

2 *Recursive State Estimation*

2.1 Introduction

At the core of probabilistic robotics is the idea of estimating state from sensor data. State estimation addresses the problem of estimating quantities from sensor data that are not directly observable, but that can be inferred. In most robotic applications, determining what to do is relatively easy if one only knew *certain* quantities. For example, moving a mobile robot is relatively easy if the exact location of the robot and all nearby obstacles are known. Unfortunately, these variables are not directly measurable. Instead, a robot has to rely on its sensors to gather this information. Sensors carry only partial information about those quantities, and their measurements are corrupted by noise. State estimation seeks to recover state variables from the data. Probabilistic state estimation algorithms compute belief distributions over possible world states. An example of probabilistic state estimation was already encountered in the introduction to this book: mobile robot localization.

The goal of this chapter is to introduce the basic vocabulary and mathematical tools for estimating state from sensor data.

- Chapter 2.2 introduces basic probabilistic concepts used throughout the book.
- Chapter 2.3 describes our formal model of robot environment interaction, setting forth some of the key terminology used throughout the book.
- Chapter 2.4 introduces *Bayes filters*, the recursive algorithm for state estimation that forms the basis of virtually every technique presented in this book.

- Chapter 2.5 discusses representational and computational issues that arise when implementing Bayes filters.

2.2 Basic Concepts in Probability

RANDOM VARIABLE

This section familiarizes the reader with the basic notation and probabilistic facts used throughout the book. In probabilistic robotics, quantities such as sensor measurements, controls, and the states of a robot and its environment are all modeled as random variables. *Random variables* can take on multiple values, and they do so according to specific probabilistic laws. Probabilistic inference is the process of calculating these laws for random variables that are derived from other random variables and the observed data.

Let X denote a random variable and x denote a specific value that X might assume. A standard example of a random variable is a *coin flip*, where X can take on the values *heads* or *tails*. If the space of all values that X can take on is discrete, as is the case if X is the outcome of a coin flip, we write

$$(2.1) \quad p(X = x)$$

to denote the probability that the random variable X has value x . For example, a fair coin is characterized by $p(X = \text{head}) = p(X = \text{tail}) = \frac{1}{2}$. Discrete probabilities sum to one, that is,

$$(2.2) \quad \sum_x p(X = x) = 1$$

Probabilities are always non-negative, that is, $p(X = x) \geq 0$.

To simplify the notation, we will usually omit explicit mention of the random variable whenever possible, and instead use the common abbreviation $p(x)$ instead of writing $p(X = x)$.

PROBABILITY DENSITY
FUNCTION
NORMAL DISTRIBUTION
GAUSSIAN

Most techniques in this book address estimation and decision making in continuous spaces. Continuous spaces are characterized by random variables that can take on a continuum of values. Unless explicitly stated, we assume that all continuous random variables possess *probability density functions* (PDFs). A common density function is that of the one-dimensional *normal distribution* with mean μ and variance σ^2 . The PDF of a normal distribution is given by the following *Gaussian* function:

$$(2.3) \quad p(x) = (2\pi\sigma^2)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right\}$$

Normal distributions play a major role in this book. We will frequently abbreviate them as $\mathcal{N}(x; \mu, \sigma^2)$, which specifies the random variable, its mean, and its variance.

MULTIVARIATE
DISTRIBUTION

The Normal distribution (2.3) assumes that x is a scalar value. Often, x will be a multi-dimensional vector. Normal distributions over vectors are called *multivariate*. Multivariate normal distributions are characterized by density functions of the following form:

$$(2.4) \quad p(x) = \det(2\pi\Sigma)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right\}$$

COVARIANCE MATRIX

Here μ is the mean vector. Σ a *positive semidefinite* and *symmetric* matrix called the *covariance matrix*. The superscript T marks the transpose of a vector. The argument in the exponent in this PDF is quadratic in x , and the parameters of this quadratic function are μ and Σ .

The reader should take a moment to realize that Equation (2.4) is a strict generalization of Equation (2.3); both definitions are equivalent if x is a scalar value and $\Sigma = \sigma^2$.

Equations (2.3) and (2.4) are examples of PDFs. Just as discrete probability distribution always sums up to 1, a PDF always integrates to 1:

$$(2.5) \quad \int p(x) dx = 1$$

However, unlike a discrete probability, the value of a PDF is not upper-bounded by 1. Throughout this book, we will use the terms *probability*, *probability density*, and *probability density function* interchangeably. We will silently assume that all continuous random variables are measurable, and we also assume that all continuous distributions actually possess densities.

JOINT DISTRIBUTION

The *joint distribution* of two random variables X and Y is given by

$$(2.6) \quad p(x, y) = p(X = x \text{ and } Y = y)$$

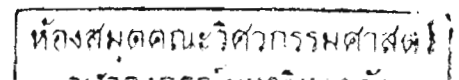
This expression describes the probability of the event that the random variable X takes on the value x and that Y takes on the value y . If X and Y are *independent*, we have

INDEPENDENCE

$$(2.7) \quad p(x, y) = p(x) p(y)$$

Often, random variables carry information about other random variables. Suppose we already know that Y 's value is y , and we would like to know the probability that X 's value is x conditioned on that fact. Such a probability will be denoted

$$(2.8) \quad p(x | y) = p(X = x | Y = y)$$



CONDITIONAL PROBABILITY and is called *conditional probability*. If $p(y) > 0$, then the conditional probability is defined as

$$(2.9) \quad p(x | y) = \frac{p(x, y)}{p(y)}$$

If X and Y are independent, we have

$$(2.10) \quad p(x | y) = \frac{p(x) p(y)}{p(y)} = p(x)$$

In other words, if X and Y are independent, Y tells us nothing about the value of X . There is no advantage of knowing the value of Y if we are interested in X . Independence, and its generalization known as conditional independence, plays a major role throughout this book.

THEOREM OF TOTAL PROBABILITY An interesting fact, which follows from the definition of conditional probability and the axioms of probability measures, is often referred to as *Theorem of total probability*:

$$(2.11) \quad p(x) = \sum_y p(x | y) p(y) \quad (\text{discrete case})$$

$$(2.12) \quad p(x) = \int p(x | y) p(y) dy \quad (\text{continuous case})$$

If $p(x | y)$ or $p(y)$ are zero, we define the product $p(x | y) p(y)$ to be zero, regardless of the value of the remaining factor.

