

Improving Recovery for Belief Bases

Frances L. Johnson and Stuart C. Shapiro

Department of Computer Science and Engineering; SNePS Research Group
Center for Multisource Information Fusion; Center for Cognitive Science
University at Buffalo, The State University of New York
201 Bell Hall, Buffalo, NY 14260-2000, USA
{flj | shapiro}@cse.buffalo.edu

Abstract

The Recovery postulate for contraction says that any beliefs lost due to the contraction of some belief p should return if p is immediately re-asserted. Recovery holds for logically closed sets of beliefs, but it does not hold for belief bases (sets of beliefs that are not logically closed). This paper discusses the Recovery aspect of the belief base optimizing operation of reconsideration (which performs hindsight belief change) and compares it to the adherence to Recovery of traditional base and belief liberation contraction operations. We also discuss the similarities and differences between the belief base manipulations for belief liberation vs. those for reconsideration — because both approaches support the concept that removing a belief from a base might allow some previously removed beliefs to return.

1 Introduction

1.1 Motivation

Any agent reasoning from a set of beliefs must be able to perform basic belief change operations, including expansion, contraction and consolidation. Briefly, expansion is adding a belief to a set without concern for any inconsistencies it might raise; contraction of a set by a belief results in a set that does not entail (cannot derive) that belief — it is the removal or retraction¹ of that belief; and consolidation of a set of beliefs produces a consistent subset of the original set. See Section 1.3 for more details.

The Recovery postulate for contraction [Alchourrón *et al.*, 1985] says that a *logically closed* set K is contained in the set that results from contraction of K by a belief p followed by union with $\{p\}$ and deductive closure.

A belief *base*, according to the foundations approach (see discussion in [Gärdenfors, 1992] and [Hansson, 1999]), is a finite set of core or base beliefs (also called hypotheses in [Martins and Shapiro, 1988]) that have independent standing

¹The term *retraction* is also used in the literature to define a specific subclass of contraction. In this paper, we use the term retraction as a synonym for removal.

and are treated differently from derived beliefs. A base is, typically, *not* logically closed.

Contraction of a belief base B by some belief p is successful if p is absent from the resulting base *and* its logical closure. Although Recovery does not hold in general for belief base contraction (due to the lack of closure before contraction), it does hold *in some specific cases*.

The research defining *belief liberation* [Booth *et al.*, 2003] and *reconsideration* [Johnson and Shapiro, 2005] supports the concept that removing a belief from a base might allow some previously removed beliefs to return.² Because of this common view, we could not present the Recovery aspect of reconsideration without first examining the similarities and differences between these works (see Section 3.4). We then use belief liberation terminology and reconsideration in separate formulations of the Recovery postulate; and we discuss and compare the specific cases that do (or do not) adhere to these formulations as well as to the traditional Recovery postulate for belief bases.

1.2 Notation and Terminology

For this paper, we use a propositional language, \mathcal{L} , which is closed under the truth functional operators $\neg, \vee, \wedge, \rightarrow$, and \leftrightarrow . Formulas of the language \mathcal{L} are denoted by lowercase letters (p, q, r, \dots). Sets and sequences are denoted by uppercase letters (A, B, C, \dots). If set A derives p , it is denoted as $A \vdash p$. Cn , is defined by $Cn(A) = \{p \mid A \vdash p\}$, and $Cn(A)$ is called the *closure* of A . A belief base B is consistent iff $B \not\vdash \perp$, where \perp denotes logical contradiction. A *belief set* (a.k.a *theory*), K , is a *logically closed* set of beliefs (i.e. $K = Cn(K)$) [Alchourrón *et al.*, 1985]. We will use B for a belief base and K for a belief set.

Given a finite belief base, B , the set of *p-kernels* of B is the set $\{A \mid A \subseteq B, A \vdash p \text{ and } (\forall A' \subsetneq A) A' \not\vdash p\}$ [Hansson, 1994]. The *p-kernels* known to derive p are called *p's origin sets* in [Martins and Shapiro, 1988].

A *nogood* in the ATMS literature [de Kleer, 1986; Forbus and de Kleer, 1993] is a minimally inconsistent set S s.t. $S \vdash \perp$, but for all $S' \subsetneq S$, $S' \not\vdash \perp$.

²This is very different from the recovery of retracted beliefs during either saturated kernel contractions [Hansson, 1994] or the second part of Hybrid Adjustment [Williams and Sims, 2000]

1.3 Background

This section briefly reviews the traditional belief change operations of expansion and contraction of a logically closed belief set K [Alchourrón *et al.*, 1985] and expansion, kernel contraction and kernel consolidation of a finite belief base B . [Hansson, 1994; 1999].

Expansion

$K + p$ (the expansion of the belief set K by the belief p) is defined as $Cn(K \cup \{p\})$.

$B + p$ (the expansion of the belief base B by the belief p) is defined as $B \cup \{p\}$.

