

PART IV

STATE UNCERTAINTY

Previous chapters have included uncertainty in the transition function, in terms of the uncertainty both in the resulting state and in the model. In this part, we extend uncertainty to include the state. Instead of observing the state exactly, we receive observations that have only a probabilistic relationship with the state. Such problems can be modeled as a *partially observable Markov decision process (POMDP)*. A common approach to solving POMDPs involves inferring a belief distribution over the underlying state at the current time step and then applying a policy that maps beliefs to actions. We will show how to update our belief distribution, given a past sequence of observations and actions. This enables us to devise exact solution methods for optimizing these belief-based policies. Unfortunately, POMDPs are intractable to optimally solve for all but the smallest of problems. We review a variety of offline approximation methods that tend to scale much better than exact methods to larger problems. We also show how to extend some of the online approximations discussed earlier in this book to accommodate partial observability. Finally, we introduce finite state controllers as an alternative policy representation and discuss methods that optimize them to solve POMDPs.