# Lecture notes on Basics of Sensor Fusion <br> Roland Hostettler ${ }^{1}$ and Simo Särkkä ${ }^{2}$ <br> ${ }^{1}$ Uppsala University, Sweden <br> ${ }^{2}$ Aalto University, Finland 

September 3, 2020

## Contents

1 Introduction ..... 1
1.1 Definition and Components of Sensor Fusion ..... 1
1.2 Model of a Drone ..... 2
1.3 Model of an Autonomous Car ..... 4
1.4 Dynamic Models of Drone and Car ..... 7
1.5 Tracking a Moving Drone or Car ..... 9
2 Sensors, Models, and Least Squares Criterion ..... 11
2.1 Sensors ..... 11
2.2 Models ..... 12
2.2.1 Basic Model ..... 13
2.2.2 Vector Model ..... 14
2.2.3 Multiple Measurements and Measurement Stacking ..... 14
2.2.4 Gaussian Measurement Noise ..... 15
2.3 Least Squares Methods ..... 16
2.3.1 Minimization of Cost Functions ..... 16
2.3.2 Least Squares ..... 16
2.3.3 Weighted Least Squares ..... 18
2.3.4 Regularized Least Squares ..... 19
3 Static Linear Models ..... 21
3.1 Linear Model ..... 21
3.1.1 Scalar Model ..... 21
3.1.2 Vector Model ..... 22
3.1.3 Affine Models ..... 23
3.2 Linear Least Squares ..... 24
3.2.1 Scalar Model ..... 24
3.2.2 General Linear Model ..... 27
3.2.3 Weighted Linear Least Squares ..... 29
3.3 Regularized Linear Least Squares ..... 30
3.4 Sequential Linear Least Squares ..... 33
4 Static Nonlinear Models ..... 37
4.1 Nonlinear Model ..... 37
4.2 Gradient Descent ..... 38
4.3 Gauss-Newton Algorithm ..... 41
4.4 Gauss-Newton Algorithm with Line Search ..... 43
4.5 Levenberg-Marquardt Algorithm ..... 45
4.6 Regularized Non-Linear Models ..... 48
4.7 Quasi-Newton Methods ..... 49
4.8 Convergence Criteria ..... 51
5 Dynamic Models ..... 53
5.1 Continuous-Time State-Space Models ..... 53
5.1.1 Deterministic Linear State-Space Models ..... 53
5.1.2 Stochastic Linear State-Space Models ..... 56
5.1.3 Nonlinear State-Space Models ..... 58
5.2 Discrete-Time State-Space Models ..... 61
5.2.1 Deterministic Linear State-Space Models ..... 61
5.2.2 Stochastic Linear Dynamic Models ..... 62
5.2.3 Nonlinear Dynamic Model ..... 63
5.3 Discretization of Linear Dynamic Models ..... 64
5.3.1 Deterministic Linear Dynamic Models ..... 64
5.3.2 Stochastic Linear Dynamic Models ..... 67
5.4 Discretization of Nonlinear Dynamic Models ..... 69
5.4.1 Discretization of Linearized Nonlinear Models ..... 69
5.4.2 Euler Discretization of Deterministic Nonlinear Models ..... 70
5.4.3 Euler-Maruyama Discretization of Stochastic Nonlinear Models ..... 70
6 Filtering ..... 73
6.1 The Filtering Approach ..... 73
6.2 Kalman Filter ..... 75
6.3 Extended Kalman Filter ..... 79
6.4 Unscented Kalman Filtering ..... 83
6.5 Iterated Extended and Unscented Kalman Filters ..... 87
6.6 Bootstrap Particle Filter ..... 87

## Preface

These lecture notes are aimed for an M.Sc. level course on sensor fusion. The goal is to provide a gentle and pragmatic introduction to the topic with an emphasis on methods that are useful for developing practical solutions to sensor fusion problems. The reader should be familiar with:

1. linear algebra and calculus,
2. ordinary differential equations, and
3. basic statistics.

The notes are based on the much more comprehensive and rigorous works on sensor fusion, estimation theory, filtering theory, stochastic process theory, and optimization theory by Gustafsson (2018), Kay (1993), Särkkä (2013), Särkkä and Solin (2019), and Nocedal and Wright (2006), and the interested reader is highly recommended to have a read in these books.

The topics covered in these notes range from simple linear observation models and least squares estimation, to modeling of dynamic systems, and inference in dynamic systems using Kalman and particle filtering. To lower the threshold to the topic, we sometimes trade mathematical rigor for easier understanding. This is in particular true for the probabilistic interpretations, especially in the beginning of the notes.

## Chapter 1

## Introduction

### 1.1 Definition and Components of Sensor Fusion

Sensor fusion refers to computational methodology which aims at combining the measurements from multiple sensors such that they jointly give more information on the measured system than any of the sensors alone. Sensor fusion has applications in many different areas of daily life and plays an important role in modern society. For example, it is used to interpret traffic scenes in autonomous cars, for navigation and localization in robotics, for controlling drones in aerial photography and deliveries, and in the analysis of biosignals in biomedical engineering.


Figure 1.1. The basic components of a sensor fusion system.

The common denominator and main objective of sensor fusion systems are that they take measurements from different sensors and estimate or infer one or more quantities of interest. For example, we can use multiple sensors to infer the position of a drone or an autonomous car - possibly along with their velocities. On a high level, this involves three main components (Figure 1.1):

- One or more $\operatorname{sensor}(s)$ that measure an observable quantity,
- model(s) that relate the observed quantity to the quantity of interest, and
- an estimation algorithm that estimates the quantity of interest by combining the model and the measured data.


Figure 1.2. A simple illustration of fusion of multiple sensor measurements made by a drone. The height is measured with one sensor (say, barometer) and the distance from a wall with another sensor (say, radar). The "fusion" of the measurements in this case simply means using both the measurements together to determine the drone's position.

### 1.2 Model of a Drone

As an illustration of (simple) sensor fusion let us take a look at the Figure 1.2 with a drone whose position $\left(p^{x}, p^{y}\right)$ we wish to determine. In this case we have two sensors: one sensor which measures the distance to a nearby wall - this sensor could be, for example, radar or ultrasound-based sensor; and a second sensor which measures the distance to the ground, that is, height - this sensor could be, for example, radar or barometer. In this case neither of the sensors alone gives the position of the drone. However, as the wall-distance sensor directly measures the horizontal position $p^{x}$ of the drone and the height sensor directly measures the vertical position $p^{y}$ of the drone, we can "fuse" the sensor measurement simply by using the measurements jointly to determine the $\left(p^{x}, p^{y}\right)$ position. In this case the fusion itself is trivial as no actual computations needs to be done to the measurements since they directly measure the components of the quantity of interest, that is, the position.

If we denote the sensor reading measuring the distance to the wall as $y_{1}$ and the height sensor measurement as $y_{2}$, then a mathematical model which relates the observed quantities (measurements) to the quantity of interests is

$$
\begin{align*}
& y_{1}=p^{x} \\
& y_{2}=p^{y} \tag{1.1}
\end{align*}
$$

and the aim of the estimation algorithm is to determine $\left(p^{x}, p^{y}\right)$ from the measurements $\left(y_{1}, y_{2}\right)$. In this case it is indeed trivial as the measurements directly give the position coordinates.

In reality, the situation is not as simple as all sensors contain measurement noise and a more proper model is actually

$$
\begin{align*}
& y_{1}=p^{x}+r_{1} \\
& y_{2}=p^{y}+r_{2} \tag{1.2}
\end{align*}
$$

which says that the measurements that we obtain are the actual positions plus measurement noises $r_{1}$ and $r_{2}$. The measurement noises are typically modeled as zeromean Gaussian random variables. In this case we cannot exactly determine the position from the sensor measurements because we only get to see noisy versions of the positions. We could, however, approximate that the noises are "small" which implies that $y_{1} \approx p^{x}, y_{2} \approx p^{y}$ and thus we can use the measurement directly as estimates of the position coordinates. In this case given only the two sensor measurements, it is hard or impossible to get rid of the noise and hence our position estimate will inherently be off by the amount of measurement noise - not even sophisticated statistical methods will be significantly better than the small-noise approximation. One way to cope with this problem is to average multiple measurements which reduces the variance of the estimate, but it obviously requires us to make more measurements with the sensor. Alternatively, by adding more sensors we can diminish the effect of measurement noise even when each of the sensors only produce a single measurement.