Kernel Contraction

The contraction of a base B [or set K] by a belief p is written as $B \sim p$ [$K \sim p$].

For this paper, $B \sim p$ is the kernel contraction [Hansson, 1994] of the belief base B by p (retraction of p from B) and, although constrained by several postulates, is basically the base resulting from the removal of at least one element from each p -kernel in B — unless $p \in Cn(\emptyset)$, in which case $B \sim p = B$. Given a belief base B , if K is the belief space for B ($K = Cn(B)$), then $K \sim p = Cn(B \sim p)$.

Kernel Consolidation

Consolidation (the removal of any inconsistency) is defined for belief *bases* only. Any inconsistent belief *set* is the set of *all* beliefs (due to closure in classical logic), making inconsistency removal a non-issue — set operations focus on preventing inconsistencies from occurring.

$B!$ (the kernel consolidation of B) is the removal of at least one element from each nogood in B s.t. $B! \subseteq B$ and $B! \not\vdash \perp$. This means that $B! =_{def} B \sim \perp$.

1.4 Recovery

Recovery does not hold for kernel contraction when elements of a p -kernel in B are retracted during the retraction of p , but are not returned as a result of the expansion by p followed by deductive closure. Not only do these base beliefs remain retracted, but derived beliefs that depend on them are also not recovered.

Example Given the base $B = \{s, d, s \rightarrow q\}$, $B \sim s \vee d = \{s \rightarrow q\}$, and $(B \sim s \vee d) + s \vee d = \{s \vee d, s \rightarrow q\}$. Not only do we not recover s or d as individual beliefs, but the derived belief q is also not recovered.

We feel the assertion of $s \vee d$ means that its earlier retraction was, in hindsight, not valid for *this current state*, so all effects of that retraction should be undone. There are various criticisms of Recovery in the literature (see [Hansson, 1999] and [Williams, 1994] for discussions and further references). We address these criticisms in [Johnson, 2005], but state our general argument below.

Our defense of Recovery is predicated on the fact that the recovered beliefs were at one time in the base as base beliefs. The recovery of those previously retracted base beliefs should occur whenever the reason that caused them to be removed is, itself, removed (or invalidated). In this case, the previously retracted beliefs should be returned to the base, specifically because they were base beliefs and the reason for disbelieving them no longer exists.

2 Reconsideration

2.1 Assuming a Linear Preference Ordering

In defining reconsideration, [Johnson and Shapiro, 2005] make the assumption that there is a recency-independent, linear preference ordering (\succeq) over all base beliefs. Thus, any base can be represented as a unique sequence of beliefs in order of descending preference: $B = p_1, p_2, \dots, p_n$, where $p_i \succ p_{i+1}$, $1 \leq i < n$. Note: $p_i \succ p_j$ means that p_i is strictly preferred over p_j (is stronger than p_j) and is true iff $p_i \succeq p_j$ and $p_j \not\prec p_i$.

2.2 The Knowledge State for Reconsideration

The knowledge state used to formalize reconsideration [Johnson and Shapiro, 2005] is a tuple with three elements. Starting with $B_0 = \emptyset$, B_n is the belief base that results from a series of expansion and consolidation operations on B_0 (and the subsequent resulting bases: B_1, B_2, B_3, \dots).³ and $B^\cup = \bigcup_{0 \leq i \leq n} B_i$. X_n is the set of base beliefs removed (and currently dis-believed: $B_n \cap X_n = \emptyset$) from these bases during the course of the series of operations: $X_n =_{def} B^\cup \setminus B_n$.

The knowledge state is a triple of the form $\langle B, B^\cup, \succeq \rangle$, where \succeq is the linear ordering of B^\cup , $X = B^\cup \setminus B$ and $Cn(\langle B, B^\cup, \succeq \rangle) = Cn(B)$. All triples are assumed to be in this form.

A numerical value for credibility of a base is calculated from the preference ordering of $B^\cup = p_1, \dots, p_n$: $\text{Cred}(B, B^\cup, \succeq) = \sum_{p_i \in B} 2^{n-i}$ (the bit vector indicating the elements in B) when $B \not\vdash \perp$. Otherwise, when $B \vdash \perp$, $\text{Cred}(B, B^\cup, \succeq) = -1$.

A linear ordering over bases (\succeq_{B^\cup}) is also defined: $B \succeq_{B^\cup} B'$ iff $\text{Cred}(B, B^\cup, \succeq) \geq \text{Cred}(B', B^\cup, \succeq)$.

2.3 Optimal Base

Given a possibly inconsistent set of base beliefs, $B^\cup = p_1, p_2, \dots, p_n$, ordered by \succeq , the base B is considered *optimal* w.r.t. B^\cup and \succeq if and only if $B \subseteq B^\cup$ and $(\forall B' \subseteq B^\cup) : B \succeq_{B^\cup} B'$. This favors retaining a single strong belief over multiple weaker beliefs.