BAYES RULE Equally important is *Bayes rule*, which relates a conditional of the type $p(x | y)$ to its “inverse,” $p(y | x)$. The rule, as stated here, requires $p(y) > 0$:

$$(2.13) \quad p(x | y) = \frac{p(y | x) p(x)}{p(y)} = \frac{p(y | x) p(x)}{\sum_{x'} p(y | x') p(x')} \quad (\text{discrete})$$

$$(2.14) \quad p(x | y) = \frac{p(y | x) p(x)}{p(y)} = \frac{p(y | x) p(x)}{\int p(y | x') p(x') dx'} \quad (\text{continuous})$$

PRIOR PROBABILITY Bayes rule plays a predominant role in probabilistic robotics (and probabilistic inference in general). If x is a quantity that we would like to infer from y , the probability $p(x)$ will be referred to as *prior probability distribution*, and y is called the *data* (e.g., a sensor measurement). The distribution $p(x)$ summarizes the knowledge we have regarding X prior to incorporating the data y .
 POSTERIOR PROBABILITY The probability $p(x | y)$ is called the *posterior probability distribution* over X . As (2.14) suggests, Bayes rule provides a convenient way to compute a posterior $p(x | y)$ using the “inverse” conditional probability $p(y | x)$ along with the prior probability $p(x)$. In other words, if we are interested in inferring

GENERATIVE MODEL

a quantity x from sensor data y , Bayes rule allows us to do so through the inverse probability, which specifies the probability of data y assuming that x was the case. In robotics, the probability $p(y | x)$ is often coined *generative model*, since it describes at some level of abstraction how state variables X cause sensor measurements Y .

An important observation is that the denominator of Bayes rule, $p(y)$, does not depend on x . Thus, the factor $p(y)^{-1}$ in Equations (2.13) and (2.14) will be the same for any value x in the posterior $p(x | y)$. For this reason, $p(y)^{-1}$ is often written as a *normalizer in Bayes rule* variable, and generically denoted η :

$$(2.15) \quad p(x | y) = \eta p(y | x) p(x)$$

The advantage of this notation lies in its brevity. Instead of explicitly providing the exact formula for a normalization constant—which can grow large very quickly in some of the mathematical derivations—we simply will use the normalization symbol η to indicate that the final result has to be normalized to 1. Throughout this book, normalizers of this type will be denoted η (or η', η'', \dots). **Important:** We will freely use the same η in different equations to denote normalizers, even if their actual values differ.

We notice that it is perfectly fine to condition any of the rules discussed so far on arbitrary random variables, such as the variable Z . For example, conditioning Bayes rule on $Z = z$ gives us:

$$(2.16) \quad p(x | y, z) = \frac{p(y | x, z) p(x | z)}{p(y | z)}$$

as long as $p(y | z) > 0$.

Similarly, we can condition the rule for combining probabilities of independent random variables (2.7) on other variables z :

$$(2.17) \quad p(x, y | z) = p(x | z) p(y | z)$$

CONDITIONAL INDEPENDENCE

Such a relation is known as *conditional independence*. As the reader easily verifies, (2.17) is equivalent to

$$(2.18) \quad p(x | z) = p(x | z, y)$$

$$(2.19) \quad p(y | z) = p(y | z, x)$$

Conditional independence plays an important role in probabilistic robotics. It applies whenever a variable y carries no information about a variable x if another variable's value z is known. Conditional independence does not

imply (absolute) independence, that is,

$$(2.20) \quad p(x, y | z) = p(x | z) p(y | z) \not\Rightarrow p(x, y) = p(x) p(y)$$

The converse is also in general untrue: absolute independence does not imply conditional independence:

$$(2.21) \quad p(x, y) = p(x) p(y) \not\Rightarrow p(x, y | z) = p(x | z) p(y | z)$$

In special cases, however, conditional and absolute independence may coincide.

EXPECTATION OF A RV

A number of probabilistic algorithms require us to compute features, or statistics, of probability distributions. The *expectation* of a random variable X is given by

$$(2.22) \quad E[X] = \sum_x x p(x) \quad (\text{discrete})$$

$$(2.23) \quad E[X] = \int x p(x) dx \quad (\text{continuous})$$

Not all random variables possess finite expectations; however, those that do not are of no relevance to the material presented in this book.

The expectation is a linear function of a random variable. In particular, we have

$$(2.24) \quad E[aX + b] = aE[X] + b$$

for arbitrary numerical values a and b . The covariance of X is obtained as follows

$$(2.25) \quad \text{Cov}[X] = E[X - E[X]]^2 = E[X^2] - E[X]^2$$

The covariance measures the squared expected deviation from the mean. As stated above, the mean of a multivariate normal distribution $\mathcal{N}(x; \mu, \Sigma)$ is μ , and its covariance is Σ .

ENTROPY

A final concept of importance in this book is *entropy*. The entropy of a probability distribution is given by the following expression:

$$(2.26) \quad H_p(x) = E[-\log_2 p(x)]$$

which resolves to

$$(2.27) \quad H_p(x) = -\sum_x p(x) \log_2 p(x) \quad (\text{discrete})$$

$$(2.28) \quad H_p(x) = -\int p(x) \log_2 p(x) dx \quad (\text{continuous})$$

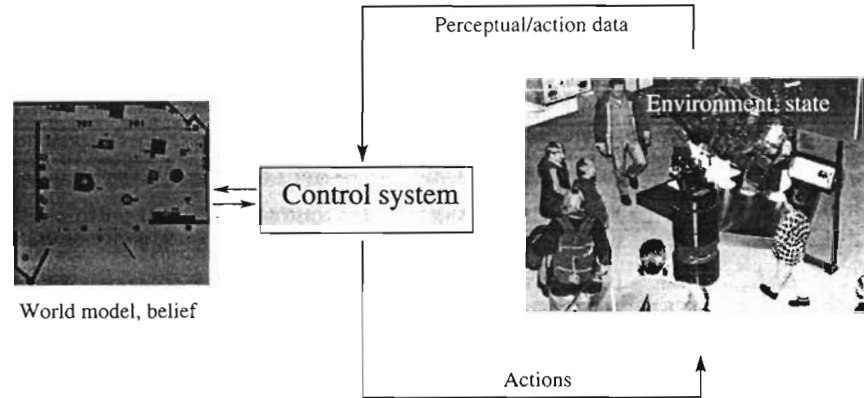


Figure 2.1 Robot environment interaction.

The concept of entropy originates in information theory. The entropy is the expected information that the value of x carries. In the discrete case, $-\log_2 p(x)$ is the number of bits required to encode x using an optimal encoding, assuming that $p(x)$ is the probability of observing x . In this book, entropy will be used in robotic information gathering, so as to express the information a robot may receive upon executing specific actions.