Let us now look at a slightly extended sensor setup show in Figure 1.3. We still have the wall-distance and height sensors, but additionally, we measure the distance from a wall that is at a $45^{\circ}$ angle with the ground. From the geometry of the problem we can now determine that the relationship of the measurements and positions should be

$$
\begin{align*}
& y_{1}=p^{x}+r_{1} \\
& y_{2}=p^{y}+r_{2} \\
& y_{3}=\frac{1}{\sqrt{2}}\left(p^{x}-x_{0}\right)+\frac{1}{\sqrt{2}} p^{y}+r_{3} \tag{1.3}
\end{align*}
$$

where $x_{0}$ is the horizontal coordinate of the rightmost point of the wall and $r_{1}, r_{2}$, and $r_{3}$ are measurement noises.

The aim of the estimation algorithm is now to determine the position $\left(p^{x}, p^{y}\right)$ given the 3 sensor measurements $\left(y_{1}, y_{2}, y_{3}\right)$. If we had no measurement noises, then any two of the measurements would exactly give us the position (after possibly solving a simultaneous pair of equations). It is indeed this over-completeness of the model that allows us to use least-squares estimation to determine the position from the measurements such that the effect of the measurements noises is significantly diminished - the more over-complete the model is, the better we can diminish the noises. Furthermore, an important factor in the above model is that the measurement $y_{3}$ is no longer a direct measurement of a position coordinate, but a linear combination of them, and in that sense it is an indirect observation of the position.


Figure 1.3. An illustration of sensor fusion with an additional sensor as compared to Figure 1.2. When we measure the drone position using 3 sensors instead of 2, the overcompleteness of the problem enables us to diminish the noises in the sensor measurements.

The model above (and in fact the more simple model also) is an example of static linear model which will be covered in Chapter 3. More specifically, the model above in Equation (1.3) can be also rewritten in vector notation by rewriting it as

$$
\underbrace{\left[\begin{array}{l}
y_{1}  \tag{1.4}\\
y_{2} \\
y_{3}
\end{array}\right]}_{\mathbf{y}}=\underbrace{\left[\begin{array}{cc}
1 & 0 \\
0 & 1 \\
\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}}
\end{array}\right]}_{\mathbf{G}} \underbrace{\left[\begin{array}{c}
p^{x} \\
p^{y}
\end{array}\right]}_{\mathbf{x}}+\underbrace{\left[\begin{array}{c}
0 \\
0 \\
-\frac{x_{0}}{\sqrt{2}}
\end{array}\right]}_{\mathbf{b}}+\underbrace{\left[\begin{array}{c}
r_{1} \\
r_{2} \\
r_{3}
\end{array}\right]}_{\mathbf{r}}
$$

which has the form

$$
\begin{equation*}
\mathbf{y}=\mathbf{G} \mathbf{x}+\mathbf{b}+\mathbf{r} \tag{1.5}
\end{equation*}
$$

and the estimation aim is to determine the unknown position vector $\mathbf{x}$ given the measurement vector $\mathbf{y}$ and the parameters $\mathbf{G}$ and $\mathbf{b}$. It is worthwhile to notice that the matrix $\mathbf{G}$ is neither square nor invertible and additionally, the noise vector $\mathbf{r}$ is unknown. Hence, we cannot just determine x by multiplying both sides with inverse of $\mathbf{G}$. Instead, we will use a least squares method to determine $\mathbf{x}$.

### 1.3 Model of an Autonomous Car

Driver-assistance and self-driving systems in cars are also important applications of sensor fusion. Let us now take a look at an autonomous car illustrated in Figure 1.4 where the car is determining its position by measuring the direction and/or range (i.e., distance) to nearby landmarks whose positions are known. This kind of measurement could be done, for example, with a camera-based computer vision system (see, e.g., Hartley and Zisserman, 2003) that detects the landmarks in the image and determines their direction and range. Another option to measure them would be, for example, by using a radar (see, e.g., Richards et al., 2010).


Figure 1.4. Illustration of positioning of an autonomous car from measurements of relative locations of landmarks (e.g., traffic signs). When only ranges or directions of the landmarks are measured, then the sensor fusion model becomes non-linear.

Even though, for example, radar or camera actually measures the direction and range from the car to the landmark, provided that both the measurements are relatively accurate, then by a simple coordinate transformation we can convert the measurement into measurement of the landmark position in the car's coordinate system. Thus we get direct measurements of the relative positions of the landmarks with respect o the car. In this case, if the landmark positions are $\left(s_{j}^{x}, s_{j}^{y}\right)$ for $j=1, \ldots, M$, and the position of the car is $\left(p^{x}, p^{y}\right)$, then the measurements are given as

$$
\begin{gather*}
y_{1}=s_{1}^{x}-p^{x}+r_{1} \\
y_{2}=s_{1}^{y}-p^{y}+r_{2} \\
\vdots \\
y_{2 M-1}=s_{M}^{x}-p^{x}+r_{2 M-1}, \\
y_{2 M}=s_{M}^{y}-p^{y}+r_{2 M}, \tag{1.6}
\end{gather*}
$$

where the odd-numbered measurements corresponds on the horizontal measurements and even-numbered the vertical measurements in the image. We can again
rewrite the model in form (1.5) as follows:

$$
\underbrace{\left[\begin{array}{c}
y_{1}  \tag{1.7}\\
y_{2} \\
\vdots \\
y_{2 M-1} \\
y_{2 M}
\end{array}\right]}_{\mathbf{y}}=\underbrace{\left[\begin{array}{cc}
-1 & 0 \\
0 & -1 \\
\vdots & \vdots \\
-1 & 0 \\
0 & -1
\end{array}\right]}_{\mathbf{G}} \underbrace{\left[\begin{array}{c}
p^{x} \\
p^{y}
\end{array}\right]}_{\mathbf{x}}+\underbrace{\left[\begin{array}{c}
s_{1}^{x} \\
s_{1}^{y} \\
\vdots \\
s_{M}^{x} \\
s_{M}^{y}
\end{array}\right]}_{\mathbf{b}}+\underbrace{\left[\begin{array}{c}
r_{1} \\
r_{2} \\
\vdots \\
r_{2 M-1} \\
r_{2 M}
\end{array}\right]}_{\mathbf{r}},
$$

where the matrix $\mathbf{G}$ is again non-square and non-invertible. However, least squares method will be applicable to estimate the car's position $\mathbf{x}$ from the noisy measurements $\mathbf{y}$.

Sometimes it happens that the sensor only measures the range from car to the landmark and does not provide information on the direction. This can happen, for example, when a certain type of radar is used. When only range measurements are available, then the range-only model can be written as

$$
\begin{align*}
y_{1}^{R} & =\sqrt{\left(s_{1}^{x}-p^{x}\right)^{2}+\left(s_{1}^{y}-p^{y}\right)^{2}}+r_{1}^{R} \\
& \vdots  \tag{1.8}\\
y_{M}^{R} & =\sqrt{\left(s_{M}^{x}-p^{x}\right)^{2}+\left(s_{M}^{y}-p^{y}\right)^{2}}+r_{M}^{R}
\end{align*}
$$

and the aim of the measurement algorithm would be to determine the position ( $p^{x}, p^{y}$ ) based on the measurements $y_{1}^{R}, \ldots, y_{M}^{R}$ in the presence of the measurement noises $r_{1}^{R}, \ldots, r_{M}^{R}$. If we now define a function

$$
\mathbf{g}(\mathbf{x})=\left[\begin{array}{c}
\sqrt{\left(s_{1}^{x}-p^{x}\right)^{2}+\left(s_{1}^{y}-p^{y}\right)^{2}}  \tag{1.9}\\
\vdots \\
\sqrt{\left(s_{M}^{x}-p^{x}\right)^{2}+\left(s_{M}^{y}-p^{y}\right)^{2}}
\end{array}\right]
$$

as well as $\mathbf{x}=\left[\begin{array}{ll}p^{x} & p^{y}\end{array}\right]^{\top}, \mathbf{y}=\left[\begin{array}{lll}y_{1}^{R} & \cdots & y_{M}^{R}\end{array}\right]^{\top}$, and $\mathbf{r}=\left[\begin{array}{lll}r_{1}^{R} & \cdots & r_{M}^{R}\end{array}\right]^{\top}$, then the model can be written as

$$
\begin{equation*}
\mathbf{y}=\mathbf{g}(\mathbf{x})+\mathbf{r} \tag{1.10}
\end{equation*}
$$

where $\mathbf{g}(\mathbf{x})$ is a non-linear function of $\mathbf{x}$.
Alternatively, let us now consider a case where we only measure the direction (i.e., bearing) of the landmark with respect to the car, but no distance. This can happen, for example, when we are using a camera without the knowledge of the true size of the landmark. The measurement in this bearings-only model can be written as

$$
\begin{align*}
& y_{1}^{D}=\operatorname{atan}_{2}\left(s_{1}^{y}-p^{y}, s_{1}^{x}-p^{x}\right)+r_{1}^{D} \\
& \quad \vdots  \tag{1.11}\\
& y_{M}^{D}=\operatorname{atan}_{2}\left(s_{M}^{y}-p^{y}, s_{M}^{x}-p^{x}\right)+r_{M}^{D},
\end{align*}
$$