As in [Johnson and Shapiro, 2005], an operation of contraction or consolidation produces the new base B' by using a global incision function⁴ that maximizes $\text{Cred}(B', B^\cup, \succeq)$ w.r.t. the operation being performed. Note: maximizing $\text{Cred}(B', B^\cup, \succeq)$ without concern for any specific operation would result in $B' = B^\cup!$.

Observation 2.1 *The consolidation of a base B is the optimal subset of that particular base (w.r.t. B^\cup and \succeq): $B! \subseteq B$ and $(\forall B' \subseteq B) : B! \succeq_{B^\cup} B'$.*

2.4 Operations on a Knowledge State

The following are operations on the knowledge state

$\mathbf{B} = \langle B, B^\cup, \succeq \rangle$.

³Adding beliefs to a finite base by way of expansion followed by consolidation is a form of non-prioritized belief change called *semi-revision* [Hansson, 1997].

⁴An incision function is the function that determines which beliefs should be removed during the operations of kernel contraction and kernel consolidation.

Expansion of B by p and its preference information, \succeq_p , is: $\mathbf{B} + \langle p, \succeq_p \rangle =_{def} \langle B + p, B^\cup + p, \succeq_1 \rangle$, where \succeq_1 is \succeq adjusted to incorporate the preference information \succeq_p — which positions p relative to other beliefs in B^\cup , while leaving the relative order of other beliefs in B^\cup unchanged. The resulting ordering is the transitive closure of these relative orderings.⁵

Contraction of B by p is:

$\mathbf{B} \sim p =_{def} \langle B \sim p, B^\cup, \succeq \rangle$.

Reconsideration of B [Johnson and Shapiro, 2005] is: $\mathbf{B} \downarrow =_{def} \langle B^\cup!, B^\cup, \succeq \rangle$.

Theorem 2.1 [Johnson and Shapiro, 2005] *The base resulting from reconsideration is optimal w.r.t. B^\cup and \succeq . Proved using Obs. 2.1.*

Observation 2.2 Reworded from [Johnson and Shapiro, 2005] *Given any knowledge state for B^\cup and \succeq , reconsideration on that state produces the optimal knowledge state: $(\forall B \subseteq B^\cup) : \langle B, B^\cup, \succeq \rangle \downarrow = \langle B_{opt}, B^\cup, \succeq \rangle$, where B_{opt} is the optimal base w.r.t. B^\cup and \succeq (because $B_{opt} = B^\cup!$).*

Optimized-addition to B (of the pair $\langle p, \succeq_p \rangle$) is:

$\mathbf{B} \downarrow \langle p, \succeq_p \rangle =_{def} (\langle B, B^\cup, \succeq \rangle + \langle p, \succeq_p \rangle) \downarrow$.

If B^\cup and \succeq are known, we adopt the shorthand writing of:

(1) $B \downarrow \langle p, \succeq_p \rangle$ to stand for $\langle B, B^\cup, \succeq \rangle \downarrow \langle p, \succeq_p \rangle$; and (2) $B^\cup \downarrow \langle p, \succeq_p \rangle$ to stand for $\langle B', B^\cup, \succeq \rangle \downarrow \langle p, \succeq_p \rangle$ for any $B' \subseteq B^\cup$. If the effect of adjusting the ordering by \succeq_p is also known, then $\langle p, \succeq_p \rangle$ can be reduced to p .

Observation 2.3 *Optimized-addition does not guarantee that the belief added will be in the optimized base — it might get removed during reconsideration.*

Example Let $B^\cup = p, p \rightarrow q, \neg q, p \rightarrow r, \neg r, m \rightarrow r, m$. And assume that $B = B^\cup! = p, p \rightarrow q, p \rightarrow r, m \rightarrow r, m$. $\langle B, B^\cup, \succeq \rangle \downarrow \langle \neg p, \succeq_{\neg p} \rangle = \langle \neg p, p \rightarrow q, \neg q, p \rightarrow r, \neg r, m \rightarrow r, m \rangle$, assuming that $\succeq_{\neg p}$ indicates $\neg p \succ p$. Notice the return of $\neg q$ and $\neg r$ to the base due to the removal of p , and the simultaneous removal of m to avoid a contradiction with $\neg r$ and $m \rightarrow r$. If, on the other hand, $\succeq_{\neg p}$ indicated $p \succ \neg p$, then the base would have remained unchanged.

Implementing Reconsideration

[Johnson and Shapiro, 2005] presents a TMS-friendly, efficient, anytime algorithm that implements reconsideration. Rather than examining *all* beliefs in B^\cup , or even X , the algorithm uses connections between beliefs and their shared no-goods, so that a retracted belief points to the removed beliefs which should be examined for possible return to the base. Therefore, the algorithm is “dependency-directed” (a phrase first coined in [Stallman and Sussman, 1977]); and it is called *dependency-directed reconsideration* (DDR).

