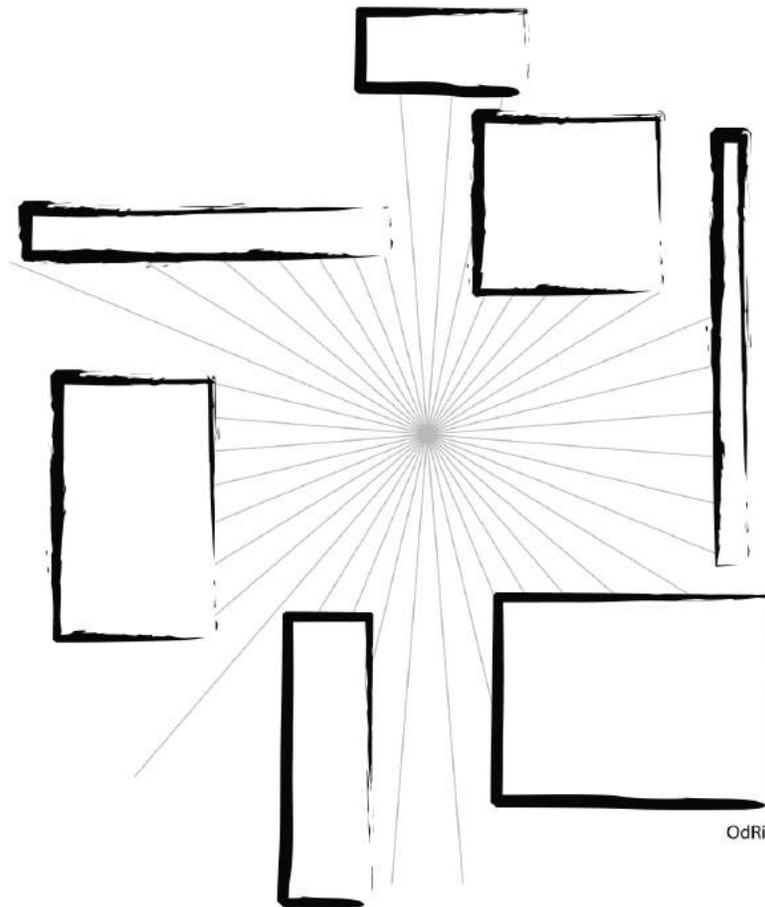


Part III

Decision-Theoretic Planning



Overview of Part III: Decision-Theoretic Planning

Planning Under Uncertainty

As in Part II, it also seems appropriate to give two names to Part III. It is officially called *decision-theoretic planning*, but it can also be considered as *planning under uncertainty*. All of the concepts in Parts I and II avoided models of uncertainties. Chapter 8 considered plans that can overcome some uncertainties, but there was no explicit modeling of uncertainty.

In this part, uncertainties generally interfere with two aspects of planning:

1. **Predictability:** Due to uncertainties, it is not known what will happen in the future when certain actions are applied. This means that future states are not necessarily predictable.
2. **Sensing:** Due to uncertainties, the current state is not necessarily known. Information regarding the state is obtained from initial conditions, sensors, and the memory of previously applied actions.

These two kinds of uncertainty are independent in many ways. Each has a different effect on the planning problem.

Making a single decision Chapter 9 provides an introduction to Part III by presenting ways to represent uncertainty in the process of making a single decision. The view taken in this chapter is that uncertainty can be modeled as interference from another decision maker. A special decision maker called *nature* will be introduced. The task is to make good decisions, in spite of actions applied by nature. Either *worst-case* or *probabilistic* models can be used to characterize nature's decision-making process. Some planning problems might involve multiple rational decision makers. This leads to game theory, which arises from the uncertainty about how other players will behave when they have conflicting goals. All of the concepts in Chapter 9 involve making a single decision; therefore, a state space is generally not necessary because there would only be one application of the state transition equation. One purpose of the chapter is to introduce and carefully evaluate the assumptions that are typically made in different forms of decision theory. This forms the basis of more complicated problems that follow, especially sequential decision making and control theory.

Uncertainty in predictability Chapter 10 takes the concepts from Chapter 9 and iterates them over multiple stages. This brings in the notions of states and state transitions, and can be considered as a blending of discrete planning concepts from Chapter 2 with the uncertainty concepts of Chapter 9. Some coverage of continuous state spaces and continuous time is also given, which extends ideas from Part II. The state transition equation is generally extended to allow future

states to depend on unknown actions taken by nature. In a game-theoretic setting, the state transitions may even depend on the actions of more than two decision makers.

For all of the models in Chapter 10, only uncertainty in predictability exists; the current state is always known. A plan is defined as a function that indicates the appropriate action to take from any current state. Plans are not formulated as a sequence of actions because future states are unpredictable, and responses to the future states may be required at the time they are achieved. Thus, for a fixed plan, the execution may be different each time: Different actions are applied and different states are reached. Plans are generally evaluated using worst-case, expected-case, or game-equilibrium analysis.

Uncertainty in sensing: The information space Chapter 11 introduces perhaps the most important concept of this book: the *information space*. If there is uncertainty in sensing the current state, then the planning problem naturally lives in an information space. An analogy can be made to the configuration space and motion planning. Before efforts to unify motion planning by using configuration space concepts [588, 657, 852], most algorithms were developed on a case-by-case basis. For example, robot manipulators and mobile robots have very different characteristics when defined in the world. However, once viewed in the configuration space, it is easier to consider general algorithms, such as those from Chapters 5 and 6.

A similar kind of unification should be possible for planning problems that involve sensing uncertainties (i.e., are unable to determine the current state). Presently, the methods in the literature are developed mainly around individual models and problems, as basic motion planning once was. Therefore, it is difficult to provide a perspective as unified as the techniques in Part II. Nevertheless, the concepts from Chapter 11 are used to provide a unified introduction to many planning problems that involve sensing uncertainties in Chapter 12. As in the case of the configuration space, some effort is required to learn the information space concepts; however, it will pay great dividends if the investment is made.

Chapter 12 presents several different problems and solutions for planning under sensing uncertainty. The problems include exploring new environments with robots, playing a pursuit-evasion game with cameras, and manipulating objects with little or no sensing. The chapter provides many interesting applications of information space concepts, but it should also leave you with the feeling that much more remains to be done. Planning in information spaces remains a challenging research problem throughout much of robotics, control theory, and artificial intelligence.