2.3 Robot Environment Interaction

ROBOT ENVIRONMENT

Figure 2.1 illustrates the interaction of a robot with its environment. The *environment*, or *world*, is a dynamical system that possesses internal state. The robot can acquire information about its environment using its sensors. However, sensors are noisy, and there are usually many things that cannot be sensed directly. As a consequence, the robot maintains an internal belief with regards to the state of its environment, depicted on the left in this figure.

The robot can also influence its environment through its actuators. The effect of doing so is often somewhat unpredictable. Hence, each control action affects both the environment state, and the robot's internal belief with regards to this state.

This interaction will now be described more formally.

2.3.1 State

STATE Environments are characterized by *state*. For the material presented in this book, it will be convenient to think of the state as the collection of all aspects of the robot and its environment that can impact the future. Certain state variables tend to change over time, such as the whereabouts of people in the vicinity of a robot. Others tend to remain static, such as the location of walls in (most) buildings. State that changes will be called *dynamic state*, which distinguishes it from *static state*, or non-changing state. The state also includes variables regarding the robot itself, such as its pose, velocity, whether or not its sensors are functioning correctly, and so on.

Throughout this book, state will be denoted x ; although the specific variables included in x will depend on the context. The state at time t will be denoted x_t . Typical state variables used throughout this book are:

- POSE
- The robot *pose*, which comprises its location and orientation relative to a global coordinate frame. Rigid mobile robots possess six such state variables, three for their Cartesian coordinates, and three for their angular orientation (pitch, roll, and yaw). For rigid mobile robots confined to planar environments, the pose is usually given by three variables, its two location coordinates in the plane and its heading direction (yaw).
 - In robot manipulation, the pose includes variables for the *configuration of the robot's actuators*. For example, they might include the joint angles of revolute joints. Each degree of freedom in a robot arm is characterized by a one-dimensional configuration at any point in time, which is part of the kinematic state of the robot. The robot configuration is often referred to as *kinematic state*.
 - The *robot velocity and the velocities of its joints* are commonly referred to as *dynamic state*. A rigid robot moving through space is characterized by up to six velocity variables, one for each pose variables. Dynamic state will play only a minor role in this book.
 - The *location and features of surrounding objects in the environment* are also state variables. An object may be a tree, a wall, or a pixel within a larger surface. Features of such objects may be their visual appearance (color, texture). Depending on the granularity of the state that is being modeled, robot environments possess between a few dozen and up to hundreds of billions of state variables (and more). Just imagine how many bits it will take to accurately describe your physical environment! For many of the

LANDMARK problems studied in this book, the location of objects in the environment will be static. In some problems, objects will assume the form of *landmarks*, which are distinct, stationary features of the environment that can be recognized reliably.

- The *location and velocities of moving objects and people* are also potential state variables. Often, the robot is not the only moving actor in its environment. Other moving entities possess their own kinematic and dynamic state.
- There are many other state variables that may impact a robot's operation. For example, whether or not a sensor is broken can be a state variable, as can be the level of battery charge for a battery-powered robot. The list of potential state variables is endless!

COMPLETE STATE A state x_t will be called *complete* if it is the best predictor of the future. Put differently, completeness entails that knowledge of past states, measurements, or controls carry no additional information that would help us predict the future more accurately. It is important to notice that our definition of completeness does not require the future to be a *deterministic* function of the state. The future may be stochastic, but no variables prior to x_t may influence the stochastic evolution of future states, unless this dependence is mediated through the state x_t . Temporal processes that meet these conditions are commonly known as *Markov chains*.

MARKOV CHAIN

The notion of state completeness is mostly of theoretical importance. In practice, it is impossible to specify a complete state for any realistic robot system. A complete state includes not just all aspects of the environment that may have an impact on the future, but also the robot itself, the content of its computer memory, the brain dumps of surrounding people, etc. Some of those are hard to obtain. Practical implementations therefore single out a small subset of all state variables, such as the ones listed above. Such a state is called *incomplete state*.

INCOMPLETE STATE

In most robotics applications, the state is continuous, meaning that x_t is defined over a continuum. A good example of a continuous state space is that of a robot pose, that is, its location and orientation relative to an external coordinate system. Sometimes, the state is discrete. An example of a discrete state space is the (binary) state variable that models whether or not a sensor is broken. State spaces that contain both continuous and discrete variables are called *hybrid* state spaces.

In most cases of interesting robotics problems, state changes over time. Time, throughout this book, will be discrete, that is, all interesting events will

take place at discrete time steps $t = 0, 1, 2, \dots$. If the robot starts its operation at a distinct point in time, we will denote this time as $t = 0$.

2.3.2 Environment Interaction

There are two fundamental types of interactions between a robot and its environment: The robot can influence the state of its environment through its actuators, and it can gather information about the state through its sensors. Both types of interactions may co-occur, but for didactic reasons we will separate them throughout this book. The interaction is illustrated in Figure 2.1.

MEASUREMENT

- **Environment sensor measurements.** Perception is the process by which the robot uses its sensors to obtain information about the state of its environment. For example, a robot might take a camera image, a range scan, or query its tactile sensors to receive information about the state of the environment. The result of such a perceptual interaction will be called a *measurement*, although we will sometimes also call it *observation* or *percept*. Typically, sensor measurements arrive with some delay. Hence they provide information about the state a few moments ago.

CONTROL ACTION

- **Control actions** change the state of the world. They do so by actively asserting forces on the robot's environment. Examples of *control actions* include robot motion and the manipulation of objects. Even if the robot does not perform any action itself, state usually changes. Thus, for consistency, we will assume that the robot *always* executes a control action, even if it chooses not to move any of its motors. In practice, the robot continuously executes controls and measurements are made concurrently.

Hypothetically, a robot may keep a record of all past sensor measurements and control actions. We will refer to such a collection as the *data* (regardless of whether they are being memorized or not). In accordance with the two types of environment interactions, the robot has access to two different data streams.

- **Environment measurement data** provides information about a momentary state of the environment. Examples of measurement data include camera images, range scans, and so on. For most parts, we will simply ignore small timing effects (e.g., most laser sensors scan environments sequentially at very high speeds, but we will simply assume the measurement corresponds to a specific point in time). The measurement data at time t will be denoted z_t .

Throughout most of this book, we simply assume that the robot takes exactly one measurement at a time. This assumption is mostly for notational convenience, as nearly all algorithms in this book can easily be extended to robots that can acquire variable numbers of measurements within a single time step. The notation

$$(2.29) \quad z_{t_1:t_2} = z_{t_1}, z_{t_1+1}, z_{t_1+2}, \dots, z_{t_2}$$

denotes the set of all measurements acquired from time t_1 to time t_2 , for $t_1 \leq t_2$.