3 Belief Liberation

3.1 Basic Notation

In this section, we summarize σ -liberation [Booth *et al.*, 2003] and compare it to reconsideration. Like reconsideration, liberation assumes a linear sequence of beliefs which is

⁵We assume that if $p \in B^\cup$, the location of p in the sequence might change — i.e. its old ordering information is removed before adding \succeq_p and performing closure — but all other beliefs remain in their same relative order.

called $\sigma = p_1, \dots, p_n$. The sequence is ordered by recency, where p_1 is the most recent information⁶ the agent has received (and has highest preference), and the set $[[\sigma]]$ is the set of all the sentences appearing in σ .

Since the ordering in this sequence is based on recency, for the remainder of this section, all comparisons between features of liberation and those of reconsideration are predicated on the assumption that *both* of their sequences are ordered by recency.⁷

3.2 A Belief Sequence Relative to K

In [Booth *et al.*, 2003] the ordering of σ is used to form the maximal consistent subset of $[[\sigma]]$ iteratively by defining the following: (1) $\mathcal{B}_0(\sigma) = \emptyset$. (2) for each $i = 0, 1, \dots, n - 1$: if $\mathcal{B}_i(\sigma) + p_{(i+1)} \not\vdash \perp$, then $\mathcal{B}_{(i+1)}(\sigma) = \mathcal{B}_i(\sigma) + p_{(i+1)}$, otherwise $\mathcal{B}_{(i+1)}(\sigma) = \mathcal{B}_i(\sigma)$. That is, each belief — from most recent to least — is added to the base only if it does *not* raise an inconsistency.

Definition 3.1 [Booth *et al.*, 2003] *Let K be a belief set and $\sigma = p_1, \dots, p_n$ a belief sequence. We say σ is a belief sequence relative to K iff $K = Cn(\mathcal{B}_n(\sigma))$.*

3.3 Removing a Belief q from K

In [Booth *et al.*, 2003] the operation of removing the belief q is defined using the following: (1) $\mathcal{B}_0(\sigma, q) = \emptyset$. (2) for each $i = 0, 1, \dots, n - 1$: if $\mathcal{B}_i(\sigma, q) + p_{(i+1)} \not\vdash q$, then $\mathcal{B}_{(i+1)}(\sigma, q) = \mathcal{B}_i(\sigma, q) + p_{(i+1)}$, otherwise $\mathcal{B}_{(i+1)}(\sigma, q) = \mathcal{B}_i(\sigma, q)$. Note that “ $\mathcal{B}_n(\sigma) = \mathcal{B}_n(\sigma, \perp)$ and $\mathcal{B}_n(\sigma, q)$ is the set-inclusion maximal amongst the subsets of $[[\sigma]]$ that do *not* imply q .” [Booth *et al.*, 2003]

Given a belief sequence σ relative to K , σ is used to define an operation \sim_σ for K such that $K \sim_\sigma q$ represents the result of removing q from K [Booth *et al.*, 2003]: $K \sim_\sigma q = Cn(\mathcal{B}_n(\sigma, q))$ if $q \notin Cn(\emptyset)$, otherwise $K \sim_\sigma q = K$.

Definition 3.2 [Booth *et al.*, 2003] *Let K be a belief set and \sim be an operator for K . Then \sim is a σ -liberation operator (for K) iff $\sim = \sim_\sigma$ for some belief sequence σ relative to K .*

Example [Booth *et al.*, 2003] Suppose $K = Cn(p \wedge q)$ and let $\sigma = p \rightarrow q, p, \neg p \wedge \neg q$ be the belief sequence relative to K — where $\neg p \wedge \neg q$ was originally blocked from inclusion in $\mathcal{B}_3(\sigma)$ by the inclusion of the more recent (and more preferred) belief p . Suppose we wish to remove p . We must first compute $\mathcal{B}_3(\sigma, p)$. We have $\mathcal{B}_0(\sigma, p) = \emptyset$, $\mathcal{B}_1(\sigma, p) = \{p \rightarrow q\} = \mathcal{B}_2(\sigma, p)$, and $\mathcal{B}_3(\sigma, p) = \{p \rightarrow q, \neg p \wedge \neg q\}$. Hence $K \sim_\sigma p = Cn(\mathcal{B}_3(\sigma, p)) = Cn(\neg p \wedge \neg q)$. Note how, when determining $\mathcal{B}_2(\sigma, p)$, p is nullified, which leads to the reinstatement, or liberation, of $\neg p \wedge \neg q$.

3.4 Comments On Liberation

Key Difference from Reconsideration

Reconsideration was intended specifically to improve adherence to Recovery for belief base contraction. The research in

⁶We have reversed the ordering from that presented in [Booth *et al.*, 2003] to avoid superficial differences when comparing their ordering with ours. We have adjusted the definitions accordingly.

⁷The differences between a recency-independent ordering and ordering by recency are discussed in Section 4.2.

belief liberation focuses on defining liberation operators for some belief set K relative to some *arbitrary* σ . The focus is on K and how it changes when a contraction is performed — whether there is *any* σ that shows that a given contraction operation is an operation of σ -liberation. The authors do not advocate maintaining any one, specific σ . Although it is clearly stated that σ -liberation does not adhere to Recovery, the similarity between σ -liberation and reconsideration prompted us to compare them in detail.