where atan $_{2}$ is the full-quadrant inverse of tangent function (e.g., function called at an2 in Matlab and Python). This model again has the form (1.10) provided that we define

$$
\mathbf{g}(\mathbf{x})=\left[\begin{array}{c}
\operatorname{atan}_{2}\left(s_{1}^{y}-p^{y}, s_{1}^{x}-p^{x}\right)  \tag{1.12}\\
\vdots \\
\operatorname{atan}_{2}\left(s_{M}^{y}-p^{y}, s_{M}^{x}-p^{x}\right)
\end{array}\right]
$$

along with $\mathbf{x}=\left[\begin{array}{ll}p^{x} & p^{y}\end{array}\right]^{\top}, \mathbf{y}=\left[\begin{array}{lll}y_{1}^{D} & \cdots & y_{M}^{D}\end{array}\right]^{\top}$, and $\mathbf{r}=\left[\begin{array}{lll}r_{1}^{D} & \cdots & r_{M}^{D}\end{array}\right]^{\top}$.
The range-only and bearings-only models above are examples of static nonlinear models that we cover in Chapter 4. The estimation of $\mathbf{x}$ in these models will be possible by solving the nonlinear least squares problem with an iterative optimization method.

### 1.4 Dynamic Models of Drone and Car

The models of the drone and car in the previous sections are static in the sense that the models do not model the time behaviour of their positions. Of course, if the position changes, then we can use new sensor measurements to recalculate the position. However, in this kind of approach, where we recompute the position from scratch always when the device moves, we lose the information contained in the sensor measurements that we did before. In order to use sensor measurements over time, we need to build a dynamic model which ties the positions at different time points together.

One simple model can be constructed by assuming that the position $\mathbf{x}=$ $\left[\begin{array}{ll}p^{x} & p^{y}\end{array}\right]^{\top}$ could change to any other position near the current position on the next time step. One way to model this is that if the position at time step $n-1$ was $\mathbf{x}_{n-1}$, then the next time step is given as

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{x}_{n-1}+\mathbf{q}_{n} \tag{1.13}
\end{equation*}
$$

where $\mathbf{q}_{n}$ is a zero-mean random variable, which is typically modeled as Gaussian. The model says that we do not know where the next position will be, but the closer positions have higher probability than the further away positions. This model is also called a random walk model, because it corresponds to random steps to random directions that are taken at each time step transition. In components of the vector $\mathbf{x}_{n}$ this model can be written as

$$
\begin{align*}
& p_{n}^{x}=p_{n-1}^{x}+q_{1, n} \\
& p_{n}^{y}=p_{n-1}^{y}+q_{2, n} \tag{1.14}
\end{align*}
$$

where we need to use two indices for denoting the vector index and the time step number.

Although a random walk model would be good for say a person or animal, both drone and car have definite velocity which itself changes according to a random
walk model. At the time step transition the velocity times the time step length $\Delta t$ is added to the position coordinates which leads to the model

$$
\begin{align*}
p_{n}^{x} & =p_{n-1}^{x}+v_{n-1}^{x} \Delta t+q_{1, n}, \\
p_{n}^{y} & =p_{n-1}^{y}+v_{n-1}^{y} \Delta t+q_{2, n}, \\
v_{n}^{x} & =v_{n-1}^{x}+q_{3, n}, \\
v_{n}^{y} & =v_{n-1}^{y}+q_{4, n}, \tag{1.15}
\end{align*}
$$

where we have added noises $q_{3, n}$ and $q_{4, n}$ which act as the noises in the velocity components.

At this point it is now convenient to change the notation so that our state vector x not only contains the position coordinates but also the velocities $\mathrm{x}=$ $\left[\begin{array}{llll}p^{x} & p^{y} & v^{x} & v^{y}\end{array}\right]^{\top}$. As the velocities are just 3rd and 4th components of $\mathbf{x}$, we can also then rewrite above model as

$$
\begin{align*}
x_{1, n} & =x_{1, n-1}+x_{3, n-1} \Delta t+q_{1, n}, \\
x_{2, n} & =x_{2, n-1}+x_{4, n-1} \Delta t+q_{2, n}, \\
x_{3, n} & =x_{3, n-1}+q_{3, n}, \\
x_{4, n} & =x_{4, n-1}+q_{4, n}, \tag{1.16}
\end{align*}
$$

which can also be written in convenient matrix form as

$$
\underbrace{\left[\begin{array}{l}
x_{1, n}  \tag{1.17}\\
x_{2, n} \\
x_{3, n} \\
x_{4, n}
\end{array}\right]}_{\mathbf{x}_{n}}=\underbrace{\left[\begin{array}{cccc}
1 & 0 & \Delta t & 0 \\
0 & 1 & 0 & \Delta t \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right]}_{\mathbf{F}} \underbrace{\left[\begin{array}{l}
x_{1, n-1} \\
x_{2, n-1} \\
x_{3, n-1} \\
x_{4, n-1}
\end{array}\right]}_{\mathbf{x}_{n-1}}+\underbrace{\left[\begin{array}{l}
q_{1, n} \\
q_{2, n} \\
q_{3, n} \\
q_{4, n}
\end{array}\right]}_{\mathbf{q}_{n}},
$$

that is,

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{F} \mathbf{x}_{n-1}+\mathbf{q}_{n} . \tag{1.18}
\end{equation*}
$$

The above model is an example of a discrete-time stochastic linear state space that we consider in Chapter 5.

The above model could also be derived by considering the state as function of continuous-time variable $\mathbf{x}(t)$ with $t$ being a time variable that can take any value $t \geq 0$. The discrete-time model can then be derived as a discretization of the continuous-time model described as a differential equation with random input. For example, the first random walk model corresponds to a continuous-time model

$$
\begin{align*}
\frac{\mathrm{d} p^{x}(t)}{\mathrm{d} t} & =w_{1}(t), \\
\frac{\mathrm{d} p^{y}(t)}{\mathrm{d} t} & =w_{2}(t), \tag{1.19}
\end{align*}
$$

where $w_{1}(t)$ and $w_{2}(t)$ are continuous time white noises.

The model with velocity components included can be derived as discretization of the second order model

$$
\begin{align*}
& \frac{\mathrm{d}^{2} p^{x}(t)}{\mathrm{d} t^{2}}=w_{1}(t), \\
& \frac{\mathrm{d}^{2} p^{y}(t)}{\mathrm{d} t^{2}}=w_{2}(t) . \tag{1.20}
\end{align*}
$$

We will come back to these in Chapter 5.

### 1.5 Tracking a Moving Drone or Car

We can now combine the dynamic model of a drone or car in Section 1.4 with a measurement model of a drone or car in Sections 1.2 or 1.3. For example, we can combine the dynamic model in (1.18) with a linear measurement model of the form (1.5), which leads to

$$
\begin{align*}
& \mathbf{x}_{n}=\mathbf{F} \mathbf{x}_{n-1}+\mathbf{q}_{n}, \\
& \mathbf{y}_{n}=\mathbf{G} \mathbf{x}_{n}+\mathbf{b}+\mathbf{r}_{n}, \tag{1.21}
\end{align*}
$$

where we get a measurement of the state at time step $n$ when the state $\mathbf{x}_{n}$ is changing according to the dynamic model. For this kind of linear (or in fact affine) dynamic models we can use the Kalman filter to compute the statistical estimate of the dynamic state $\mathbf{x}_{n}$ at each time step $n$. This methodology is covered in Chapter 6.

We can also combine the dynamic model (1.18) with a non-linear measurement model of the form (1.10), which leads to a model of the form

$$
\begin{align*}
& \mathbf{x}_{n}=\mathbf{F} \mathbf{x}_{n-1}+\mathbf{q}_{n} \\
& \mathbf{y}_{n}=\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n} \tag{1.22}
\end{align*}
$$

For this kind of models we can use extended Kalman filter (EKF), unscented Kalman filter (UKF), or particle filter, which will also be covered in Chapter 6. Additionally, these methods can also be used in the case that the dynamics are modeled with nonlinear models instead of a linear model as we have above.

## Chapter 2

## Sensors, Models, and Least Squares Criterion

### 2.1 Sensors

A sensor is a device that provides a measurement related to the quantity of interest. Usually, a sensor is implemented as a device which converts a physical phenomenon into an electrical signal (Wilson, 2005) which is then further transformed into digital form. The measurement may be direct or indirect: In direct measurements, the quantity observed by the sensor is the quantity of interest, for example, the temperature in a thermometer. Sensors that measure indirectly provide a measurement of a quantity that is only related to the quantity of interest. For example, we might only be able to measure the distance to an object when we are actually interested in the position. The indirect measurement can also be related to dynamical properties of the quantity, for example, an accelerometer measures a body's acceleration which is the second derivative of the position.

