Knowledge State Reconsideration: Hindsight Belief Revision*

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

As a knowledge representation and reasoning (KRR) system gathers and reasons about information, it has to update its belief space to maintain consistency. Some belief change operations it can perform include expansion (addition with no consistency checking), contraction (aka removal or retraction), revision (consistent prioritized addition), and consolidation (elimination of *any and all* inconsistencies). Whether belief change operations are performed on theories (Alchourrón, Gärdenfors, & Makinson 1985) or bases (Nebel 1989; Hansson 1991; 1993b), with ideal agents or those that are resource-bounded (Wassermann 1999; Williams 1997), there is no doubt that the order of operations typically affects the makeup of the resulting belief base.

If a KRR system gains *new* information that, in hindsight, might have altered the outcome of an earlier belief change decision, the earlier decision should be re-examined. We call this operation *reconsideration*, and the result is an optimal belief base regardless of the order of previous belief change operations.

For this paper, we assume a global decision function is used in the belief change operations, and it will favor retaining the most preferred beliefs as determined by a preference ordering ($\succ$) that is irreflexive, anti-symmetric and transitive (referred to here as an IAT-preference ordering).[1] Any base can be represented as a sequence of beliefs in order of decending preference: $B = p_1, p_2, \ldots, p_n$, where p_i is preferred over p_{i+1} ($p_i \succ p_{i+1}$).

This work proposes maintaining a knowledge state (KS) which is defined as follows: $KS = \langle B, B^{\cup}, \succ\rangle$, where B is the current belief base[2], $B^{\cup}$ is the set of all beliefs that have ever been in the belief base at any time (effectively, the union of all past and current bases), and $\succ$ is the preference ordering relation for the beliefs in $B^{\cup}$.

*This work was supported in part by the US Army Communications and Electronics Command (CECOM), Ft. Monmouth, NJ through Contract #DAAB-07-01-D-G001 with Booze·Allen & Hamilton. The authors appreciate the insights and feedback of the SNePS Research Group.

[1]An ordering that is not anti-symmetric is more realistic, but requires more discussion than this format allows.

[2]Assuming a foundations approach where base beliefs have independent justification (Hansson 1993a).

We define reconsideration of a knowledge state KS (denoted $KS\circledast$) as: $KS\circledast = \langle B^{\cup}!, B^{\cup}, \succ\rangle$, where ! is a consolidation operation (Hansson 1991; 1997).

The goal of reconsideration is to obtain an optimal base. Given a $KS = \langle B, B^{\cup}, \succ\rangle$, $B = p_1, p_2, \ldots, p_n$ is optimal if it is consistent and there is no $B' = q_1, q_2, \ldots, q_m$ s.t. $B' \subseteq B^{\cup}$, B' is consistent, and either $B \subset B'$ or $\exists q_i$ s.t $q_i \succ p_i$ and $p_1, p_2, \ldots, p_{i-1} = q_1, q_2, \ldots, q_{i-1}$.

We assume a system using reconsideration would use the non-prioritized revision of semi-revision (Hansson 1997), where B revised by the belief p is defined as $(B \cup \{p\})!$. A belief being added must survive on its own merits.

As a result of reconsideration, beliefs retracted in an earlier belief change operation might be recovered, and some current beliefs might be retracted.

Example1: If we believe that Ty graduated Tufts Medical School (p) and infer that he is well-educated (q , using the base belief $p \rightarrow q$), we might disregard someone's comment that Ty does not seem well-educated ($\neg q$). Upon being told that p is a false statement, reconsideration would return $\neg q$ to our belief base, since there is no *current* reason to disbelieve it.

This is much like the recovery postulate for contraction ($\div$) followed by expansion ($+$) of a deductively closed belief set K (aka theory): $K \subseteq (K \div p) + p$ (Alchourrón, Gärdenfors, & Makinson 1985).

For those opposed to the recovery postulate: consider that any beliefs recovered through reconsideration were, originally, *base* beliefs with *independent justification*. They were removed solely because of some additional base belief that has since been removed. In hindsight, it is wise to return them to the base provided no other obstruction exists.

Believing that Ty is not well-educated ($\neg q$) just because he did not graduate Tufts Med School is reasonable, because we were *told* $\neg q$ and no longer have evidence to the contrary. If the order of operations that added then removed p had been reversed, $\neg q$ would never have been removed.

It should also be noted that reconsideration works equally well for relevance logics, where believing both p and q does not typically infer $p \rightarrow q$.