- **Control data** carry information about the *change of state* in the environment. In mobile robotics, a typical example of control data is the velocity of a robot. Setting the velocity to 10 cm per second for the duration of five seconds suggests that the robot's pose, after executing this motion command, is approximately 50 cm ahead of its pose before command execution. Thus, control conveys information regarding the change of state.

ODOMETER

An alternative source of control data are *odometers*. Odometers are sensors that measure the revolution of a robot's wheels. As such they convey information about the change of state. Even though odometers are sensors, we will treat odometry as control data, since they measure the effect of a control action.

Control data will be denoted u_t . The variable u_t will always correspond to the change of state in the time interval $(t-1; t]$. As before, we will denote sequences of control data by $u_{t_1:t_2}$, for $t_1 \leq t_2$:

$$(2.30) \quad u_{t_1:t_2} = u_{t_1}, u_{t_1+1}, u_{t_1+2}, \dots, u_{t_2}$$

Since the environment may change even if a robot does not execute a specific control action, the fact that time passed by constitutes, technically speaking, control information. We therefore assume that there is exactly one control data item per time step t , and include as legal action "*do-nothing*".

The distinction between measurement and control is a crucial one, as both types of data play fundamentally different roles in the material yet to come. Environment perception provides information about the environment's state, hence it tends to increase the robot's knowledge. Motion, on the other hand, tends to induce a loss of knowledge due to the inherent noise

in robot actuation and the stochasticity of robot environments. By no means is our distinction intended to suggest that actions and perceptions are separated in time. Rather, perception and control takes place concurrently. Our separation is strictly for convenience.

2.3.3 Probabilistic Generative Laws

The evolution of state and measurements is governed by probabilistic laws. In general, the state x_t is generated stochastically from the state x_{t-1} . Thus, it makes sense to specify the probability distribution from which x_t is generated. At first glance, the emergence of state x_t might be conditioned on all past states, measurements, and controls. Hence, the probabilistic law characterizing the evolution of state might be given by a probability distribution of the following form: $p(x_t | x_{0:t-1}, z_{1:t-1}, u_{1:t})$. Notice that through no particular motivation we assume here that the robot executes a control action u_1 first, and then takes a measurement z_1 .

An important insight is the following: If the state x is complete then it is a sufficient summary of all that happened in previous time steps. In particular, x_{t-1} is a sufficient statistic of all previous controls and measurements up to this point in time, that is, $u_{1:t-1}$ and $z_{1:t-1}$. From all the variables in the expression above, only the control u_t matters if we know the state x_{t-1} .

In probabilistic terms, this insight is expressed by the following equality:

$$(2.31) \quad p(x_t | x_{0:t-1}, z_{1:t-1}, u_{1:t}) = p(x_t | x_{t-1}, u_t)$$

The property expressed by this equality is an example of *conditional independence*. It states that certain variables are independent of others if one knows the values of a third group of variables, the conditioning variables. Conditional independence will be exploited pervasively in this book. It is the primary reason why many of the algorithms presented in the book are computationally tractable.

One might also want to model the process by which measurements are being generated. Again, if x_t is complete, we have an important conditional independence:

$$(2.32) \quad p(z_t | x_{0:t}, z_{1:t-1}, u_{1:t}) = p(z_t | x_t)$$

In other words, the state x_t is sufficient to predict the (potentially noisy) measurement z_t . Knowledge of any other variable, such as past measurements, controls, or even past states, is irrelevant if x_t is complete.

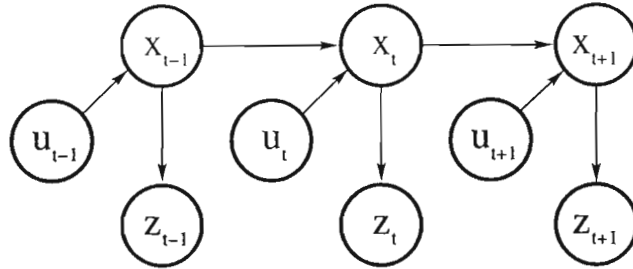


Figure 2.2 The dynamic Bayes network that characterizes the evolution of controls, states, and measurements.

STATE TRANSITION
PROBABILITY

This discussion leaves open as to what the two resulting conditional probabilities are: $p(x_t | x_{t-1}, u_t)$ and $p(z_t | x_t)$. The probability $p(x_t | x_{t-1}, u_t)$ is the *state transition probability*. It specifies how environmental state evolves over time as a function of robot controls u_t . Robot environments are stochastic, which is reflected by the fact that $p(x_t | x_{t-1}, u_t)$ is a probability distribution, not a deterministic function. Sometimes the state transition distribution does not depend on the time index t , in which case we may write it as $p(x' | u, x)$, where x' is the successor and x the predecessor state.

MEASUREMENT
PROBABILITY

The probability $p(z_t | x_t)$ is called the *measurement probability*. It also may not depend on the time index t , in which case it shall be written as $p(z | x)$. The measurement probability specifies the probabilistic law according to which measurements z are generated from the environment state x . It is appropriate to think of measurements as noisy projections of the state.

The state transition probability and the measurement probability together describe the dynamical stochastic system of the robot and its environment. Figure 2.2 illustrates the evolution of states and measurements, defined through those probabilities. The state at time t is stochastically dependent on the state at time $t-1$ and the control u_t . The measurement z_t depends stochastically on the state at time t . Such a temporal generative model is also known as *hidden Markov model* (HMM) or *dynamic Bayes network* (DBN).

2.3.4 Belief Distributions

BELIEF

Another key concept in probabilistic robotics is that of a *belief*. A belief reflects the robot's internal knowledge about the state of the environment. We already discussed that state cannot be measured directly. For example, a robot's pose might be $x_t = \langle 14.12, 12.7, 45^\circ \rangle$ in some global coordinate sys-

INFORMATION STATE

tem, but it usually cannot know its pose, since poses are not measurable directly (not even with GPS!). Instead, the robot must infer its pose from data. We therefore distinguish the true state from its internal *belief* with regards to that state. Synonyms for belief in the literature are the terms *state of knowledge* and *information state* (not to be confused with the information vector and information matrix discussed below).

Probabilistic robotics represents beliefs through conditional probability distributions. A belief distribution assigns a probability (or density value) to each possible hypothesis with regards to the true state. Belief distributions are posterior probabilities over state variables conditioned on the available data. We will denote belief over a state variable x_t by $bel(x_t)$, which is an abbreviation for the posterior

$$(2.33) \quad bel(x_t) = p(x_t | z_{1:t}, u_{1:t})$$

This posterior is the probability distribution over the state x_t at time t , conditioned on all past measurements $z_{1:t}$ and all past controls $u_{1:t}$.