There are many different types of sensors for measuring a wide range of phenomena available today. These include electronic sensors, (micro-)mechanical sensors, or virtual sensors and a (non-exhaustive) overview of a few commonly used sensors, their measurements, and some of their application areas is shown in Table 2.1.

Sensors are characterized by many different properties that affect their performance. These include limitations such as measurement range, requirements and sensitivity for the environmental conditions (temperature, humidity, etc.), sampling frequency, as well as measurement noise, biases, and drifts. Some of these factors have to be taken into account when designing the hardware whereas others can be accurately dealt with when designing the estimation algorithm. In this course, we will focus on the latter aspects of the sensors, that is, we will introduce tools that are appropriate for taking into account the properties that can be handled by sensor fusion.

Table 2.1. Examples of common sensors

| Sensor | Measurement | Application Examples |
| :--- | :--- | :--- |
| Accelerometer | Gravity, acceleration | Inertial navigation, activity track- <br> ing, screen rotation |
| Gyroscope <br> Magnetometer | Rotational velocity <br> Magnetic field strength | Inertial navigation, activity tracking <br> Inertial navigation, digital compass, <br> object tracking |
| Radar | Range, bearing, speed | Target tracking, autonomous vehi- <br> cles |
| LIDAR | Range, bearing, speed | Target tracking, autonomous vehi- <br> cles, robotics <br> Robotics |
| Ultrasound | Range | Security systems, autonomous vehi- <br> cles, robotics |
| Camera | Visual scene | Inertial navigation, autonomous ve- <br> hicles, robotics |
| Barometer | Air pressure | Autonomous vehicles, aerospace <br> applications |
| GNSS | Position | Condition monitoring, scales |

The measurement obtained by a sensor will be denoted using the symbol $y_{n}$ or $\mathbf{y}_{n}$. The subscript $n$ denotes the $n$th measurement, which may refer to an arbitrary measurement index, step number in time series, a sensor identification number in a sensor network, depending on the context. The regular notation $\left(y_{n}\right)$ is used to denote scalar measurements (e.g., from a temperature sensor) whereas the bold symbol ( $\mathbf{y}_{n}$ ) denotes a vector observation (e.g., as measured by a tri-axial accelerometer).

### 2.2 Models

A model is a mathematical formulation that relates the quantity of interest to the measurements in a systematic way. Additionally, in the dynamic case, models are also used to describe how the quantity of interest evolves over time. We already saw examples of models in Chapter 1.

As mentioned earlier, sensors suffer from several limitations that we need to account for in models. In particular, it is important to take uncertainties in the measurements such as sensor noise or measurement noise into account. Unfortunately, measurement noise is difficult to quantify analytically. Instead, a suitable approach is to use statistical modeling. It is based on the observations that the measurement noise exhibits statistical properties that can be quantified. While the actual value of the noise can not be observed or measured, its statistical properties can, and this
allows us to design sensor fusion methods that are able to (statistically) minimize the effect of the noise. A common assumption is that it is zero mean, that is, on average, it is zero (but each individual realization is not).

### 2.2.1 Basic Model

Next, we introduce a very simplistic model, which turns out to be quite generic, despite its simplicity. First, assume that a scalar measurement $y_{n}$ of the quantity of interests $\mathbf{x}$ (e.g., the position) is available. The objective is then to relate the measurement $y_{n}$ to the unknown quantity $\mathbf{x}$ such that the uncertainty is taken into account. We know that the measurement might be some function of the unknown quantity plus measurement noise. Hence, we can write the model as

$$
\begin{equation*}
y_{n}=g_{n}(\mathbf{x})+r_{n} \tag{2.1}
\end{equation*}
$$

Equation (2.1) is called a sensor model, measurement model, or observation model. The generic anatomy of a measurement model is that

- the measured quantity $y_{n}$ is found on the left hand side of the equation ${ }^{1}$, and
- on the right hand side there are two terms, a function $g_{n}(\mathbf{x})$ that relates the quantities of interest $\mathbf{x}$ to the measurement $y_{n}$ and a noise term $r_{n}$.

As mentioned above, the measurement noise $r_{n}$ is assumed to be a random variable. As such, $r_{n}$ follows some kind of probability density function (pdf), that is, we have that

$$
\begin{equation*}
r_{n} \sim p\left(r_{n}\right) \tag{2.2}
\end{equation*}
$$

which reads as " $r_{n}$ is distributed according to $p\left(r_{n}\right)$ " where $p\left(r_{n}\right)$ denotes the corresponding pdf. At this point, we will not concern ourselves with any particular form of $p\left(r_{n}\right)$ but only assume that the $r_{n} \mathrm{~s}$ are zero-mean, independent random variables with variance $\sigma_{r, n}^{2}$. In other words, we have that

$$
\begin{align*}
\mathrm{E}\left\{r_{n}\right\} & =0  \tag{2.3a}\\
\operatorname{var}\left\{r_{n}\right\} & =\mathrm{E}\left\{r_{n}^{2}\right\}-\left(\mathrm{E}\left\{r_{n}\right\}\right)^{2}=\sigma_{r, n}^{2}  \tag{2.3b}\\
\operatorname{Cov}\left\{r_{m}, r_{n}\right\} & =\mathrm{E}\left\{r_{m} r_{n}\right\}-\mathrm{E}\left\{r_{m}\right\} \mathrm{E}\left\{r_{n}\right\}=0, \quad(m \neq n), \tag{2.3c}
\end{align*}
$$

where $\mathrm{E}\left\{r_{n}\right\}$ and $\operatorname{var}\left\{r_{n}\right\}$ denote the expected value and variance of $r_{n}$, respectively. Note that the measurement noise variance may vary for the different measurements, as indicated by the subscript $n$ on $\sigma_{r, n}^{2}$ (e.g., if the different measurements are obtained by different sensors).

[^0]
### 2.2.2 Vector Model

The basic model only considers scalar measurements. However, many sensors actually provide vector-valued measurements, for example, accelerometers provide full 3D-acceleration measurements at each sampling instant. The basic model (2.1) can easily be extended to account for such vector-valued measurements. In this case, the vector measurement model can be written as

$$
\begin{equation*}
\mathbf{y}_{n}=\mathbf{g}_{n}(\mathbf{x})+\mathbf{r}_{n} \tag{2.4}
\end{equation*}
$$

where $\mathbf{y}_{n}$ is an $d_{y}$-dimensional column vector. In this case, the measurement noise $\mathbf{r}_{n}$ becomes a multivariate random variable with pdf

$$
\begin{equation*}
\mathbf{r}_{n} \sim p\left(\mathbf{r}_{n}\right) \tag{2.5}
\end{equation*}
$$

As for the scalar case, $p\left(\mathbf{r}_{n}\right)$ can have any arbitrary form. For simplicity, we again assume that the $\mathbf{r}_{n}$ s are zero-mean, independent random variables with covariance matrix $\mathbf{R}_{n}$, that is,

$$
\begin{align*}
\mathrm{E}\left\{\mathbf{r}_{n}\right\} & =0,  \tag{2.6a}\\
\operatorname{Cov}\left\{\mathbf{r}_{n}\right\} & =\mathrm{E}\left\{\mathbf{r}_{n} \mathbf{r}_{n}^{\top}\right\}-\mathrm{E}\left\{\mathbf{r}_{n}\right\} \mathrm{E}\left\{\mathbf{r}_{n}\right\}^{\top}=\mathbf{R}_{n},  \tag{2.6b}\\
\operatorname{Cov}\left\{\mathbf{r}_{m}, \mathbf{r}_{n}\right\} & =\mathrm{E}\left\{\mathbf{r}_{m} \mathbf{r}_{n}^{\top}\right\}-\mathrm{E}\left\{\mathbf{r}_{m}\right\} \mathrm{E}\left\{\mathbf{r}_{n}\right\}^{\top}=0, \quad(m \neq n) \tag{2.6c}
\end{align*}
$$

As it can be seen, the scalar model in (2.1) is just a special case of the vector model in (2.4), where the measurement and noise are scalars (i.e., $d_{y}=1$ ). Hence, it is often more general to work with vector model, but sometimes more instructive to work with the scalar model.

### 2.2.3 Multiple Measurements and Measurement Stacking

In order to be able to fuse the measurements of one or more sensors, we need multiple measurements. These may either be from one sensor at different time instants (repeated measurements), from different sensors, or both. If we have obtained a total of $N$ measurements $y_{1}, y_{2}, \ldots, y_{N}$, we will refer to the set of all measurements using the notation $y_{1: N}=\left\{y_{1}, y_{2}, \ldots, y_{N}\right\}$ for the scalar case and $\mathbf{y}_{1: N}=\left\{\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{N}\right\}$ for the vector case.

