traction of knowledge bases from the domain experts. The second phase, often called ***knowledge base refinement***, progressively refines the initial knowledge base into one that performs better. This phase usually adds, deletes, or alters some rule-components of rules in the knowledge base, so as to improve the system's *empirical adequacy*. That is, the refined knowledge base solves its intended problems satisfactorily. A knowledge base refinement problem is a special kind of optimization problem whose objective is to optimize the system performance over a given set of cases, subject to certain constraints, e.g., syntactic and seman-

from it, i.e., both A and ¬A (for some proposition A) can be derived. Early research [Nyugen et al., 1985, Suwa et al., 1982] focused on the static comparison of pairs of rules with conflicting or identical conclusions, obtaining only a partial checking because such pairwise analysis cannot detect all inconsistency in a knowledge base. KB-Reducer [Ginsberg, 1988] does a total checking by performing a complete prior analysis of the knowledge base. It computes all different sets of input supporting a given fact by employing the ATMS technique.

In moving from the propositional logic of boolean truth