Base Belief Change and Optimized Recovery[1]

Frances Johnson [2] and Stuart C. Shapiro
University at Buffalo, CSE Dept., Buffalo, NY, USA flj|shapiro@cse.buffalo.edu

Abstract.
Optimized Recovery (OR) adds belief base optimization to the traditional Recovery postulate—improving Recovery adherence without sacrificing adherence to the more accepted postulates or to the foundations approach. Reconsideration and belief liberation systems both optimize a knowledge base through consolidation of a chain of base beliefs; and recovered base beliefs are returned *to the base*. The effects match an iterated revision axiom and show benefits for pre-orders, as well. Any system that can resolve an inconsistent belief base can produce these results.

Keywords. Base belief revision, knowledge base optimization, reconsideration, recovery, truth maintenance system (TMS)

1. Introduction

1.1. Motivation

This paper shows how existing belief change operations can be used to optimize a belief base and improve the return of previously retracted base beliefs *to the base*—offering the best aspects of the Recovery postulate while retaining a true foundations approach as is required for implemented systems.

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 retraction[1] of that belief; and consolidation of a finite set of beliefs produces a consistent subset of the original set. These are discussed in more detail in Section 1.3.

The Recovery postulate for belief theory contraction [1] states that a *logically closed* belief theory K is contained in the belief theory that results from contraction of K by a belief p followed by union with $\{p\}$ and deductive closure. One feature of Recovery is

[1]This paper is a revised version of a 2005 IJCAI workshop paper [13].

[2]Correspondence to: Frances Johnson, University at Buffalo, Department of Computer Science and Engineering, 201 Bell Hall, Buffalo, NY 14260-2000, USA. Tel.: +1 716-998-8394; Fax: +1 716-645-3464; E-mail: flj@cse.buffalo.edu

[1]The 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.

that any beliefs lost, due to contraction by p, should become reasserted when p is returned and closure is performed; it is this feature of recovery that is a key element of our paper.

Although a belief base can be infinite and deductively closed (theoretically), we are focusing on the base of an *implemented system* which consists solely of input to the system (e.g., observations, sensor readings, rules). This follows the foundations approach (see discussion in [7] and [10]), where base beliefs have independent standing and are treated differently from inferred beliefs, and results in a finite base that is not deductively closed.[2] We refer to the closure of a belief base as its belief *space*.

The well accepted Success postulate [10]) requires that contraction of a belief base B by some belief p is successful when p is absent from the resulting base *and* its logical closure (unless $p \in Cn(\emptyset)$). Recovery does not hold in general for belief *base* contraction, because additional base beliefs removed during contraction by some belief p might not return (or be derivable) when p is returned to the base. This is because the base is *not* deductively closed prior to the contraction.

There are base contraction operations that *do* satisfy Recovery by inserting "Recovery-enhancing" beliefs into the base during contraction. These inserted beliefs, however, do not come from an input source, so this technique deviates from the foundations approach as we are applying it for an implemented system. Adding new base beliefs during contraction also violates the Inclusion postulate [1,10], which states that the result of contracting a belief theory/base should be a subset ($\subseteq$) of that theory/base. This is discussed further in Section 5.2.

The research defining *belief liberation* [2] and *reconsideration* [11,12] supports the concept that removing a belief from a base might allow some previously removed beliefs to return.[3] Both liberation and reconsideration research discuss a re-optimization of the current belief base when the sequence of all base beliefs has been altered. It is this re-optimization that is instrumental in defining Optimized Recovery and in providing the recovery feature that has been missing in base belief change. Our discussion of various recovery formulations is accompanied by a table that illustrates all cases where these formulations do (or do not) hold for base belief change.

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 $\leftrightarrow$. Atoms of the language $\mathcal{L}$ are denoted by lowercase letters $(p, q, r, \ldots)$. Sets and sequences are denoted by uppercase letters $(A, B, C, \ldots)$. 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 set of beliefs S is consistent iff $S \nvdash \perp$, where $\perp$ denotes logical contradiction. A *belief theory*, K, is a *logically closed* set of beliefs (i.e. $K = Cn(K)$) [1]. We will use B for a belief base and K for a belief theory.

Note that we use the term *set* to refer to *any* set of beliefs—whether finite or infinite, deductively closed or not.

[2]Note that we do *not* consider the rare case of a finite base that is considered "closed" if it contains at least one (but not all) of each logically equivalent belief in its deductive closure. This is rare and unlikely in a real-world implementation of any appreciable size.