Similarities to Reconsideration

Assume $B^\cup = [[\sigma]]$ and is ordered by recency, and we refer to the belief set associated with σ as K_σ .

$\mathcal{B}_n(\sigma)$ is the maximal consistent subset of $[[\sigma]]$ — i.e. $\mathcal{B}_n(\sigma) = [[\sigma]]! = B^\cup!$. Similarly, $\mathcal{B}_n(\sigma, p)$ is the kernel contraction of $[[\sigma]]$ by p . In other words, $\mathcal{B}_n(\sigma, p) = B^\cup \sim p$.⁸ Thus, $K \sim_\sigma p = Cn(B^\cup \sim p)$.

If $B = B^\cup! = \mathcal{B}_n(\sigma)$, then we can define σ_B to be a recency ordering of *just* the beliefs in $\mathcal{B}_n(\sigma)$, and $K_\sigma = K_{\sigma_B}$. Now we can define contraction of an optimal knowledge state in terms of contraction for σ -liberation: $B \sim p = (K_\sigma \sim_{\sigma_B} p) \cap B$ and $Cn(\langle B, B^\cup, \succeq \rangle \sim p) = K_\sigma \sim_{\sigma_B} p$.

Let us define σ -addition (adding a belief to σ) as follows: $\sigma + p$ is adding the belief p to the sequence $\sigma = p_1, \dots, p_n$ to produce the new sequence $\sigma_1 = p, p_1, \dots, p_n$.⁹

If σ is the sequence for B^\cup , then the optimized addition of p to any knowledge state for B^\cup results in a base equivalent to the base for p added to σ : Given $B^\cup +_\downarrow p = \langle B', B^\cup + p, \succeq' \rangle$, then $B' = \mathcal{B}_{n+1}(\sigma + p)$.¹⁰

Likewise, σ -addition followed by recalculation of the belief set is equivalent to optimized-addition followed by closure: $K_{\sigma+p} = Cn(B^\cup +_\downarrow p)$.

Cascading Belief Status Effects of Liberation

It is important to realize that there is a potential cascade of belief status changes (both liberations *and* retractions) as the belief set resulting from a σ -liberation operation of retracting a belief p is determined; and these changes cannot be anticipated by looking at *only* the nogoods and kernels for p .

Example Let $\sigma = p \rightarrow q, p, \neg p \wedge \neg q, r \rightarrow p \vee q, r, \neg r$. Then, $\mathcal{B}_6(\sigma) = \{p \rightarrow q, p, r \rightarrow p \vee q, r\}$. Note that $r \in K_\sigma$ and $\neg r \notin K_\sigma$. $K \sim_\sigma p = Cn(\{p \rightarrow q, \neg p \wedge \neg q, r \rightarrow p \vee q, \neg r\})$. Even though r is not in a p -kernel in $[[\sigma]]$, $r \notin K \sim_\sigma p$. Likewise, $\neg r$ is liberated even though $\nexists N$ s.t. N is a nogood in $[[\sigma]]$ and $\{\neg r, p\} \subseteq N$.

Reconsideration has an identical effect. If $B^\cup = \sigma$, and $B = B^\cup! = \mathcal{B}_6(\sigma)$, then $\langle B, B^\cup, \succeq \rangle +_\downarrow \langle \neg p, \succeq_{\neg p} \rangle$, where $\succeq_{\neg p}$ indicates $\neg p \succ p$, would result in the base $B_1 = \{\neg p, p \rightarrow q, \neg p \wedge \neg q, r \rightarrow p \vee q, \neg r\}$.

⁸Note: specifically *not* $\mathcal{B}_n(\sigma, p) = B^\cup! \sim p$.

⁹This is also the technique described in [Chopra *et al.*, 2001] — though, again, we have reversed the order.

¹⁰Our notation for the base associated with a σ -addition is not inconsistent with the notation of [Booth *et al.*, 2003] for the base associated with a σ -liberation operation. Addition changes the sequence, so we are determining the base for the *new* sequence ($\sigma + p$): $\mathcal{B}(\sigma + p)$. The operation of σ -liberation changes the *base* used to determine the belief set (from $\mathcal{B}(\sigma)$ to $\mathcal{B}(\sigma, p)$), but the sequence σ remains unchanged.

4 Improving Recovery for Belief Bases

4.1 Comparing Recovery-like Formulations

Let $\mathbf{B} = \langle B, B^\cup, \succeq \rangle$, s.t. $B^\cup = [[\sigma]]$, $B = B^\cup! = \mathcal{B}_n(\sigma)$ and $K = K_\sigma = Cn(B)$, \mathbf{P} = the set of p -kernels in B , $\mathbf{p} = \langle p, \succeq_p \rangle$, $\mathbf{B}_1 = \langle B_1, B_1^\cup, \succeq_1 \rangle = (\mathbf{B} \sim p) +_\downarrow \mathbf{p}$, and $X_1 = B_1^\cup \setminus B_1$. The first element in any knowledge state triple is recognized as the currently believed base of that triple (e.g. B in \mathbf{B}), and is the default set for any shorthand set notation formula using that triple (e.g. $A \subseteq \mathbf{B}$ means $A \subseteq B$).

