8. Conclusions and Outlook

In this book, we have developed various approaches to what we have called *case-based inference*. The idea of CBI is to exploit experience in the form of a memory of observed cases (a case base consisting of input–output tuples) in order to predict a set of promising candidate outputs given a new query input. The corresponding inference schemes are based on suitable formalizations of the heuristic assumption that similar inputs yield similar outputs. Proceeding from a very simple, constraint-based model of this hypothesis, more sophisticated versions have been developed within different formal frameworks of approximate reasoning and reasoning under uncertainty. Let us again highlight the following properties of our approaches:

– For many of the CBI inference schemes, it was possible to derive interesting theoretical properties, for example the fact that a prediction covers the true outcome with high probability. From a case-based reasoning point of view, such CBI methods support a "reliable" retrieval of candidate solutions and, hence, contribute to the formal foundation of an important step within a CBR process.

– As our inference schemes hardly assume more than the specification of similarity measures for inputs and outputs, they are quite general and widely applicable. In particular, since no kind of transitivity is assumed for the similarity measures, the structure of the input and output space might be weaker than that of a metric space. This is a point of great practical relevance for CBR, where inputs and outputs can be complex objects. It also means that predictions can be derived in many situations where standard methods (e.g. from statistics) are not applicable.

– Our inference schemes are applicable for any pair of similarity measures, even if these measures are not defined in an optimal way. That is, the predictions remain correct, even though they might become rather imprecise. This, however, should not be seen as a disadvantage. On the contrary, these methods do not pretend a precision or credibility of case-based predictions which is actually not justified. Instead, imprecise predictions can be taken as an indication that either CBR is not appropriate for the application, or at least that the similarity measures are not well specified.

Most experiments conducted in this book have focused on prediction problems like classification and regression, for which benchmark data is available and predictive accuracy can easily be measured. Of course, from a (case-based) problem solving point of view, prediction appears to be the most simple problem class, mainly because there is no need for adapting the predicted solution. Still, an open question concerns the integration of our CBI methods into more complex CBR systems, that is, the use of these methods for more general types of problem solving.

A interesting idea in this regard is to apply CBI in the context of "search-oriented" CBR. In fact, according to the view of transformational adaptation taken in [30], case-based problem solving can be cast as a search process. Within the related model, (potential) cases correspond to search states and adaptation operators play the role of search operators. Now, the key idea is to use CBI in order to complement this model in a reasonable way. In fact, in [30] the authors note that, according to their approach, CBR could principally be realized by enumerating the search space completely. Understandably, they look at this idea with reservation, immediately pointing to the enormous complexity it brings about. Our approach applies exactly to this problem: CBI supports problem solving by predicting a promising subset of search states (outputs), thereby focusing search to promising regions of the search space and thus providing important information to a search method which is applied for actually finding a solution. From the perspective of CBR, this approach might not merely be seen as an application. In conjunction with the ideas presented in [30], it could contribute in a more general way to a formal framework of CBR in which (transformational) adaptation is realized as a search process and (case-based) experience is used in order to concentrate on promising regions of the related search space.

Indeed, in [30], the concept of similarity is integrated into problem solving by means of a, say, "ideal" similarity measure. By pointing to optimal initial search states, this measure somehow guarantees the retrieval of cases which can be adapted easily. Needless to say, finding such measures will be difficult in practice, if not impossible. As mentioned previously, our CBI methods take a different (more pragmatic) approach: They take any similarity measure as a *given* input, even if this measure is not "ideal", and then derive a *set* of *promising* search states rather than *the optimal* initial state.