[3]This is very different from the recovery of retracted beliefs during either saturated kernel contractions [8] or the second part of Hybrid Adjustment [16]

Given a belief base, B, the set of *p-kernels* of B is the set $\{A \mid A \subseteq B, A \vdash p$ and $(\forall A' \subsetneq A) A' \not\vdash p\}$ [8].

Truth maintenance systems (TMSs) [6] retain the information about how a belief is derived, distinguishing between base and derived beliefs. An assumption-based truth maintenance system (ATMS) [5] stores the minimal set of base beliefs underlying a derivation. A *nogood* in the ATMS literature is an inconsistent set of base beliefs. We will define a $\perp$-kernel (falsum-kernel) as a minimally inconsistent nogood: a set S s.t. $S \vdash \perp$, but for all $S' \subsetneq S$, $S' \not\vdash \perp$.

1.3. Background

This section briefly reviews the traditional belief change operations of expansion and contraction of a logically closed belief theory K [1] and expansion, kernel contraction and kernel consolidation of a finite belief base B. [8,10].

1.3.1. Expansion

$K + p$ (the expansion of the belief theory 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 $B \cup \{p\}$.

1.3.2. Kernel Contraction

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

For this paper, $B \sim p$ is the kernel contraction [8] 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 $B \sim p = B$.

A decision function determines which beliefs should be removed during kernel contraction.[4] Although minimal damage to a knowledge base is a desirable feature of a decision function, it often comes with increased computational cost; when choosing a decision function for an implemented system, the tradeoff between minimizing damage and minimizing complexity must be considered.

Given a belief base B, if the belief theory K is the belief space for B ($K = Cn(B)$), then the contraction of this belief space by p through the use of kernel contraction is defined as $K \sim p =_{def} Cn(B \sim p)$.

1.3.3. Kernel Consolidation

Consolidation (the removal of any inconsistency) is defined for belief *bases* only. Any inconsistent belief *theory* is the set of *all* beliefs (due to deductive closure), so operations on belief theories focus on *preventing* inconsistencies, as opposed to resolving them.

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

[4] An example of six different decision functions can be seen in the six different adjustment strategies used by SATEN [16].

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 1 Given the base $B = \{s, d, s \to q\}$, $B \sim s \vee d = \{s \to q\}$, and $(B \sim s \vee d) + s \vee d = \{s \vee d, s \to q\}$. 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 [10] and [15] for discussions and further references); their argument is that Recovery is not as essential an axiom for contraction as the other axioms, which we do not dispute.

We do, however, prefer to adhere to Recovery whenever possible, predicated on the fact that 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 such a case, the previously retracted beliefs should be returned *to the base*, 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, we make the assumption that there is a linear preference ordering ($\succeq$) over all base beliefs [11,12]. See [14] (also in these proceedings) for a discussion of reconsideration on non-linear pre-orders. Although the beliefs may be ordered by recency, we assume a different ordering may be used. Thus, any base can be represented as a unique sequence of beliefs in order of descending preference: $B = p_1, p_2, \ldots, p_n$, 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 $p_j \not\succeq p_i$.

2.2. The Knowledge State for Reconsideration

The knowledge state used to formalize reconsideration [11,12] 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: $B_1, B_2, B_3, \ldots$).[5], 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: $X_n =_{def} B^{\cup} \setminus B_n$.

[5] Adding beliefs to a finite base by way of expansion followed by consolidation is a form of non-prioritized belief change called *semi-revision* [9].

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 $B^\cup = p_1, \ldots, p_n$: Cred$(B, B^\cup, \succeq)$ = $\sum_{p_i \in B} 2^{n-i}$ (the bit vector indicating the elements in B) when $B \nvdash \perp$. Otherwise, when $B \vdash \perp$, Cred$(B, B^\cup, \succeq)$ = -1.

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

2.3. *Optimal Base*

Given a possibly inconsistent set of base beliefs, $B^\cup = p_1, p_2, ..., p_n$, 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 [11,12], an operation of contraction or consolidation produces the new base B' by using a global decision function that maximizes Cred$(B', B^\cup, \succeq)$ w.r.t. the operation being performed. Note: maximizing Cred$(B', B^\cup, \succeq)$ without concern for any specific operation would result in $B' = B^\cup!$.