The reader may notice that we silently assume that the belief is taken *after* incorporating the measurement z_t . Occasionally, it will prove useful to calculate a posterior *before* incorporating z_t , just after executing the control u_t . Such a posterior will be denoted as follows:

$$(2.34) \quad \overline{bel}(x_t) = p(x_t | z_{1:t-1}, u_{1:t})$$

PREDICTION

This probability distribution is often referred to as *prediction* in the context of probabilistic filtering. This terminology reflects the fact that $\overline{bel}(x_t)$ predicts the state at time t based on the previous state posterior, before incorporating the measurement at time t . Calculating $bel(x_t)$ from $\overline{bel}(x_t)$ is called *correction* or the *measurement update*.

2.4 Bayes Filters

2.4.1 The Bayes Filter Algorithm

BAYES FILTER

The most general algorithm for calculating beliefs is given by the *Bayes filter* algorithm. This algorithm calculates the belief distribution bel from measurement and control data. We will first state the basic algorithm and then will elucidate it with a numerical example. After that, we will derive it mathematically from the assumptions made so far.

Table 2.1 depicts the basic Bayes filter in pseudo-algorithmic form. The Bayes filter is recursive, that is, the belief $bel(x_t)$ at time t is calculated from

```

1:   Algorithm Bayes_filter( $bel(x_{t-1}), u_t, z_t$ ):
2:     for all  $x_t$  do
3:        $\overline{bel}(x_t) = \int p(x_t | u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1}$ 
4:        $bel(x_t) = \eta p(z_t | x_t) \overline{bel}(x_t)$ 
5:     endfor
6:     return  $bel(x_t)$ 

```

Table 2.1 The general algorithm for Bayes filtering.

UPDATE RULE OF A
BAYES FILTER

the belief $bel(x_{t-1})$ at time $t - 1$. Its input is the belief bel at time $t - 1$, along with the most recent control u_t and the most recent measurement z_t . Its output is the belief $bel(x_t)$ at time t . Table 2.1 only depicts a single iteration of the Bayes Filter algorithm: the *update rule*. This update rule is applied recursively, to calculate the belief $bel(x_t)$ from the belief $bel(x_{t-1})$, calculated previously.

The Bayes filter algorithm possesses two essential steps. In line 3, it processes the control u_t . It does so by calculating a belief over the state x_t based on the prior belief over state x_{t-1} and the control u_t . In particular, the belief $\overline{bel}(x_t)$ that the robot assigns to state x_t is obtained by the integral (sum) of the product of two distributions: the prior assigned to x_{t-1} , and the probability that control u_t induces a transition from x_{t-1} to x_t . The reader may recognize the similarity of this update step to Equation (2.12). As noted above, this update step is called the control update, or *prediction*.

The second step of the Bayes filter is called the *measurement update*. In line 4, the Bayes filter algorithm multiplies the belief $\overline{bel}(x_t)$ by the probability that the measurement z_t may have been observed. It does so for each hypothetical posterior state x_t . As will become apparent further below when actually deriving the basic filter equations, the resulting product is generally not a probability. It may not integrate to 1. Hence, the result is normalized, by virtue of the normalization constant η . This leads to the final belief $bel(x_t)$, which is returned in line 6 of the algorithm.

To compute the posterior belief recursively, the algorithm requires an initial belief $bel(x_0)$ at time $t = 0$ as boundary condition. If one knows the value of x_0 with certainty, $bel(x_0)$ should be initialized with a point mass distribution that centers all probability mass on the correct value of x_0 , and

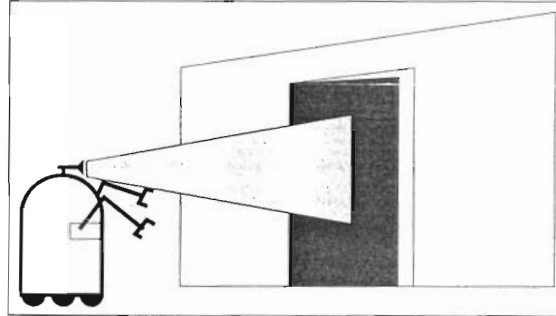


Figure 2.3 A mobile robot estimating the state of a door.

assigns zero probability anywhere else. If one is entirely ignorant about the initial value x_0 , $bel(x_0)$ may be initialized using a uniform distribution over the domain of x_0 (or a related distribution from the Dirichlet family of distributions). Partial knowledge of the initial value x_0 can be expressed by non-uniform distributions; however, the two cases of full knowledge and full ignorance are the most common ones in practice.

The algorithm Bayes filter can only be implemented in the form stated here for very simple estimation problems. In particular, we either need to be able to carry out the integration in line 3 and the multiplication in line 4 in closed form, or we need to restrict ourselves to finite state spaces, so that the integral in line 3 becomes a (finite) sum.

2.4.2 Example

Our illustration of the Bayes filter algorithm is based on the scenario in Figure 2.3, which shows a robot estimating the state of a door using its camera. To make this problem simple, let us assume that the door can be in one of two possible states, open or closed, and that only the robot can change the state of the door. Let us furthermore assume that the robot does not know the state of the door initially. Instead, it assigns equal prior probability to the two possible door states:

$$bel(X_0 = \text{open}) = 0.5$$

$$bel(X_0 = \text{closed}) = 0.5$$

Let us now assume the robot's sensors are noisy. The noise is characterized by the following conditional probabilities:

$$p(Z_t = \text{sense_open} \mid X_t = \text{is_open}) = 0.6$$

$$p(Z_t = \text{sense_closed} \mid X_t = \text{is_open}) = 0.4$$

and

$$\begin{aligned} p(Z_t = \text{sense_open} \mid X_t = \text{is_closed}) &= 0.2 \\ p(Z_t = \text{sense_closed} \mid X_t = \text{is_closed}) &= 0.8 \end{aligned}$$

These probabilities suggest that the robot's sensors are relatively reliable in detecting a *closed* door, in that the error probability is 0.2. However, when the door is open, it has a 0.4 probability of an erroneous measurement.