Sometimes, it is also helpful to write the complete set of measurements together with the measurement model more compactly. This can be achieved by stacking the measurements into a single measurement vector, which yields a much more compact model of the form

$$
\begin{equation*}
\mathbf{y}=\mathbf{g}(\mathbf{x})+\mathbf{r} \tag{2.7}
\end{equation*}
$$

For scalar measurements of the form (2.1), we have that

$$
\mathbf{y}=\left[\begin{array}{c}
y_{1}  \tag{2.8}\\
y_{2} \\
\vdots \\
y_{N}
\end{array}\right], \mathbf{g}(\mathbf{x})=\left[\begin{array}{c}
g_{1}(\mathbf{x}) \\
g_{2}(\mathbf{x}) \\
\vdots \\
g_{N}(\mathbf{x})
\end{array}\right], \text { and } \mathbf{r}=\left[\begin{array}{c}
r_{1} \\
r_{2} \\
\vdots \\
r_{N}
\end{array}\right]
$$

Furthermore, the overall noise covariance $\mathbf{R}$ for this model is a diagonal matrix of the form

$$
\mathbf{R}=\operatorname{Cov}\{\mathbf{r}\}=\left[\begin{array}{cccc}
\sigma_{r, 1}^{2} & 0 & \ldots & 0  \tag{2.9}\\
0 & \sigma_{r, 2}^{2} & & \vdots \\
\vdots & & \ddots & 0 \\
0 & \ldots & 0 & \sigma_{r, N}^{2}
\end{array}\right]
$$

For vector measurements, $\mathbf{y}, \mathbf{g}(\mathbf{x})$, and $\mathbf{r}$ are constructed in the same way as for the scalar case in (2.8), that is,

$$
\mathbf{y}=\left[\begin{array}{c}
\mathbf{y}_{1}  \tag{2.10}\\
\mathbf{y}_{2} \\
\vdots \\
\mathbf{y}_{N}
\end{array}\right], \mathbf{g}(\theta)=\left[\begin{array}{c}
g_{1}(\mathbf{x}) \\
g_{2}(\mathbf{x}) \\
\vdots \\
g_{N}(\mathbf{x})
\end{array}\right], \text { and } \mathbf{r}=\left[\begin{array}{c}
\mathbf{r}_{1} \\
\mathbf{r}_{2} \\
\vdots \\
\mathbf{r}_{N}
\end{array}\right]
$$

The measurement noise covariance is a block-diagonal matrix with the individual covariance matrices $\mathbf{R}_{n}$ on the diagonal, that is,

$$
\mathbf{R}=\operatorname{Cov}\{\mathbf{r}\}=\left[\begin{array}{cccc}
\mathbf{R}_{1} & 0 & \ldots & 0  \tag{2.11}\\
0 & \mathbf{R}_{2} & & \vdots \\
\vdots & & \ddots & 0 \\
0 & \ldots & 0 & \mathbf{R}_{N}
\end{array}\right]
$$

### 2.2.4 Gaussian Measurement Noise

A common assumption is that the measurement noise is zero-mean Gaussian noise. In this case, the pdf of the noise is the (multivariate) Gaussian distribution given by

$$
\begin{equation*}
p(\mathbf{r})=\frac{1}{(2 \pi)^{M / 2}|\mathbf{R}|^{1 / 2}} \exp \left(-\frac{1}{2} \mathbf{r}^{\top} \mathbf{R}^{-1} \mathbf{r}\right) \tag{2.12}
\end{equation*}
$$

where $|\mathbf{R}|$ denotes the matrix determinant and we have assumed that the measurement noise is a vector with $M$ dimensions. Equation (2.12) may also be written more compactly as

$$
\begin{equation*}
p(\mathbf{r})=\mathcal{N}(\mathbf{r} ; 0, \mathbf{R}) \tag{2.13}
\end{equation*}
$$

More generally, $\mathcal{N}(\mathbf{z} ; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the multivariate pdf of an $M$-dimensional Gaussian random variable $\mathbf{z}$ with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ defined as

$$
\begin{equation*}
\mathcal{N}(\mathbf{z} ; \boldsymbol{\mu}, \boldsymbol{\Sigma})=\frac{1}{(2 \pi)^{M / 2}|\boldsymbol{\Sigma}|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu})^{\top} \boldsymbol{\Sigma}^{-1}(\mathbf{z}-\boldsymbol{\mu})\right) \tag{2.14}
\end{equation*}
$$

Assuming the measurement noise to be Gaussian is indeed reasonable in many (but not all) applications. Furthermore, it also has some important implications on the algorithms that we will derive, in particular with respect to the estimators' statistical properties. However, we will not focus on these aspects throughout most of this course.

### 2.3 Least Squares Methods

The estimation algorithm or inference algorithm combines all the available measurements by using the measurement models to estimate the parameters (or states) in a way that is optimal with respect to some criterion - which is defined by a cost function. The cost function is here selected to be the least squares cost function. By combining the information from multiple measurements and sensors, we can exploit the diversity of the measurements to obtain a better parameter estimate. Hence, sensor fusion algorithms can essentially be seen as an application of estimation theory, which provides a statistically sound and flexible framework.

### 2.3.1 Minimization of Cost Functions

In this course (with the exception of the particle filtering method introduced in Section 6.6), we sacrifice the generality and rigor of estimation theory for simplicity and only consider algorithms that minimize a cost function $J(\mathbf{x})$. These types of algorithms are a good starting point for many sensor fusion applications and have historically proved useful. They also have sound statistical interpretations. However, we will not discuss these explicitly, but for the interested student, there are a couple of exercises that delve into this topic.

Mathematically, minimizing a cost function $J(\mathbf{x})$ to find an estimate of the quantity of interest $\hat{\mathbf{x}}$ (the hat indicates that $\hat{\mathbf{x}}$ is an estimate of the unknown quantity $\mathbf{x}$ ), can be written as the optimization problem

$$
\begin{equation*}
\hat{\mathbf{x}}=\arg \min _{\mathbf{x}} J(\mathbf{x}) \tag{2.15}
\end{equation*}
$$

which reads as "the estimate $\hat{\mathbf{x}}$ is the argument $\mathbf{x}$ that minimizes $J(\mathbf{x})$ ". Typical cost functions minimize a function of the error

$$
\begin{equation*}
e_{n}=y_{n}-g_{n}(\mathbf{x}) \tag{2.16}
\end{equation*}
$$

between the measurement and the output predicted by the estimate. Two typical functions are the absolute error (Figure 2.1a)

$$
\begin{equation*}
\left|e_{n}\right|=\left|y_{n}-g_{n}(\mathbf{x})\right| \tag{2.17}
\end{equation*}
$$

which penalizes all errors equally, and the quadratic cost function (Figure 2.1b)

$$
\begin{equation*}
e_{n}^{2}=\left(y_{n}-g_{n}(\mathbf{x})\right)^{2} \tag{2.18}
\end{equation*}
$$

which favors smaller errors over larger ones.

### 2.3.2 Least Squares

In practice, the quadratic cost is much more common. It implicitly imposes a certain smoothness onto the problem by penalizing large errors more than small


Figure 2.1. Examples of cost functions: (a) Absolute error, and (b) quadratic error.
ones. Additionally, it is closely related to the case where the measurement noise is assumed to be Gaussian as discussed in Section 2.2. Minimizing the quadratic cost function is also referred to as the least squares method. As we will see, many of the most common estimation algorithms can be formulated based on this approach.

Typically, the aim of the estimation algorithm must be to minimize the error over all measurements $y_{1: N}$ rather than only one measurement $y_{n}$. Hence, the resulting cost function is rather the sum over all the errors $e_{n}$. For the quadratic cost (2.18), the cost function is thus

$$
\begin{align*}
J_{\mathrm{LS}}(\mathbf{x}) & =\sum_{n=1}^{N} e_{n}^{2} \\
& =\sum_{n=1}^{N}\left(y_{n}-g_{n}(\mathbf{x})\right)^{2} \tag{2.19}
\end{align*}
$$

Similarly, for vector measurements, the quadratic error of a single measurement can be written as

$$
\begin{equation*}
e_{n}^{2}=\left(\mathbf{y}_{n}-\mathbf{g}_{n}(\mathbf{x})\right)^{\top}\left(\mathbf{y}_{n}-\mathbf{g}_{n}(\mathbf{x})\right) \tag{2.20}
\end{equation*}
$$

and thus, the cost function becomes

$$
\begin{equation*}
J_{\mathrm{LS}}(\mathbf{x})=\sum_{n=1}^{M}\left(\mathbf{y}_{n}-\mathbf{g}_{n}(\mathbf{x})\right)^{\top}\left(\mathbf{y}_{n}-\mathbf{g}_{n}(\mathbf{x})\right) \tag{2.21}
\end{equation*}
$$

Finally, we can write (2.19) and (2.21) more compactly by using the vector notation introduced in Section 2.2. This yields the compact cost function

$$
\begin{equation*}
J_{\mathrm{LS}}(\mathbf{x})=(\mathbf{y}-\mathbf{g}(\mathbf{x}))^{\top}(\mathbf{y}-\mathbf{g}(\mathbf{x})) \tag{2.22}
\end{equation*}
$$

A method that determines the estimate $\hat{\mathbf{x}}$ by minimizing the quadratic cost function in Equation (2.22) is called a least squares ( $L S$ ) method. The formulations (2.19) or (2.21) and (2.22) are equivalent. Hence, we will use these formulations interchangeably, depending on the context.