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

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.[6]

Contraction of B by p is: $\mathbf{B} \sim p =_{def} \langle B \sim p, B^\cup, \succeq\rangle$.

Reconsideration of B [11,12] is: $\mathbf{B}!\!\!\cup =_{def} \langle B^\cup!, B^\cup, \succeq\rangle$.

Theorem 1 *[11,12] The base resulting from reconsideration is optimal w.r.t. $B^\cup$ and $\succeq$. Proved using Obs. 1.*

Observation 2 Reworded from [12] *Given* any *knowledge state for $B^\cup$ and $\succeq$, reconsideration on that state produces the* optimal knowledge state*: $(\forall B \subseteq B^\cup)$: $\langle B, B^\cup, \succeq\rangle!\!\!\cup = \langle B_{opt}, B^\cup, \succeq\rangle$, 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 $\langle p, \succeq_p\rangle$) [11] is:

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

---

[6]We 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.

If $B^{\cup}$ and $\succeq$ are known, a shorthand expression is used: (1) $B +_{\cup!} \langle p, \succeq_p \rangle$ to stand for $\langle B, B^{\cup}, \succeq \rangle +_{\cup!} \langle p, \succeq_p \rangle$; and (2) $B^{\cup} +_{\cup!} \langle p, \succeq_p \rangle$ to stand for $\langle B', B^{\cup}, \succeq \rangle +_{\cup!} \langle p, \succeq_p \rangle$ 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 3 *Optimized-addition does not guarantee that the belief added will be in the optimized base — it might get removed during reconsideration.*

Example 2 Let $B^{\cup} = p, p \rightarrow q, \neg q, p \rightarrow r, \neg r, m \rightarrow r, m$. And assume that $B = B^{\cup}! = p, p \rightarrow q, p \rightarrow r, m \rightarrow r, m$. $\langle B, B^{\cup}, \succeq \rangle +_{\cup!} \langle \neg p, \succeq_{\neg p} \rangle = \neg p, p \rightarrow q, \neg q, p \rightarrow r, \neg r, m \rightarrow r$, 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.

3. Belief Liberation

3.1. Basic Notation

In this section, we summarize σ-liberation [2] and compare it to reconsideration. Like reconsideration, liberation assumes a linear sequence of beliefs $\sigma = p_1, \ldots, p_n$. The sequence is ordered by recency, where p_1 is the most recent information[7] 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.[8]

3.2. A Belief Sequence Relative to K

In [2] 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 $i = 0, 1, \ldots, n-1$: if $\mathcal{B}_i(\sigma) + p_{(i+1)} \nvdash \perp$, 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 1 [2] Let K be a belief theory and $\sigma = p_1, \ldots, p_n$ 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 [2] the operation of removing the belief q is defined using the following: (1) $\mathcal{B}_0(\sigma, q) = \emptyset$. (2) for each $i = 0, 1, \ldots, n-1$: if $\mathcal{B}_i(\sigma, q) + p_{i+1} \nvdash q$, 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 "$\mathcal{B}_n(\sigma) = \mathcal{B}_n(\sigma, \perp)$

[7] We have reversed the ordering from that presented in [2] to avoid superficial differences with the ordering for reconsideration. We have adjusted the definitions accordingly.

[8] We discuss the effects of a recency-independent ordering in Section 5.2.

and $\mathcal{B}_n(\sigma, q)$ is the set-inclusion maximal amongst the subsets of $[[\sigma]]$ that do *not* imply q."[2]

Given a belief sequence σ relative to K, σ is used to define an operation $\sim_\sigma$ for K such that $K \sim_\sigma q$ represents the result of removing q from K [2]: $K \sim_\sigma q = Cn(\mathcal{B}_n(\sigma, q))$ if $q \notin Cn(\emptyset)$, otherwise $K \sim_\sigma q = K$.

Definition 2 [2] Let K be a belief theory 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.

4. Comparing Reconsideration and Liberation

4.1. The Sequence σ is Used for Defining Liberation

The research in belief liberation focuses on defining liberation operators for some belief theory 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 indicates that a given contraction operation is an operation of σ-liberation. Liberation research does not advocate maintaining any one, specific σ. It is clearly stated that σ-liberation does not adhere to Recovery, but if you maintain a belief base as a recency-ordered sequence, σ, then liberation terminology can be directly related to that of reconsideration and re-optimization.