Finally, let us assume the robot uses its manipulator to push the door open. If the door is already open, it will remain open. If it is closed, the robot has a 0.8 chance that it will be open afterwards:

$$\begin{aligned} p(X_t = \text{is_open} \mid U_t = \text{push}, X_{t-1} = \text{is_open}) &= 1 \\ p(X_t = \text{is_closed} \mid U_t = \text{push}, X_{t-1} = \text{is_open}) &= 0 \\ p(X_t = \text{is_open} \mid U_t = \text{push}, X_{t-1} = \text{is_closed}) &= 0.8 \\ p(X_t = \text{is_closed} \mid U_t = \text{push}, X_{t-1} = \text{is_closed}) &= 0.2 \end{aligned}$$

It can also choose not to use its manipulator, in which case the state of the world does not change. This is stated by the following conditional probabilities:

$$\begin{aligned} p(X_t = \text{is_open} \mid U_t = \text{do_nothing}, X_{t-1} = \text{is_open}) &= 1 \\ p(X_t = \text{is_closed} \mid U_t = \text{do_nothing}, X_{t-1} = \text{is_open}) &= 0 \\ p(X_t = \text{is_open} \mid U_t = \text{do_nothing}, X_{t-1} = \text{is_closed}) &= 0 \\ p(X_t = \text{is_closed} \mid U_t = \text{do_nothing}, X_{t-1} = \text{is_closed}) &= 1 \end{aligned}$$

Suppose at time $t = 1$, the robot takes no control action but it senses an open door. The resulting posterior belief is calculated by the Bayes filter using the prior belief $bel(X_0)$, the control $u_1 = \text{do_nothing}$, and the measurement sense_open as input. Since the state space is finite, the integral in line 3 turns into a finite sum:

$$\begin{aligned} \overline{bel}(x_1) &= \int p(x_1 \mid u_1, x_0) bel(x_0) dx_0 \\ &= \sum_{x_0} p(x_1 \mid u_1, x_0) bel(x_0) \\ &= p(x_1 \mid U_1 = \text{do_nothing}, X_0 = \text{is_open}) bel(X_0 = \text{is_open}) \\ &\quad + p(x_1 \mid U_1 = \text{do_nothing}, X_0 = \text{is_closed}) bel(X_0 = \text{is_closed}) \end{aligned}$$

We can now substitute the two possible values for the state variable X_1 . For the hypothesis $X_1 = \text{is_open}$, we obtain

$$\begin{aligned}
 \overline{bel}(X_1 = \text{is_open}) &= p(X_1 = \text{is_open} \mid U_1 = \text{do_nothing}, X_0 = \text{is_open}) \\
 &\quad bel(X_0 = \text{is_open}) \\
 &\quad + p(X_1 = \text{is_open} \mid U_1 = \text{do_nothing}, X_0 = \text{is_closed}) \\
 &\quad bel(X_0 = \text{is_closed}) \\
 &= 1 \cdot 0.5 + 0 \cdot 0.5 = 0.5
 \end{aligned}$$

Likewise, for $X_1 = \text{is_closed}$ we get

$$\begin{aligned}
 \overline{bel}(X_1 = \text{is_closed}) &= p(X_1 = \text{is_closed} \mid U_1 = \text{do_nothing}, X_0 = \text{is_open}) \\
 &\quad bel(X_0 = \text{is_open}) \\
 &\quad + p(X_1 = \text{is_closed} \mid U_1 = \text{do_nothing}, X_0 = \text{is_closed}) \\
 &\quad bel(X_0 = \text{is_closed}) \\
 &= 0 \cdot 0.5 + 1 \cdot 0.5 = 0.5
 \end{aligned}$$

The fact that the belief $\overline{bel}(x_1)$ equals our prior belief $bel(x_0)$ should not surprise, as the action `do_nothing` does not affect the state of the world; neither does the world change over time by itself in our example.

Incorporating the measurement, however, changes the belief. Line 4 of the Bayes filter algorithm implies

$$bel(x_1) = \eta p(Z_1 = \text{sense_open} \mid x_1) \overline{bel}(x_1)$$

For the two possible cases, $X_1 = \text{is_open}$ and $X_1 = \text{is_closed}$, we get

$$\begin{aligned}
 bel(X_1 = \text{is_open}) &= \eta p(Z_1 = \text{sense_open} \mid X_1 = \text{is_open}) \overline{bel}(X_1 = \text{is_open}) \\
 &= \eta 0.6 \cdot 0.5 = \eta 0.3
 \end{aligned}$$

and

$$\begin{aligned}
 bel(X_1 = \text{is_closed}) &= \eta p(Z_1 = \text{sense_open} \mid X_1 = \text{is_closed}) \overline{bel}(X_1 = \text{is_closed}) \\
 &= \eta 0.2 \cdot 0.5 = \eta 0.1
 \end{aligned}$$

The normalizer η is now easily calculated:

$$\eta = (0.3 + 0.1)^{-1} = 2.5$$

Hence, we have

$$\begin{aligned} \text{bel}(X_1 = \text{is_open}) &= 0.75 \\ \text{bel}(X_1 = \text{is_closed}) &= 0.25 \end{aligned}$$

This calculation is now easily iterated for the next time step. As the reader easily verifies, for $u_2 = \text{push}$ and $z_2 = \text{sense_open}$ we get

$$\begin{aligned} \overline{\text{bel}}(X_2 = \text{is_open}) &= 1 \cdot 0.75 + 0.8 \cdot 0.25 = 0.95 \\ \overline{\text{bel}}(X_2 = \text{is_closed}) &= 0 \cdot 0.75 + 0.2 \cdot 0.25 = 0.05 \end{aligned}$$

and

$$\begin{aligned} \text{bel}(X_2 = \text{is_open}) &= \eta 0.6 \cdot 0.95 \approx 0.983 \\ \text{bel}(X_2 = \text{is_closed}) &= \eta 0.2 \cdot 0.05 \approx 0.017 \end{aligned}$$

At this point, the robot believes that with 0.983 probability the door is open.

At first glance, this probability may appear to be sufficiently high to simply accept this hypothesis as the world state and act accordingly. However, such an approach may result in unnecessarily high costs. If mistaking a closed door for an open one incurs costs (e.g., the robot crashes into a door), considering both hypotheses in the decision making process will be essential, as unlikely as one of them may be. Just imagine flying an aircraft on auto pilot with a perceived chance of 0.983 for not crashing!

2.4.3 Mathematical Derivation of the Bayes Filter

The correctness of the Bayes filter algorithm is shown by induction. To do so, we need to show that it correctly calculates the posterior distribution $p(x_t | z_{1:t}, u_{1:t})$ from the corresponding posterior one time step earlier, $p(x_{t-1} | z_{1:t-1}, u_{1:t-1})$. The correctness follows then by induction under the assumption that we correctly initialized the prior belief $\text{bel}(x_0)$ at time $t = 0$.