Table 1 shows the cases where different Recovery formulations hold — and where they do *not* hold. There is a column for each formulation and a row for each case. The traditional Recovery postulate for bases ($Cn(B) \subseteq Cn((B \sim p) + p)$) is shown in column (TR). In column (LR), the recovery postulate for σ -liberation retraction followed by expansion (Liberation-recovery, LR, our term) is: $K \subseteq ((K \sim_\sigma p) + p)$.

In column (OR), the recovery-like formulation for kernel contraction followed by optimized-addition is: $K \subseteq Cn((B \sim p) +_\downarrow \mathbf{p})$ (called Optimized-recovery, OR). We also claim $B \subseteq ((B \sim p) +_\downarrow \mathbf{p})$, which is more strict than Recovery: base beliefs are recovered in the base, itself, not just its closure. For column (OR-i), we assume that the ordering for B^\cup and B_1^\cup is recency. For column (OR-ii), we assume that the ordering is *not* recency-based, $p \in B^\cup$ (not applicable for Case 3), and optimized-addition returns p to the sequence in its *original* place (i.e. $\succeq = \succeq_1$). Note that (OR) is not a true Recovery axiom for some contraction operation; because it can be rewritten as $K \subseteq Cn(\langle (B \sim p) + p, \succeq \rangle)$, where reconsideration is performed *after* the expansion but *before* the closure to form the new belief space.

YES means the formulation *always* holds for that given case; NO means it does not *always* hold; NA means the given case is not possible for that column's conditions. The second entry indicates whether the base/set is optimal w.r.t. $B_1^\cup (= B^\cup + p = \sigma + p)$ and its linear order. If not optimal, then a designation for consistency is indicated. Recall that optimality requires consistency.

Theorem 4.1 *Expansion of an optimal knowledge state by a belief that is consistent with the base (and is not being relocated to a lower position in the ordering) results in a new and optimal knowledge state: Given $\mathbf{B} = \langle B, B^\cup, \succeq \rangle$, where $B = B^\cup!$ and $X = B^\cup \setminus B$, then ($\forall p$ s.t. $B + p \not\vdash \perp$): $\mathbf{B} + \langle p, \succeq_p \rangle = \langle B + p, B^\cup + p, \succeq' \rangle \downarrow = \langle B + p, B^\cup + p, \succeq' \rangle$. (Provided: if $p = p_i \in B^\cup = p_j \in (B^\cup + p)$, then $j \leq i$; otherwise \downarrow might remove p .)*

Proof: $B = B^\cup!$. ($\forall x \in X$) : $B + x \vdash \perp$ and ($\nexists B' \subseteq B$) s.t. both $(B \setminus B') + x \not\vdash \perp$ and ($\forall b \in B'$) $x \succeq b$. Therefore, since $B + p \not\vdash \perp$, then $\forall B'' \subseteq (B^\cup + p)$: $(B + p) \succeq_{B^\cup + p} B''$. \square

Case 1 In this simple case, $\{p\}$ is the sole p -kernel in B . For all formulations, p is removed then returned to the base, therefore all cases hold.

Case 2 Since there are p -kernels in B that consist of beliefs other than p , beliefs other than p must be retracted during contraction of p . For (TR), if $B = \{p \wedge q\}$, then $B \sim p = \emptyset$ and $(B \sim p) + p = \{p\}$. Therefore, $K \not\subseteq Cn((B \sim p) + p)$, and (TR) does not hold. For (LR), if $\sigma = p \wedge q$, then $K \sim_\sigma p = \emptyset$ and $(K \sim_\sigma p) + p = Cn(\{p\})$. So, (LR) also does

		(TR)	(LR)	(OR)	
Case		$K \subseteq Cn((B \sim p) + p)$	$K \subseteq Cn((K \sim_\sigma p) + p)$	$K \subseteq Cn((B \sim p) +_\downarrow \mathbf{p})$ but also $B \subseteq (B \sim p) +_\downarrow \mathbf{p}$	
		ordered by recency	ordered by recency	(i) ordered by recency	(ii) $p \in B^\cup$ and $\succeq_1 = \succeq$
1.	$p \in Cn(B);$ $\mathbf{P} = \{\{p\}\}$	YES optimal	YES possibly inconsistent	YES optimal	YES optimal
2.	$p \in Cn(B);$ $\mathbf{P} \setminus \{\{p\}\} \neq \emptyset$	NO consistent	NO possibly inconsistent	YES optimal	YES optimal
3.	$p \notin Cn(B);$ $B + p \not\vdash \perp$	YES optimal	YES optimal	YES optimal	NA
4.	$p \notin Cn(B);$ $B + p \vdash \perp$	YES inconsistent	YES inconsistent	NO optimal	YES optimal