Example2: If we start with the base $B1$ as described in Figure 1, the addition of $\neg p$ (where $\neg p \succ p$) forces the retraction of p. NOTE: Most systems would stop here. Reconsideration produces the following changes: (1) $\neg q$ re-

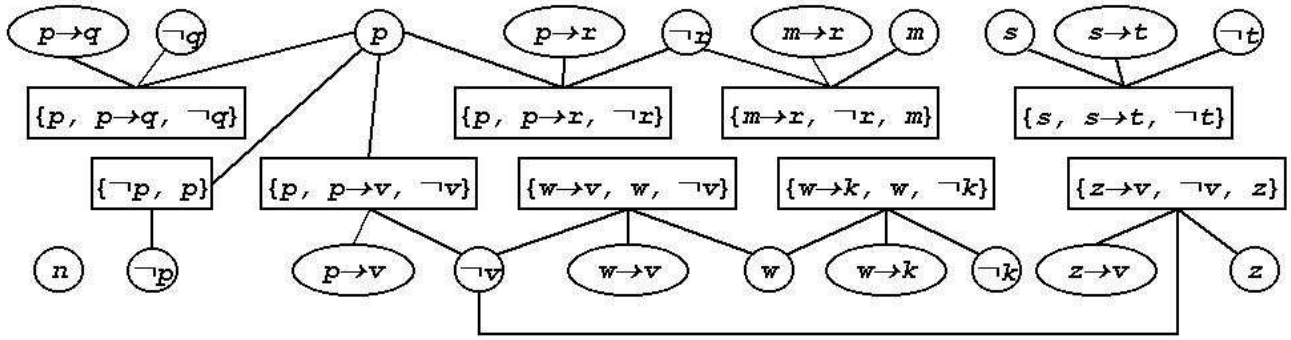

Figure 1: A graph showing the elements of $B^{\cup}$ (circles/ovals) of a KS connected to their inconsistent sets (rectangles). If $\neg p$ was not yet added to the base, and we had an IAT-preference ordering that ordered the beliefs in the following sequence from most to least preferred: $p, p \rightarrow q, p \rightarrow r, m \rightarrow r, s \rightarrow t, w \rightarrow v, w \rightarrow k, p \rightarrow v, z \rightarrow v, n, \neg r, w, s, \neg v, m, z, \neg q, \neg t, \neg k$, then the optimal base would be $B1 = \{p, p \rightarrow q, p \rightarrow r, m \rightarrow r, s \rightarrow t, w \rightarrow v, w \rightarrow k, p \rightarrow v, z \rightarrow v, n, w, s, m, z\}$. The semi-revision addition of $\neg p$ (preferred over p) followed by reconsideration is described in Example2.

turns to the base, and (2) $\neg r$ returns to the base with the simultaneous removal of m[3].

The "recovery for bases" flavor of reconsideration is seen if we discover that the ordering of p and $\neg p$ was incorrect: actually $p \succ \neg p$. Reconsideration would return p, remove $\neg p$, remove $\neg q$, and remove $\neg r$ while returning m to the base. The resulting base would be $B1$.

One algorithm for reconsideration in a TMS (Forbus & De Kleer 1993) might be to place all inconsistent sets on a priority queue in decreasing order of their culprits (least preferred beliefs). Processing each set, S, in turn (provided $S \subseteq$ the current base) by retracting its culprit would result in a literal reconsideration ($B^{\cup}!$). All beliefs not retracted during this process would be in the base of the resulting KS.

The implementation of reconsideration, however, does not require an actual consolidation over all $B^{\cup}$. Recalling Example2, a dependency-directed reconsideration (DDR) algorithm that determines that $\neg v$ cannot return to the base (due to its being the culprit for the inconsistent set $\{w \rightarrow v, w, \neg v\}$) would prune off the examination of the inconsistent sets containing $\neg k$ and z. The inconsistent set containing s could also be ignored, because it is not connected to p in any way. This latter case is representative of the possibly thousands of unrelated inconsistent sets for a typical belief base which *would* be checked during the literal $B^{\cup}!$ operation, but are ignored by DDR.

For some KRR systems, maintaining $B^{\cup}$ and its inconsistent sets would require additional memory usage, but TMS systems are already storing this information. Specifically, ATMS systems do this to enable reasoning in multiple contexts (Forbus & De Kleer 1993; Martins & Shapiro 1983). Some concession to limited resources may be needed (e.g. permanently remove oldest retracted beliefs).

Although incorporating reconsideration into a system's belief change operations will increase runtime, it is worth implementing, because (1) it allows sequential operations of belief change without concern for operation order and (2) it provides an optimal base from which to reason.

[3]Because $\neg r \succ m$.

References

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

Forbus, K. D., and De Kleer, J. 1993. *Building Problem Solvers*. Cambridge, MA: MIT Press.

Hansson, S. O. 1991. *Belief Base Dynamics*. Ph.D. Dissertation, Uppsala University.

Hansson, S. O. 1993a. Reversing the Levi identity. *Journal of Philosophical Logic* 22:637–669.

Hansson, S. O. 1993b. Theory contraction and base contraction unified. *Journal of Symbolic Logic* 58:602–625.

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

Martins, J. P., and Shapiro, S. C. 1983. Reasoning in multiple belief spaces. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann. 370–373.

Nebel, B. 1989. A knowledge level analysis of belief revision. In Brachman, R. J.; Levesque, H. J.; and Reiter, R., eds., *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, 301–311.

Wassermann, R. 1999. *Resource-bounded Belief Revision*. Ph.D. Dissertation, University of Amsterdam, Amsterdam, The Netherlands.

Williams, M.-A. 1997. Anytime belief revision. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann. 74–79.