### 2.3.3 Weighted Least Squares

The least squares method assumes that all measurements are equally reliable. In other words, each measurement is weighed into the estimate equally. However, not every measurement might actually carry the same amount of information: Some measurements might be more important (in some sense) than others and thus, should contribute more to the solution than others.

This can be achieved by using the weighted least squares ( $W L S$ ) cost function instead. By introducing the positive weights $w_{n}$ for each term in (2.19), the weighted least squares cost function for the scalar case becomes

$$
\begin{equation*}
J_{\mathrm{WLS}}(\mathbf{x})=\sum_{n=1}^{N} w_{n}\left(y_{n}-g_{n}(\mathbf{x})\right)^{2} . \tag{2.23}
\end{equation*}
$$

Similarly, for the vector case (2.21), we can introduce a positive definite ${ }^{2}$ weighing $\mathbf{W}_{n}$ matrix and weight each term using this matrix. The resulting cost function becomes

$$
\begin{equation*}
J_{\mathrm{WLS}}(\mathbf{x})=\sum_{n=1}^{N}\left(\mathbf{y}_{n}-\mathbf{g}_{n}(\mathbf{x})\right)^{\top} \mathbf{W}_{n}\left(\mathbf{y}_{n}-\mathbf{g}_{n}(\mathbf{x})\right) \tag{2.24}
\end{equation*}
$$

Again, (2.23)-(2.24) can be written using the compact notation as

$$
\begin{equation*}
J_{\mathrm{WLS}}(\mathbf{x})=(\mathbf{y}-\mathbf{g}(\mathbf{x}))^{\top} \mathbf{W}(\mathbf{y}-\mathbf{g}(\mathbf{x})), \tag{2.25}
\end{equation*}
$$

where the overall weighing matrix is given by either

$$
\mathbf{W}=\left[\begin{array}{cccc}
w_{1} & 0 & \ldots & 0 \\
0 & w_{2} & & \vdots \\
\vdots & & \ddots & 0 \\
0 & \ldots & 0 & w_{N}
\end{array}\right] \text { or } \mathbf{W}=\left[\begin{array}{cccc}
\mathbf{W}_{1} & 0 & \ldots & 0 \\
0 & \mathbf{W}_{2} & & \vdots \\
\vdots & & \ddots & 0 \\
0 & \ldots & 0 & \mathbf{W}_{N}
\end{array}\right]
$$

for the scalar or vector case, respectively. The forms (2.24)-(2.25) are indeed very general and they will be the basis for many of the estimation algorithms derived in this course.

An important question is the choice of the weights $w_{n}$ (or $\mathbf{W}_{n}$ ). This can, in principle, be arbitrary as long as they are positive (or positive definite). In practice, however, a principled choice is to use

$$
w_{n}=1 / \sigma_{r, n}^{2} \text { and } \mathbf{W}_{n}=\mathbf{R}_{n}^{-1}
$$

that is, use the inverse (co)variance of the measurement noise as the weighting factor. The rationale behind this choice is as follows: The covariance is a measure

[^1]for the amount of uncertainty in our measurement due to the measurement noise. A large covariance means that there is a large uncertainty and vice-versa. Hence, by weighing the measurements using the inverse of the covariance, measurements with high uncertainty are given lower weights and measurements with low uncertainty are given higher weights.

### 2.3.4 Regularized Least Squares

In plain (weighted) least squares we find an estimate that best explains the measurements. This amounts to selection of the estimate that minimizes the cost function. However, sometimes it is also useful to regularize the estimate itself by, for example, forcing it to be "small" or in some predefined part of space. This is advantageous in particular when the estimate is not unique - in that case a regularization term can be used to select the estimate that correspond to the area where we believe that the estimate should be. This kind of "prior belief" is called "prior information" and indeed regularization terms correspond to prior distributions of the unknown quantity in estimation theory.

The general form of regularized least squares (ReLS) that we use is

$$
\begin{equation*}
J_{\operatorname{ReLS}}(\mathbf{x})=(\mathbf{y}-\mathbf{g}(\mathbf{x}))^{\top} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{g}(\mathbf{x}))+(\mathbf{x}-\mathbf{m})^{\top} \mathbf{P}^{-1}(\mathbf{x}-\mathbf{m}), \tag{2.26}
\end{equation*}
$$

which has a quadratic regularization term defined by parameters $\mathbf{m}$ and $\mathbf{P}$. We have on purpose selected the weighting term to correspond to the measurement noise covariance $\mathbf{W}=\mathbf{R}^{-1}$, because in that case the regularization term has the interpretation of corresponding to a Gaussian distribution with mean $\mathbf{m}$ and covariance $\mathbf{P}$. What it says is that our prior belief is that the estimate should have a Gaussian distribution with mean $\mathbf{m}$ and covariance $\mathbf{P}$. With covariance going to infinity, our prior belief vanishes and the problem becomes an ordinary (weighted) least squares problem.

20 CHAPTER 2. SENSORS, MODELS, AND LEAST SQUARES CRITERION

## Chapter 3

## Static Linear Models

In this chapter, we focus on static, linear models. These models are quite versatile already and have several important properties, one being that we can find a closed form estimation algorithm.

### 3.1 Linear Model

### 3.1.1 Scalar Model

The simplest measurement model arises when the measurement is a scaled and noisy version of a scalar quantity of interest. In this case, the scalar measurement $y_{n}$ can be written as

$$
\begin{equation*}
y_{n}=g x+r_{n} \tag{3.1}
\end{equation*}
$$

where we assume that the scale factor $g$ is known.
Example 3.1. To measure the wall-distance of the drone in Figure 1.2 we could use, for example, radar which essentially measures the time difference between the sent signal and the signal reflected from the wall. This time difference is ideally equal to the double distance divided by the speed of light $\tau=2 p^{x} / c$, where $c=299792458 \mathrm{~m} / \mathrm{s}$. If we do this measurent $N$ times, and each measurement is corrupted with noise, then the measurement model is

$$
\begin{equation*}
y_{n}=\frac{2}{c} p^{x}+r_{n}, \quad n=1, \ldots, N \tag{3.2}
\end{equation*}
$$

which has the form (3.1) with $g=2 / c$.
With (3.1) in mind, we can construct similar models where the scalar measurement linearly depends on a set of $K$ unknown quantities $x_{1}, \ldots, x_{K}$ :

$$
\begin{equation*}
y_{n}=g_{1} x_{1}+g_{2} x_{2}+\cdots+g_{K} x_{K}+r_{n} \tag{3.3}
\end{equation*}
$$

with known scale factors $g_{1}, \ldots, g_{K}$ and noise term $r_{n}$. We can rewrite (3.3) in matrix form as

$$
y_{n}=\left[\begin{array}{llll}
g_{1} & g_{2} & \ldots & g_{K}
\end{array}\right]\left[\begin{array}{c}
x_{1}  \tag{3.4}\\
x_{2} \\
\vdots \\
x_{K}
\end{array}\right]+r_{n}=\mathbf{g x}+r_{n}
$$

where $\mathbf{g}=\left[\begin{array}{llll}g_{1} & g_{2} & \ldots & g_{K}\end{array}\right]$ is the vector of constants and $\mathbf{x}=\left[\begin{array}{llll}x_{1} & x_{2} & \ldots & x_{K}\end{array}\right]^{\top}$ is the vector of unknowns. For further generality, we can replace $\mathbf{g}$ above with $\mathbf{g}_{n}$ which would allow for different set of coefficients for each measurement, which yields

$$
\begin{equation*}
y_{n}=\mathbf{g}_{n} \mathbf{x}+r_{n} \tag{3.5}
\end{equation*}
$$

For a total of $N$ measurements $y_{1}, y_{2}, \ldots, y_{N}$, we can use measurement stacking as discussed in Section 2.2 to obtain the compact model

$$
\begin{equation*}
\mathbf{y}=\mathbf{G x}+\mathbf{r} \tag{3.6}
\end{equation*}
$$

where $\mathbf{y}$ and $\mathbf{r}$ are the stacked measurements and measurement noise, respectively. Moreover,

$$
\mathbf{G}=\left[\begin{array}{c}
\mathbf{g}_{1}  \tag{3.7}\\
\mathbf{g}_{2} \\
\vdots \\
\mathbf{g}_{N}
\end{array}\right]
$$

is a matrix that contains the known factors $\mathbf{g}_{1}, \ldots, \mathbf{g}_{N}$.