Table 1: This table indicates whether each of three different Recovery formulations (TR, LR and OR) holds in each of four different cases (which comprise all possible states of belief). $K = Cn(B)$ and $\mathbf{p} = \langle p, \succeq_p \rangle$. YES means the formulation *always* holds for that given case; NO means it does not *always* hold; NA means the given case is not possible for that column’s conditions. See the text for a detailed description. Note: If requiring contraction for consistency maintenance only, a column for adherence to either $B \subseteq (B +_\downarrow \neg p) +_\downarrow p$ (ordered by recency) or $K_\sigma \subseteq K_{(\sigma + \neg p) + p}$ would match (OR-i).¹¹

not hold. For (OR), since $p \in Cn(B)$, then $B + p \not\vdash \perp$. Thus $B_1 = B + p$ (from Theorem 4.1), so $B \subseteq B_1$, and (OR) holds.

Case 3 Since $p \notin Cn(B)$ and $B + p \not\vdash \perp$, we know $p \notin B^\cup$ — otherwise, $(B+p) \succ_{B^\cup} B$ and $B \neq B^\cup$! as it was defined. Column (OR-ii) has NA (for “Not Applicable”) as its entry, because (OR-ii) assumes that $p \in B^\cup$. For the other columns, $B \sim p = B$, $K \sim_\sigma p = K = Cn(B)$, and $\mathbf{B} \sim p = \mathbf{B}$. Clearly, (TR) holds and (LR) holds. (OR-i) also holds (Theorem 4.1).

Case 4 Because $p \notin Cn(B)$, $B \sim p = B$ and $K \sim_\sigma p = K = Cn(B)$. Since $B + p \vdash \perp$ and both (TR) and (LR) produce inconsistent spaces, they both hold. For (OR), $\mathbf{B} \sim p = \mathbf{B}$. For (OR-i), the optimized-addition puts p at the most preferred end of the new sequence (most recent), so $p \in B_1$ forcing weaker elements of B to be retracted for consistency maintenance during reconsideration (recall $B + p \vdash \perp$). Therefore (OR-i) does not hold.¹¹ For (OR-ii), optimized-addition returns p to the same place in the sequence that it held in B^\cup (recall $B_1^\cup = B^\cup$ and $\succeq = \succeq_1$). Therefore, $\mathbf{B} = \mathbf{B}_1$ and (OR-ii) holds.

4.2 Discussion

When comparing the traditional base recovery adherence (in column TR) to optimized recovery adherence (shown in the OR columns), the latter produces improved adherence, because:

1. when the retraction of p is truly “undone” (column (OR-ii)), B is recovered in all applicable cases;
2. using a recency-based ordering (OR-i), B is recovered in all cases where $p \in Cn(B)$;
3. if expansion by p traditionally makes the final base inconsistent (TR,4), although B is not recovered, the final base *is* consistent *and* optimal (OR-i,4).

¹¹Producing an optimal base is preferred to adhering to a recovery-like formulation by having an inconsistent base.

Reconsideration eliminates the results of any preceding contraction, because B^\cup is unaffected by contraction: $(\mathbf{B} \sim p)_\downarrow = \mathbf{B}_\downarrow$. Likewise, optimized-addition also eliminates the results of any preceding contraction: $\forall q : (\mathbf{B} \sim q) +_\downarrow \mathbf{p} = \mathbf{B} +_\downarrow \mathbf{p}$.

If we consider contraction *for consistency-maintenance only* (assuming ordering by recency), the recovery-like formulation $\mathbf{B} \subseteq (\mathbf{B} +_\downarrow \neg p) +_\downarrow \mathbf{p}$ would have column entries identical to those in the column under (OR-i). Likewise, the entries in a column for $K_\sigma \subseteq K_{(\sigma + \neg p) + p}$ would also be identical to the entries for column (OR-i).¹²

We also note that the improved Recovery aspect that reconsideration provides does not involve the addition of *extra* beliefs to the belief base. A belief base can “adhere” to Recovery if the contraction operation to remove p also inserts $p \rightarrow q$ into the base, for every belief q that is removed during that retraction of p . However, this deviates from our assumption of a foundations approach, where the base beliefs represent the base input information from which the system or agent should reason. Not only would this technique insert unfounded *base beliefs*, but the recovery of previously removed beliefs would only show up in the belief *space*; whereas reconsideration actually returns the removed beliefs to the belief *base*.