4.2. Similarities

Assume $B^\cup = [[\sigma]]$ and is ordered by recency. We refer to the belief theory 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, $\mathcal{B}_n(\sigma, p) = B^\cup \sim p$.[9] 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 $\sigma = p_1, \ldots, p_n$ to produce the new sequence $\sigma_1 = p, p_1, \ldots, p_n$.[10]

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 $B^\cup +_{\circlearrowleft}^{!} p = \langle B', B^\cup + p, \succeq'\rangle$, then $B' = \mathcal{B}_{n+1}(\sigma + p)$.[11]

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

---

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

[10]This is also the technique described in [3].

[11]This notation for the base associated with a σ-addition is not inconsistent with the notation in [2] 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 theory (from $\mathcal{B}(\sigma)$ to $\mathcal{B}(\sigma, p)$), but the sequence σ remains unchanged.

		(TR)	(LR)	(OR)	
Case		$K \subseteq Cn((B \sim p) + p)$	$K \subseteq Cn((K \sim_\sigma p) + p)$	$K \subseteq Cn((B \sim p) +_{\cup!} \mathbf{p})$ but also $B \subseteq (B \sim p) +_{\cup!} \mathbf{p}$	
				(i)	(ii)
		ordered by recency	ordered by recency	ordered by recency	$p \in B^{\cup}$ 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.	$p \notin Cn(B)$; $B + p \nvdash \bot$	YES **optimal**	YES **optimal**	YES **optimal**	NA
4.	$p \notin Cn(B)$; $B + p \vdash \bot$	YES inconsistent	YES inconsistent	NO **optimal**	YES **optimal**

Table 1. This table indicates whether each of three recovery formulations (TR, LR and OR) always holds in each of four different cases (comprising all possible states of belief). $K = Cn(B)$ and $\mathbf{p} = \langle p, \succeq_p \rangle$. See the text for a detailed description. If contraction is used for consistency maintenance only, a column for adherence to either $B \subseteq (B +_{\cup!} \neg p) +_{\cup!} p$ (ordered by recency) or $K_\sigma \subseteq K_{(\sigma + \neg p) + p}$ would match (OR-i).[11]

4.3. *Cascading Belief Status Effects*

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

Example 3 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 $\bot$-kernel in $[[\sigma]]$ and $\{\neg r, p\} \subseteq N$.

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

5. Improving Recovery for Belief Bases

5.1. *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, $\mathbf{p} = \langle p, \succeq_p \rangle$, $\mathbf{B}_1 = \langle B_1, B_1^{\cup}, \succeq_1 \rangle = (\mathbf{B} \sim p) +_{\cup!} \mathbf{p}$, 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) is: $K \subseteq ((K \sim_\sigma p) + p)$.

In column (OR), the recovery-like formulation for kernel contraction followed by optimized-addition is: $K \subseteq Cn((B \sim p) +_{\cup!} \mathbf{p})$ (called Optimized-recovery, OR).

OR can also be written as $B \subseteq ((B \sim p) +_{\Downarrow} \mathbf{p})$, which is more strict than Recovery: base beliefs are actually recovered *in the base*, itself, not just its closure.

Essentially, OR an axiom about contraction followed by optimized-addition—as opposed to the regular Recovery axiom, which describes the results of contraction followed by expansion. In either case, recovering retracted beliefs is a desirable feature of contraction *followed by* either expansion or optimized-addition.

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 $K \subseteq Cn(\ ((B \sim p) + p)!_{\Downarrow}\ )$, where the re-optimizing operation of 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/theory 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 2 *Expansion of an optimal knowledge state by a belief that is consistent with the base (and is* not *being* re*located 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.* $B + p \nvdash \bot)$*:* $\mathbf{B} + \langle p, \succeq_p \rangle = \langle B + p, B^{\cup} + p, \succeq' \rangle !_{\Downarrow} = \langle B + p, B^{\cup} + p, \succeq' \rangle$*. (Provided: if* $p = p_i \in B^{\cup} = p_j \in (B^{\cup} + p)$*, then* $j \leq i$*; otherwise* $\Downarrow$ *might remove p.)*
Proof: $B = B^{\cup}!$ *.* $(\forall x \in X) : B + x \vdash \bot$ *and* $(\nexists B' \subseteq B)$ *s.t. both* $(B \setminus B') + x \nvdash \bot$ *and* $(\forall b \in B')x \succeq b$*. Since* $B + p \nvdash \bot$*, then* $\forall B'' \subseteq (B^{\cup} + p) : (B + p) \succeq_{B^{\cup}+p} B''$*.* ■

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 formulations hold.