Our derivation requires that the state x_t is complete, as defined in Chapter 2.3.1, and it requires that controls are chosen at random. The first step of our derivation involves the application of Bayes rule (2.16) to the target posterior:

$$\begin{aligned} (2.35) \quad p(x_t | z_{1:t}, u_{1:t}) &= \frac{p(z_t | x_t, z_{1:t-1}, u_{1:t}) p(x_t | z_{1:t-1}, u_{1:t})}{p(z_t | z_{1:t-1}, u_{1:t})} \\ &= \eta p(z_t | x_t, z_{1:t-1}, u_{1:t}) p(x_t | z_{1:t-1}, u_{1:t}) \end{aligned}$$

We now exploit the assumption that our state is complete. In Chapter 2.3.1, we defined a state x_t to be complete if no variables prior to x_t may influence

the stochastic evolution of future states. In particular, if we (hypothetically) knew the state x_t and were interested in predicting the measurement z_t , no past measurement or control would provide us additional information. In mathematical terms, this is expressed by the following conditional independence:

$$(2.36) \quad p(z_t | x_t, z_{1:t-1}, u_{1:t}) = p(z_t | x_t)$$

Such a statement is another example of *conditional independence*. It allows us to simplify (2.35) as follows:

$$(2.37) \quad p(x_t | z_{1:t}, u_{1:t}) = \eta p(z_t | x_t) p(x_t | z_{1:t-1}, u_{1:t})$$

and hence

$$(2.38) \quad \text{bel}(x_t) = \eta p(z_t | x_t) \overline{\text{bel}}(x_t)$$

This equation is implemented in line 4 of the Bayes filter algorithm in Table 2.1.

Next, we expand the term $\overline{\text{bel}}(x_t)$, using (2.12):

$$(2.39) \quad \begin{aligned} \overline{\text{bel}}(x_t) &= p(x_t | z_{1:t-1}, u_{1:t}) \\ &= \int p(x_t | x_{t-1}, z_{1:t-1}, u_{1:t}) p(x_{t-1} | z_{1:t-1}, u_{1:t}) dx_{t-1} \end{aligned}$$

Once again, we exploit the assumption that our state is complete. This implies if we know x_{t-1} , past measurements and controls convey no information regarding the state x_t . This gives us

$$(2.40) \quad p(x_t | x_{t-1}, z_{1:t-1}, u_{1:t}) = p(x_t | x_{t-1}, u_t)$$

Here we retain the control variable u_t , since it does *not* predate the state x_{t-1} . In fact, the reader should quickly convince herself that $p(x_t | x_{t-1}, u_t) \neq p(x_t | x_{t-1})$.

Finally, we note that the control u_t can safely be omitted from the set of conditioning variables in $p(x_{t-1} | z_{1:t-1}, u_{1:t})$ for randomly chosen controls. This gives us the recursive update equation

$$(2.41) \quad \overline{\text{bel}}(x_t) = \int p(x_t | x_{t-1}, u_t) p(x_{t-1} | z_{1:t-1}, u_{1:t-1}) dx_{t-1}$$

As the reader easily verifies, this equation is implemented by line 3 of the Bayes filter algorithm in Table 2.1.

To summarize, the Bayes filter algorithm calculates the posterior over the state x_t conditioned on the measurement and control data up to time t . The derivation assumes that the world is Markov, that is, the state is complete.

Any concrete implementation of this algorithm requires three probability distributions: The initial belief $p(x_0)$, the measurement probability $p(z_t | x_t)$, and the state transition probability $p(x_t | u_t, x_{t-1})$. We have not yet specified these densities for actual robot systems. But we will soon: Chapter 5 is entirely dedicated to $p(x_t | u_t, x_{t-1})$ and Chapter 6 to $p(z_t | x_t)$. We also need a representation for the belief $bel(x_t)$, which will be discussed in Chapters 3 and 4.

2.4.4 The Markov Assumption

MARKOV ASSUMPTION

A word is in order on the *Markov assumption*, or the *complete state assumption*, since it plays such a fundamental role in the material presented in this book. The Markov assumption postulates that past and future data are independent if one knows the current state x_t . To see how severe an assumption this is, let us consider our example of mobile robot localization. In mobile robot localization, x_t is the robot's pose, and Bayes filters are applied to estimate the pose relative to a fixed map. The following factors may have a systematic effect on sensor readings. Thus, they induce violations of the Markov assumption:

- Unmodeled dynamics in the environment not included in x_t (e.g., moving people and their effects on sensor measurements in our localization example),
- inaccuracies in the probabilistic models $p(z_t | x_t)$ and $p(x_t | u_t, x_{t-1})$ (e.g., an error in the map for a localizing robot),
- approximation errors when using approximate representations of belief functions (e.g., grids or Gaussians, which will be discussed below), and
- software variables in the robot control software that influence multiple controls (e.g., the variable "target location" typically influences an entire sequence of control commands).

In principle, many of these variables can be included in state representations. However, incomplete state representations are often preferable to more complete ones to reduce the computational complexity of the Bayes filter algorithm. In practice, Bayes filters have been found to be surprisingly robust to such violations. As a general rule of thumb, however, one should exercise care when defining the state x_t , so that the effect of unmodeled state variables has close-to-random effects.

2.5 Representation and Computation

In probabilistic robotics, Bayes filters are implemented in several different ways. As we will see in the next two chapters, there exist quite a variety of techniques and algorithms that are all derived from the Bayes filter. Each such technique relies on different assumptions regarding the measurement and state transition probabilities and the initial belief. These assumptions then give rise to different types of posterior distributions, and the algorithms for computing them have different computational characteristics. As a general rule of thumb, exact techniques for calculating beliefs exist only for highly specialized cases; in general robotics problems, beliefs have to be approximated. The nature of the approximation has important ramifications on the complexity of the algorithm. Finding a suitable approximation is usually a challenging problem, with no unique best answer for all robotics problems.

When choosing an approximation, one has to trade off a range of properties:

1. **Computational efficiency.** Some approximations, such as linear Gaussian approximations that will be discussed further below, make it possible to calculate beliefs in time polynomial in the dimension of the state space. Others may require exponential time. Particle-based techniques, discussed further below, have an *any-time* characteristic, enabling them to trade off accuracy with computational efficiency.
2. **Accuracy of the approximation.** Some approximations can approximate a wider range of distributions more tightly than others. For example, linear Gaussian approximations are limited to unimodal distributions, whereas histogram representations can approximate multi-modal distributions, albeit with limited accuracy. Particle representations can approximate a wide array of distributions, but the number of particles needed to attain a desired accuracy can be large.
3. **Ease of implementation.** The difficulty of implementing probabilistic algorithms depends on a variety of factors, such as the form of the measurement probability $p(z_t | x_t)$ and the state transition probability $p(x_t | u_t, x_{t-1})$. Particle representations often yield surprisingly simple implementations for complex nonlinear systems—one of the reasons for their recent popularity.