If the linear ordering is *recency-independent* and $\succeq_1 \neq \succeq$, then there are cases where Optimized-recovery does *not* hold even though the resulting base will still be optimal. For Case 1, if p is re-inserted into the ordering at a weaker spot, it might be retracted during reconsideration if it is re-asserted in a position that is weaker than the conflicting elements of some pre-existing nogood *and* the incision function favors retracting p . This could also happen in case 2, unless the elements of some p -kernel are *all* high enough in the order to force the retraction of the beliefs conflicting with p . In Case 3 all Recovery formulations *always* hold. In Case 4, if p is inserted into the final ordering at a strong enough position, it could

¹²The (OR-i) results for $\mathbf{B} \subseteq (\mathbf{B} +_\downarrow \neg p) +_\downarrow \mathbf{p}$ and $K_\sigma \subseteq K_{(\sigma + \neg p) + p}$ show adherence to (R3) in [Chopra *et al.*, 2002].

survive the reconsideration step of optimized-addition — in which case, (OR) would not hold.

5 Conclusions and Future Work

From Section 3.4, we see that a system that can implement σ -liberation can also implement reconsideration and vice versa.

Kernel consolidation of a finite belief base has adherence results for Optimized-recovery (OR) that are preferred over the adherence results for the traditional Recovery postulate for base contraction. Thus, reconsideration imparts a Recovery aspect to belief bases.

Although σ -liberation retraction was never intended to adhere to Recovery, if we require that contractions are for consistency maintenance only, it adheres identically as well as kernel contraction adheres to OR.

Ongoing work involves formalizing reconsideration for non-linear belief orderings (those lacking comparability and/or anti-symmetry) and exploring adherence to the recovery-like postulates in [Chopra *et al.*, 2002] (altered using optimized-addition).

Acknowledgments

The authors are grateful for the support, insights and feedback of William J. Rapaport, Carl Alphonse, Ken Regan, David R. Pierce, Jan Chomicki, Samir Chopra, Thomas Meyer, and the SNePS Research Group.

References

- [Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, June 1985.
- [Booth *et al.*, 2003] Richard Booth, Samir Chopra, Aditya Ghose, and Thomas Meyer. Belief liberation (and retraction). In *TARK '03: Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 159–172. ACM Press, 2003. Full version in *Studia Logica* 79(1):47-72, 2005.
- [Chopra *et al.*, 2001] Samir Chopra, Konstantinos Georgatos, and Rohit Parikh. Relevance sensitive non-monotonic inference on belief sequences. *Journal of Applied Non-Classical Logics*, 11(1-2):131–150, 2001.
- [Chopra *et al.*, 2002] Samir Chopra, Aditya Ghose, and Thomas Meyer. Iterated revision and recovery: a unified treatment via epistemic states. In Frank van Harmelen, editor, *ECAI 2002: 15th European Conference on Artificial Intelligence*, number 77 in *Frontiers in Artificial Intelligence and Applications*, pages 541–545, Amsterdam, The Netherlands, 2002. IOS Press.
- [de Kleer, 1986] Johan de Kleer. An assumption-based truth maintenance system. *Artificial Intelligence*, 28(2):127–162, 1986.
- [Forbus and de Kleer, 1993] K. D. Forbus and J. de Kleer. *Building Problem Solvers*. MIT Press, Cambridge, MA, 1993.
- [Gärdenfors, 1992] Peter Gärdenfors. *Belief Revision*. Cambridge Computer Tracts. Cambridge University Press, Cambridge, 1992.
- [Hansson, 1994] Sven Ove Hansson. Kernel contraction. *J. Symb. Logic*, 59(3):845–859, 1994.
- [Hansson, 1997] Sven Ove Hansson. Semi-revision. *Journal of Applied Non-Classical Logic*, 7:151–175, 1997.
- [Hansson, 1999] Sven Ove Hansson. *A Textbook of Belief Dynamics*, volume 11 of *Applied Logic*. Kluwer, Dordrecht, The Netherlands, 1999.
- [Johnson and Shapiro, 2005] Frances L. Johnson and Stuart C. Shapiro. Dependency-directed reconsideration: Belief base optimization for truth maintenance systems. In *Proceedings of AAAI-05*, Menlo Park, CA, 2005. AAAI Press. In press.
- [Johnson, 2005] Frances L. Johnson. *Belief Change in a Deductively Open Belief Space*. PhD thesis, Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, NY, 2005. Forthcoming.
- [Martins and Shapiro, 1988] João P. Martins and Stuart C. Shapiro. A model for belief revision. *Artificial Intelligence*, 35:25–79, 1988.
- [Stallman and Sussman, 1977] Richard Stallman and Gerald Jay Sussman. Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis. *Artificial Intelligence*, 9(2):135–196, 1977.
- [Williams and Sims, 2000] Mary-Anne Williams and Aidan Sims. SATEN: An object-oriented web-based revision and extraction engine. In Chitta Baral and Miroslaw Truszyński, editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning NMR'2000*, 2000. CoRR article: cs.AI/0003059.
- [Williams, 1994] M.-A. Williams. On the logic of theory base change. In C. MacNish, D. Pearce, and L. Moniz Pereira, editors, *Logics in Artificial Intelligence (JELIA)*, pages 86–105. Springer, Berlin, Heidelberg, 1994.