### 3.1.2 Vector Model

The scalar model can also be extended to the case where each measurement itself is a vector. In this case, a single vector-valued measurement $\mathbf{y}_{n}$ can be written as

$$
\mathbf{y}_{n}=\left[\begin{array}{cccc}
g_{11} & g_{12} & \ldots & g_{1 K}  \tag{3.8}\\
g_{21} & g_{22} & \ldots & g_{2 K} \\
\vdots & \vdots & \ddots & \vdots \\
g_{d_{y} 1} & g_{d_{y} 2} & \cdots & g_{d_{y} K}
\end{array}\right]\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{K}
\end{array}\right]+\mathbf{r}_{n}
$$

We can now collect the terms $g_{i j}$ into a matrix and rewrite the above model in form

$$
\begin{equation*}
\mathbf{y}_{n}=\mathbf{G}_{n} \mathbf{x}+\mathbf{r}_{n} \tag{3.9}
\end{equation*}
$$

where $\mathbf{G}_{n}$ contains the collected terms, and similarly to previous section, we have already allowed it to vary for each measurement.

Again, by stacking the measurements $\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{N}$ into a single vector $\mathbf{y}$, we obtain the compact model

$$
\begin{equation*}
\mathbf{y}=\mathbf{G} \mathbf{x}+\mathbf{r} \tag{3.10}
\end{equation*}
$$

where the matrix $\mathbf{G}$ is given by

$$
\mathbf{G}=\left[\begin{array}{c}
\mathbf{G}_{1}  \tag{3.11}\\
\mathbf{G}_{2} \\
\vdots \\
\mathbf{G}_{N}
\end{array}\right] .
$$

Comparing (3.6) and (3.10), it can be seen that both the scalar and the vector model can be written in the same form, the only difference being the matrix $\mathbf{G}$. Furthermore, that common model is completely linear in all the unknowns x and thus, this model is called the linear model.

### 3.1.3 Affine Models

An important class of models is a slight extension of a linear model, called affine model, where the model includes an additional known bias $\mathbf{b}$ :

$$
\begin{equation*}
\mathbf{y}=\mathbf{G} \mathbf{x}+\mathbf{b}+\mathbf{r} \tag{3.12}
\end{equation*}
$$

Fortunately, it turns out that the same algorithms that are used for linear models (without a bias) can also be directly used for model with a bias. This is because we can now rewrite the above model as

$$
\begin{equation*}
\mathbf{y}-\mathbf{b}=\mathbf{G} \mathbf{x}+\mathbf{r} \tag{3.13}
\end{equation*}
$$

and consider $\tilde{\mathbf{y}}=\mathbf{y}-\mathbf{b}$ as the measurement is a linear model. That is, we start our estimation procedure by computing $\tilde{\mathbf{y}}$ by subtracting the biases from the measurements and then apply linear model algorithm to the modified model

$$
\begin{equation*}
\tilde{\mathbf{y}}=\mathbf{G x}+\mathbf{r} \tag{3.14}
\end{equation*}
$$

Example 3.2. The drone model in Equations (1.3) and (1.4) has this kind of form. Thus, we can convert the model into linear model by defining de-biased measurements as

$$
\begin{align*}
& \tilde{y}_{1}=y_{1} \\
& \tilde{y}_{2}=y_{2} \\
& \tilde{y}_{3}=y_{3}+\frac{1}{\sqrt{2}} x_{0} \tag{3.15}
\end{align*}
$$

which reduces (1.4) to

$$
\underbrace{\left[\begin{array}{l}
\tilde{y}_{1}  \tag{3.16}\\
\tilde{y}_{2} \\
\tilde{y}_{3}
\end{array}\right]}_{\tilde{\mathbf{y}}}=\underbrace{\left[\begin{array}{cc}
1 & 0 \\
0 & 1 \\
\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}}
\end{array}\right]}_{\mathbf{G}} \underbrace{\left[\begin{array}{c}
p^{x} \\
p^{y}
\end{array}\right]}_{\mathbf{x}}+\underbrace{\left[\begin{array}{c}
r_{1} \\
r_{2} \\
r_{3}
\end{array}\right]}_{\mathbf{r}},
$$

which is indeed a linear model.

Example 3.3. For the autonomous car model in Equations (1.6) and (1.7) we can eliminate the bias by defining

$$
\begin{gather*}
\tilde{y}_{1}=y_{1}-s_{1}^{x} \\
\tilde{y}_{2}=y_{2}-s_{1}^{y} \\
\vdots \\
\tilde{y}_{2 M-1}=y_{2 M-1}-s_{M}^{x},  \tag{3.17}\\
\tilde{y}_{2 M}
\end{gather*}=y_{2 M}-s_{M}^{y} .
$$

### 3.2 Linear Least Squares

The linear model introduced in the previous section in (3.10) can be seen to be a linear equation system. Unfortunately, in addition to the $K$ unknown parameters of interest in $\mathbf{x}$, there are also a $N d_{y}$ unknown noise values, which yields a total of $K+N d_{y}$ unknowns. Since we only have $N d_{y}$ measurements ( $\approx$ equations) the equation system is under-determined and can not be solved. Even adding more measurements does not solve this problem since every new measurement adds a new unknown (the noise values).

Fortunately, we have already introduced an alternative approach: Minimizing the error cost as discussed in Section 2.3. In this section, we will pursue the least squares approach for the linear model introduced in the previous section.

### 3.2.1 Scalar Model

We start our discussion based on the scalar model (3.1). For this model, the error for each measurement is given by

$$
\begin{equation*}
e_{n}=y_{n}-g x \tag{3.18}
\end{equation*}
$$

Hence, the least squares cost function becomes

$$
\begin{equation*}
J_{\mathrm{LS}}(x)=\sum_{n=1}^{N}\left(y_{n}-g x\right)^{2} \tag{3.19}
\end{equation*}
$$

The least squares estimate $\hat{x}_{\text {LS }}$ is the value of $x$ that minimizes the cost function (3.19). The minima of the cost function with respect to $x$ are found by setting the derivative of $J_{\mathrm{LS}}(x)$ (with respect to $x$ ) to zero and solving it for $x$.

The derivative is given by

$$
\begin{align*}
\frac{\partial J_{\mathrm{LS}}(x)}{\partial x} & =\frac{\partial}{\partial x} \sum_{n=1}^{N}\left(y_{n}-g x\right)^{2} \\
& =\sum_{n=1}-2 g\left(y_{n}-g x\right) \\
& =-\sum_{n=1} 2 g y_{n}+\sum_{n=1}^{N} 2 g^{2} x \\
& =-2 g \sum_{n=1} y_{n}+2 N g^{2} x . \tag{3.20}
\end{align*}
$$

Setting (3.20) to zero yields

$$
0=-2 g \sum_{n=1} y_{n}+2 N g^{2} x,
$$

and solving for $x$ finally yields the estimator

$$
\begin{equation*}
\hat{x}_{\mathrm{LS}}=\frac{1}{N g} \sum_{n=1}^{N} y_{n} . \tag{3.21}
\end{equation*}
$$

Two important aspects of an estimator are its mean and variance. The mean indicates whether an estimator (on average) converges to the true parameter value $x$ and the covariance indicates how confident we are about the estimated value. For the scalar least squares estimator in (3.21), the mean is given by

$$
\begin{aligned}
\mathrm{E}\left\{\hat{x}_{\mathrm{LS}}\right\} & =\mathrm{E}\left\{\frac{1}{N g} \sum_{n=1}^{N} y_{n}\right\} \\
& =\frac{1}{N g} \sum_{n=1}^{N} \mathrm{E}\left\{g x+r_{n}\right\} \\
& =\frac{1}{N g} \sum_{n=1}^{N} g x+\mathrm{E}\left\{r_{n}\right\} \\
& =\frac{1}{N g} N g x+\sum_{n=1}^{N} \mathrm{E}\left\{r_{n}\right\} \\
& =x+\sum_{n=1}^{N} \mathrm{E}\left\{r_{n}\right\} .
\end{aligned}
$$

Hence, if (and only if) the measurement noise is zero-mean, that is, $\mathrm{E}\left\{r_{n}\right\}=0$, the mean reduces to

$$
\begin{equation*}
\mathrm{E}\left\{\hat{x}_{\mathrm{LS}}\right\}=x . \tag{3.22}
\end{equation*}
$$