The next two chapters will introduce concrete implementable algorithms, which fare quite differently relative to the criteria described above.

2.6 Summary

In this section, we introduced the basic idea of Bayes filters in robotics, as a means to estimate the state of an environment and the robot.

- The interaction of a robot and its environment is modeled as a coupled dynamical system, in which the robot manipulates its environment by choosing controls, and in which it can perceive the environment through its sensors.
- In probabilistic robotics, the dynamics of the robot and its environment are characterized in the form of two probabilistic laws: the state transition distribution, and the measurement distribution. The state transition distribution characterizes how state changes over time, possibly as the effect of robot controls. The measurement distribution characterizes how measurements are governed by states. Both laws are probabilistic, accounting for the inherent uncertainty in state evolution and sensing.
- The *belief* of a robot is the posterior distribution over the state of the environment (including the robot state) given all past sensor measurements and all past controls. The *Bayes filter* is the principal algorithm for calculating the belief in robotics. The Bayes filter is recursive; the belief at time t is calculated from the belief at time $t - 1$.
- The Bayes filter makes a *Markov assumption* according to which the state is a complete summary of the past. This assumption implies the belief is sufficient to represent the past history of the robot. In robotics, the Markov assumption is usually only an approximation. We identified conditions under which it is violated.
- Since the Bayes filter is not a practical algorithm, in that it cannot be implemented on a digital computer, probabilistic algorithms use tractable approximations. Such approximations may be evaluated according to different criteria, relating to their accuracy, efficiency, and ease of implementation.

The next two chapters discuss two popular families of recursive state estimation techniques that are both derived from the Bayes filter.

2.7 Bibliographical Remarks

The basic statistical material in this chapter is covered in most introductory textbooks to probability and statistics. Some early classical texts by DeGroot (1975), Subrahmaniam (1979), and Thorp (1966) provide highly accessible introductions into this material. More advanced treatments can be found in (Feller 1968; Casella and Berger 1990; Tanner 1996), and in (Devroye et al. 1996; Duda et al. 2000). The robot environment interaction paradigm is common in robotics. It is discussed from the AI perspective by Russell and Norvig (2002).

2.8 Exercises

1. A robot uses a range sensor that can measure ranges from $0m$ and $3m$. For simplicity, assume that actual ranges are distributed uniformly in this interval. Unfortunately, the sensor can be faulty. When the sensor is faulty, it constantly outputs a range below $1m$, regardless of the actual range in the sensor's measurement cone. We know that the prior probability for a sensor to be faulty is $p = 0.01$.

Suppose the robot queried its sensor N times, and every single time the measurement value is below $1m$. What is the posterior probability of a sensor fault, for $N = 1, 2, \dots, 10$. Formulate the corresponding probabilistic model.

2. Suppose we live at a place where days are either sunny, cloudy, or rainy. The weather transition function is a Markov chain with the following transition table:

		tomorrow will be...		
		sunny	cloudy	rainy
today it's...	sunny	.8	.2	0
	cloudy	.4	.4	.2
	rainy	.2	.6	.2

- (a) Suppose Day 1 is a sunny day. What is the probability of the following sequence of days: Day2 = *cloudy*, Day3 = *cloudy*, Day4 = *rainy*?
- (b) Write a simulator that can randomly generate sequences of "weathers" from this state transition function.
- (c) Use your simulator to determine the stationary distribution of this Markov chain. The stationary distribution measures the probability that a random day will be sunny, cloudy, or rainy.
- (d) Can you devise a closed-form solution to calculating the stationary distribution based on the state transition matrix above?

- (e) What is the entropy of the stationary distribution?
- (f) Using Bayes rule, compute the probability table of yesterday's weather given today's weather. (It is okay to provide the probabilities numerically, and it is also okay to rely on results from previous questions in this exercise.)
- (g) Suppose we added seasons to our model. The state transition function above would only apply to the Summer, whereas different ones would apply to Winter, Spring, and Fall. Would this violate the Markov property of this process? Explain your answer.
3. Suppose that we cannot observe the weather directly, but instead rely on a sensor. The problem is that our sensor is noisy. Its measurements are governed by the following measurement model:

		our sensor tells us...		
		sunny	cloudy	rainy
the actual weather is...	sunny	.6	.4	0
	cloudy	.3	.7	0
	rainy	0	0	1

- (a) Suppose Day 1 is sunny (this is known for a fact), and in the subsequent four days our sensor observes *cloudy, cloudy, rainy, sunny*. What is the probability that Day 5 is indeed sunny as predicted by our sensor?
- (b) Once again, suppose Day 1 is known to be sunny. At Days 2 through 4, the sensor measures *sunny, sunny, rainy*. For each of the Days 2 through 4, what is the most likely weather on that day? Answer the question in two ways: one in which only the data available to the day in question is used, and one in hindsight, where data from future days is also available.
- (c) Consider the same situation (Day 1 is sunny, the measurements for Days 2, 3, and 4 are *sunny, sunny, rainy*). What is the most likely sequence of weather for Days 2 through 4? What is the probability of this most likely sequence?
4. In this exercise we will apply Bayes rule to Gaussians. Suppose we are a mobile robot who lives on a long straight road. Our location x will simply be the position along this road. Now suppose that initially, we believe to be at location $x_{\text{init}} = 1,000m$, but we happen to know that this estimate

is uncertain. Based on this uncertainty, we model our initial belief by a Gaussian with variance $\sigma_{\text{init}}^2 = 900m^2$.

To find out more about our location, we query a GPS receiver. The GPS tells us our location is $z_{\text{GPS}} = 1,100m$. This GPS receiver is known to have an error variance of $\sigma_{\text{init}}^2 = 100m^2$.

- (a) Write the probability density functions of the prior $p(x)$ and the measurement $p(z | x)$.
- (b) Using Bayes rule, what is the posterior $p(x | z)$? Can you prove it to be Gaussian?
- (c) How likely was the measurement $x_{\text{GPS}} = 1,100m$ given our prior, and knowledge of the error probability of our GPS receiver?

Hint: This is an exercise in manipulating quadratic expressions.

5. Derive Equations (2.18) and (2.19) from (2.17) and the laws of probability stated in the text.
6. Prove Equation (2.25). What are the implications of this equality?