Case 2 Since there are p-kernels in B that consist of beliefs other than p, some base beliefs other than p must be retracted during contraction by p. Returning these removed base beliefs is the recovery feature that is the central focus of this paper. 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 not hold. For (OR), since $p \in Cn(B)$, then $B + p \nvdash \bot$. Thus $B_1 = B + p$ (from Theorem 2), so $B \subseteq B_1$, and (OR) holds.

Case 3 Since $p \notin Cn(B)$ and $B + p \nvdash \bot$, 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 2).

Case 4 Because $p \notin Cn(B)$, $B \sim p = B$ and $K \sim_{\sigma} p = K = Cn(B)$. Since $B + p \vdash \bot$ 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 $B + p \vdash \bot$). Therefore (OR-

i) does not hold.[12] 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.

5.2. Discussion

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

1. if ordering by recency (OR-i), B is recovered in all cases where $p \in Cn(B)$;
2. any beliefs removed due to contraction by p are returned (OR,1;OR,2)
3. if expansion by p would make the final base inconsistent (TR,4), B is not recovered if recency ordered, but the final base *is* consistent *and* optimal (OR-i,4).
4. when the retraction of p is truly “undone" (column (OR-ii)), B is recovered in all applicable cases;

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

If we consider contraction *for consistency-maintenance only* (assuming ordering by recency), the recovery-like formulation $\mathbf{B} \subseteq (\mathbf{B} +_{!_{\cup}} \neg p) +_{!_{\cup}} \mathbf{p}$ 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). These results show adherence to (R3) in [4]: if $\neg p \notin K$, then $K \subseteq (K * \neg p) * p$, where $*$ is prioritized revision (consistent addition of a belief requiring the belief to be in the resulting belief theory) [1]. In the case where the ordering is not by recency, R3 still holds provided (1) $p \succ \neg p$ in the final ordering and (2) if p was in the original ordering, it is not weaker in the new ordering.

We also note that the improved recovery compliance that reconsideration provides does not involve the addition of *new* beliefs to the belief base during contraction. Belief base contraction 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*[13], 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*.

An additional benefit is that the belief removed (whether through contraction or revision by a contradicting belief) need not be reasserted in its original syntactic form. Any logically equivalent assertion will have the same effect: provided the newly asserted belief survives re-optimization, it and all those beliefs just retracted will be returned to the base. In fact, any belief that is inconsistent with the removed belief’s negation will have this same effect (assuming it is consistent with the starting base).

---

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

[13] The new beliefs are not from some input source, but derived from the contraction operation. This violates the foundations approach as well as the Inclusion postulate (as discussed in Section 1.1).

Example 4 Given a knowledge state triple with $B^{\cup} = B = \{s, d, s \rightarrow q\}$, $B \sim s \vee d = \{s \rightarrow q\}$, as described in Example 1, $(B \sim s \vee d) +_! (s \vee d) = (B +_! \neg(s \vee d)) +_! (s \vee d) = \{s \vee d, s, d, s \rightarrow q\}$ and q is derivable. Similarly, if the belief that is asserted last (and strongest) is merely inconsistent with $\neg(s \vee d)$, the recovery of retracted beliefs is performed just the same: $(B +_! \neg(s \vee d)) +_! (p \wedge (\neg s \rightarrow (d \wedge m)))$ results in a final base $B_2 = \{p \wedge (\neg s \rightarrow (d \wedge m))), s, d, s \rightarrow q\}$ and q is derivable (as is m).

If the linear ordering is *not* based on recency and $\succeq_1 \neq \succeq$, then there are cases where Optimized-recovery does *not* hold even though the resulting base will still be optimal—those cases where p does *not* survive the optimization process. 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 one of its pre-existing $\perp$-kernels *and* the decision 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 survive the reconsideration step of optimized-addition — in which case, (OR) would not hold. These exceptions are typical of any re-ordering of beliefs.

The benefits of reconsideration are not limited to linear orderings. A discussion of reconsideration on pre-orders is offered in [11] and [14] along with a table showing reconsideration using the six adjustment strategies implemented in SATEN [16], where five bases are improved—three to optimal, showing full recovery and adhering to (R3).[14]

Assuming an implemented TMS system retains its $\perp$-kernels, reconsideration (and its recovery-like benefits) can be implemented using an efficient, anytime algorithm called dependency-directed reconsideration (DDR) [11,12]. Examining a small subset of $B^{\cup}$ in a series of steps, the process can be suspended whenever reasoning, acting or belief change need to be performed. The system performs these operations on the most credible base it has *at that time*. DDR can be re-called later to continue its optimization, which will be adjusted to take the interleaved operations into account.

