E_2)$ and that there are consistent pairs of those.
 If both B_1 and B_2 are \sqsubseteq -minimal in $\mathcal{C}_D(E_1)$ and $\mathcal{C}_D(E_2)$ respectively then we have nothing to prove.
 Let us assume that B_1 is *not* \sqsubseteq -minimal in $\mathcal{C}_D(E_1)$. This means that for any \sqsubseteq -minimal tuple $B'_1 \in \mathcal{C}_D(E_1)$ it will be the case that $dt(B'_1) \sqsubset dt(B_1)$ (on account of the totality of the preorders). Irrespective of the \sqsubseteq -minimality status of B_2 , it will also be the case that for any \sqsubseteq -minimal tuple B'_2 in $\mathcal{C}_D(E_2)$, $dt(B'_2) \sqsubseteq dt(B_2)$.
 Since the choice of B'_1 and B'_2 is arbitrary among the minimal tuples of $\mathcal{C}_D(E_1)$ and $\mathcal{C}_D(E_2)$ we will choose them so that $\bigwedge B'_1 \wedge \bigwedge B'_2$ is consistent, as per assumption. By using $C3$ we obtain $dt(B'_1 \sqcup B'_2) \sqsubset dt(B_1 \sqcup B_2)$. But then, $B_1 \sqcup B_2$ cannot be minimal in $\mathcal{C}_D(E_1 \sqcup E_2)$ anymore, a contradiction. Thus, there exist $\psi_1 = \bigwedge B_1$ and $\psi_2 = \bigwedge B_2$ such that $\psi_1 \wedge \psi_2 = \chi$ and $\psi_1 \in \Delta^{\text{Gmax}}(E_1)$ and $\psi_2 \in \Delta^{\text{Gmax}}(E_2)$.