Chapter 5

Summary and Conclusions

In this thesis we have explored the development of decision-theoretic methods in AI, focusing on their development within computer game playing (CGP). We saw that the general theoretical concepts of probability and expected utility were vital for moving CGP techniques beyond the earlier *ad hoc*, ill-justified techniques, and towards more well-founded methods, that explicitly tried to approximate the normative recommendations of decision theory.

However, just as in other fields, when CGP researchers began to apply decision theory, they quickly encountered what might be considered its primary flaw: its lack of a direct recommendation for what to do when there is not sufficient time to compute the decision-theoretic optimal action. Russell and Wefald were spurred by this problem to begin investigations into a theory of rational metareasoning, which would hopefully provide a principled way to design agents despite the limited rationality problem.

Their research, and the even more sophisticated later techniques by Baum and Smith, inspired us to explore whether the techniques being developed for game-playing had finally reached a point of generality and sophistication such that they could be fruitfully reapplied to a more practical problem. We chose decision analysis as the general problem framework to investigate, because decision trees share many features in common with game trees. We wished to take the Shannon paradigm of growing and evaluating a partial tree that has been so successful in game-playing and see if, augmented with the new decision-theoretic tree growth techniques, it could be equally useful when performing the sort of very large decision-analyses that are done in medical and other domains.

We decided to begin our investigations with a more accessible domain than that of medical decision making, namely that of Spider solitaire. In contrast to two-player games, this game is such that a classical decision analysis may be performed on its game tree to choose an optimal move. Unfortunately, after some time spent considering how best to

represent the moves and states in a Spider game to allow an efficient search, it was realized that most of the work in designing the program was going to be in finding the right way of working with the different levels of abstraction at which reasoning about Spider could be performed. The salient feature of decision-making in Spider was that it required exploring the right kinds of abstractions, not that it required delicate control of search within a fixed level of abstraction. The CGP search control techniques were essentially still inappropriate; Spider demanded a different kind of representation than the flat state-graph upon which all the CGP search control algorithms were based.

Although this was a disappointment with regard to the state of the art of limited rationality methods in decision-theoretic AI, it opened the door to a whole new area of investigation which is just now beginning to be explored. That area is the principled control of search through large, complex spaces of abstract states. The exciting possibility for future research is twofold: First, that the consideration of issues of abstraction, rather than an over-focus on concrete states, will allow the development of more general, robust, and powerful theoretical models for limited rationality (some investigations of which I have begun in chapter 3), and second, that in the meantime while such theories are being developed, the use of probabilistic and decision-theoretic techniques will aid greatly in the development of AI algorithms for reasoning with abstractions (which I have begun exploring in chapter 4). Other researchers, too, are already beginning to slowly explore limited rationality and abstraction techniques. Eventually there will be a symbiosis between LR ideas and abstraction ideas that will hasten the development of both into a common foundation for general reasoning.

As for computer game-playing, I believe that it has proven to be a fruitful breeding ground for many interesting ideas, but historically it has involved an overemphasis on particular games, and not enough of a focus on techniques designed to be readily generalizable to other domains. However, this problem can probably be fixed by exploring a greater variety of games beyond the classic chess-like games, such as games with much larger branching factors, such as Go ([Pell-90], [Brugmann-93]), and games with imperfect information, such as Solitaire and Bridge ([Gamback-et-al-91], [Smith-Nau-93]). Other ways we might try to improve CGP's relevance to AI are by concentrating on more general game-playing tasks such as Metagame [Pell-93], and by just resolving as individual

researchers to focus more effort on developing generalizable methods, as Russell and Wefald and Baum and Smith have done.

In one sense, recent CGP research has been largely disjoint from AI, in that many of the AI people that were working on games eventually either turned into brute-force chess hackers, or left games because their more general methods could not compete with the discouragingly successful brute-force ones. But now, some AI people are coming back to games, as evidenced by the [Games-93] proceedings. However, as this happens, we should always keep in mind that although games are interesting and fun, a game-playing technique is worthless if it is not developed with an eye towards eventually generalizing it to work on problems of more general utility.

Although the research reported in this thesis has not so far achieved as much in terms of conclusive results as was originally hoped, we believe that this thesis still serves as a valuable contribution, surveying and tying together several lines of research, synthesizing a broad overall perspective on the situation in computer game-playing and decision-theoretic AI, and pointing out some very interesting and promising directions, namely limited rationality and abstract search, for further advances in decision-theoretic artificial intelligence.