6. Conclusions and Future Work

Optimized Recovery (OR) adds belief base optimization to the traditional Recovery postulate allowing a system to experience the restoration of p-kernels *in the base* when re-asserting p without sacrificing adherence to the other more accepted postulates (such as Success and Inclusion) or to the foundations approach. Reconsideration (determining the base for a belief sequence) optimizes a base through consolidation of a chain of base beliefs. The effects match the iterated revision axiom (R3) and show benefits for total pre-orders, as well. Any system that implements consolidation can produce these results. The anytime algorithm for DDR can be implemented in a TMS.

Future work includes exploring how this research relates to other iterated belief change axioms and improving the current implementation of reconsideration in an existing ATMS so that it can handle non-linear orderings.

[14]Consolidation (!) is called *theory extraction* in [16]. We assume $a \succ \neg a$.
SATEN website: `http://magic.it.uts.edu.au/systems/saten.html`

Acknowledgments

The authors are grateful for the support, insights and feedback of William J. Rapaport, Carl Alphonce, Ken Regan, David R. Pierce, Jan Chomicki, Samir Chopra, Thomas Meyer, and the SNePS Research Group. Special thanks go out to Sven Ove Hansson, the outside reader for Fran's dissertation [11].

References

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

[2] R. Booth, S. Chopra, A. Ghose, and T. Meyer. Belief liberation (and retraction). *Studia Logica*, 79 (1):47–72, 2005.

[3] S. Chopra, K. Georgatos, and R. Parikh. Relevance sensitive non-monotonic inference on belief sequences. *Journal of Applied Non-Classical Logics*, 11(1-2):131–150, 2001.

[4] S. Chopra, A. Ghose, and T. Meyer. Iterated revision and recovery: a unified treatment via epistemic states. In F. 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.

[5] J. de Kleer. An assumption-based truth maintenance system. *Artificial Intelligence*, 28(2):127–162, 1986.

[6] K. D. Forbus and J. de Kleer. *Building Problem Solvers*. MIT Press, Cambridge, MA, 1993.

[7] P. Gärdenfors. *Belief Revision*. Cambridge Computer Tracts. Cambridge University Press, Cambridge, 1992.

[8] S. O. Hansson. Kernel contraction. *J. Symb. Logic*, 59(3):845–859, 1994.

[9] S. O. Hansson. Semi-revision. *Journal of Applied Non-Classical Logic*, 7:151–175, 1997.

[10] S. O. Hansson. *A Textbook of Belief Dynamics*, volume 11 of *Applied Logic*. Kluwer, Dordrecht, The Netherlands, 1999.

[11] F. L. Johnson. *Dependency-Directed Reconsideration: An Anytime Algorithm for Hindsight Knowledge-Base Optimization*. PhD thesis, Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, NY, January 2006.

[12] F. L. Johnson and S. C. Shapiro. Dependency-directed reconsideration: Belief base optimization for truth maintenance systems. In *Proceedings of the Twentieth National Conference on Artificial Intellicence (AAAI-05)*, pages 313–320, Menlo Park, CA, 2005. AAAI Press.

[13] F. L. Johnson and S. C. Shapiro. Improving recovery for belief bases. In L. Morgenstern and M. Pagnucco, editors, *IJCAI-05 Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC'05): Working Notes*, pages 65–70, Edinburgh, 2005. IJCAII.

[14] F. L. Johnson and S. C. Shapiro. Reconsideration on non-linear base orderings. In *Proceedings of STAIRS'06 at ECAI'06*, Amsterdam, 2006. IOS Press.

[15] M.-A. Williams. On the logic of theory base change. In C. MacNish, D. Pearce, and L. M. Pereira, editors, *Logics in Artificial Intelligence(JELIA)*, pages 86–105. Springer, Berlin, Heidelberg, 1994.

[16] M.-A. Williams and A. Sims. SATEN: An object-oriented web-based revision and extraction engine. In C. Baral and M. Truszyński, editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning NMR'2000*, 2000. CoRR article: cs.AI/0003059.