In this case, $\hat{x}$ converges to the true parameter value and the estimator is said to be an unbiased estimator. Furthermore, for the zero-mean noise case, the variance is given by

$$
\begin{aligned}
\operatorname{var}\left\{\hat{x}_{\mathrm{LS}}\right\} & =\mathrm{E}\left\{\left(\hat{x}_{\mathrm{LS}}-\mathrm{E}\left\{\hat{x}_{\mathrm{LS}}\right\}\right)^{2}\right\} \\
& =\mathrm{E}\left\{\left(\frac{1}{N g} \sum_{n=1}^{N} y_{n}-x\right)^{2}\right\} \\
& =\mathrm{E}\left\{\left(\frac{1}{N g} \sum_{n=1}^{N}\left(g x+r_{n}\right)-x\right)^{2}\right\} \\
& =\mathrm{E}\left\{\left(\frac{1}{N g}\left(N g x+\sum_{n=1}^{N} r_{n}\right)-x\right)^{2}\right\} \\
& =\mathrm{E}\left\{\left(\frac{1}{N g} \sum_{n=1}^{N} r_{n}\right)^{2}\right\} \\
& =\frac{1}{N^{2} g^{2}} \mathrm{E}\left\{\left(\sum_{n=1}^{N} r_{n}\right)^{2}\right\}
\end{aligned}
$$

and finally assuming that $\operatorname{var}\left\{r_{n}\right\}=\sigma^{2}$ we get

$$
\begin{equation*}
\operatorname{var}\left\{\hat{x}_{\mathrm{LS}}\right\}=\frac{\sigma^{2}}{N g^{2}} \tag{3.23}
\end{equation*}
$$

where we have also assumed that the individual $r_{n} \mathrm{~s}$ are independent.
The error in the estimator is measured by the standard deviation which is the square root of the variance

$$
\begin{equation*}
\operatorname{std}\left\{\hat{x}_{\mathrm{LS}}\right\}=\frac{1}{\sqrt{N}} \frac{\sigma}{|g|} \tag{3.24}
\end{equation*}
$$

which shows that the standard error of the estimate diminishes as $1 / \sqrt{N}$ with the number of measurements. Given the estimator and its standard deviation we can also compute the confidence interval of (loosely speaking) the true parameter value. The typically used confidence interval is the $95 \%$ interval

$$
\begin{equation*}
\left[\hat{x}_{\mathrm{LS}}-1.96 \operatorname{std}\left\{\hat{x}_{\mathrm{LS}}\right\}, \hat{x}_{\mathrm{LS}}+1.96 \operatorname{std}\left\{\hat{x}_{\mathrm{LS}}\right\}\right], \tag{3.25}
\end{equation*}
$$

which under certain conditions (Gaussianity, etc.) contains the true parameter value with a probability of $95 \%$. For visualization purposes the coefficient 1.96 is often replaced with constant 2 as it is an approximation anyway. Furthermore, in multivariate case the formula for confidence interval is a bit more complicated, but using the multiplier 2 still provides a reasonable approximation to the confidence interval.

Example 3.4. Let us consider the wall-distance measurement model in Example 3.1 and assume that we estimate $p^{x}$ by using (3.21) with $N$ measurements:

$$
\begin{equation*}
\hat{x}_{1}=\frac{c}{2 N} \sum_{n=1}^{N} y_{n} \tag{3.26}
\end{equation*}
$$

This estimator is unbiased. Further assume that the standard deviation of the measurement is $\sigma=10^{-9} \mathrm{~s}$ (1 nanosecond). Then by (3.24), the standard deviation of the estimator is

$$
\begin{equation*}
\operatorname{std}\left\{\hat{x}_{1}\right\}=\frac{1}{\sqrt{N}} \frac{c \sigma}{2} \tag{3.27}
\end{equation*}
$$

With a single measurement $(N=1)$ we get the error of 15 cm whereas by averaging 100 measurements the error drops to 1.5 cm .

### 3.2.2 General Linear Model

We can now generalize the result (3.21) for general linear models of the form (3.10). Using the vector form of the least squares criterion in (2.22) together with the model (3.10) yields the cost function

$$
\begin{equation*}
J_{\mathrm{LS}}(\mathbf{x})=(\mathbf{y}-\mathbf{G} \mathbf{x})^{\top}(\mathbf{y}-\mathbf{G} \mathbf{x}) \tag{3.28}
\end{equation*}
$$

To minimize this cost function, we can follow the same steps as in the scalar case. First, we calculate the gradient of (3.28) with respect to the parameters x. Since $\mathbf{x}$ is a vector, this can be done by calculating the partial derivative with respect to each individual entry $x_{k}$ in $\mathbf{x}$ individually. A more compact alternative is to use vector calculus. In the vector calculus notation we have

$$
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial \mathbf{x}}=\left[\begin{array}{c}
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial x_{1}}  \tag{3.29}\\
\frac{\partial \mathrm{JS}_{\mathrm{LS}}(\mathbf{x})}{\partial x_{2}} \\
\vdots \\
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial x_{K}}
\end{array}\right]
$$

which is the gradient of $J_{\mathrm{LS}}(\mathbf{x})$ which in some literature is denoted as " $\nabla J_{\mathrm{LS}}(\mathbf{x})$ ". We now have the following computation rules for the vector derivatives (Petersen and Pedersen, 2012):

$$
\begin{aligned}
\frac{\partial \mathbf{x}^{\top} \mathbf{a}}{\partial \mathbf{x}} & =\frac{\partial \mathbf{a}^{\top} \mathbf{x}}{\partial \mathbf{x}}=\mathbf{a} \\
\frac{\partial \mathbf{x}^{\top} \mathbf{A} \mathbf{x}}{\partial \mathbf{x}} & =2 \mathbf{A} \mathbf{x}
\end{aligned}
$$

which then gives

$$
\begin{aligned}
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial \mathbf{x}} & =\frac{\partial}{\partial \mathbf{x}}(\mathbf{y}-\mathbf{G} \mathbf{x})^{\top}(\mathbf{y}-\mathbf{G} \mathbf{x}) \\
& =\frac{\partial}{\partial \mathbf{x}}\left(\mathbf{y}^{\top} \mathbf{y}-\mathbf{y}^{\top} \mathbf{G} \mathbf{x}-\mathbf{x}^{\top} \mathbf{G}^{\top} \mathbf{y}+\mathbf{x}^{\top} \mathbf{G}^{\top} \mathbf{G} \mathbf{x}\right) \\
& =-2 \mathbf{G}^{\top} \mathbf{y}+2 \mathbf{G}^{\top} \mathbf{G} \mathbf{x}
\end{aligned}
$$

Then, setting the gradient to zero and solving for x yields

$$
\begin{equation*}
\hat{\mathbf{x}}_{\mathrm{LS}}=\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{y} \tag{3.30}
\end{equation*}
$$

for the least squares estimator. Note that (3.30) is valid for both scalar measurements of the form (3.3) as well as (3.9), with the appropriate choice of the matrix G.

Similar to the scalar case, we can calculate the statistical properties of the estimator (3.30). The mean is found from

$$
\begin{aligned}
\mathrm{E}\left\{\hat{\mathbf{x}}_{\mathrm{LS}}\right\} & =\mathrm{E}\left\{\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{y}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top}(\mathbf{G} \mathbf{x}+\mathbf{r})\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{G} \mathbf{x}+\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{r}\right\} \\
& =\mathbf{x}+\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathrm{E}\{\mathbf{r}\},
\end{aligned}
$$

and we can see that if (and only if) $\mathrm{E}\{\mathbf{r}\}=0$, we have that

$$
\begin{equation*}
\mathrm{E}\left\{\hat{\mathrm{x}}_{\mathrm{LS}}\right\}=\mathbf{x} \tag{3.31}
\end{equation*}
$$

and the estimator is unbiased. Furthermore, for the zero-mean noise case, the variance can be derived from

$$
\begin{aligned}
\operatorname{Cov}\left\{\hat{\mathbf{x}}_{\mathrm{LS}}\right\} & =\mathrm{E}\left\{\left(\hat{\mathbf{x}}_{\mathrm{LS}}-\mathrm{E}\left\{\hat{\mathbf{x}}_{\mathrm{LS}}\right\}\right)\left(\hat{\mathbf{x}}_{\mathrm{LS}}-\mathrm{E}\left\{\hat{\mathbf{x}}_{\mathrm{LS}}\right\}\right)^{\top}\right\} \\
& =\mathrm{E}\left\{\left(\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top}(\mathbf{G} \mathbf{x}+\mathbf{r})-\mathbf{x}\right)\left(\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top}(\mathbf{G} \mathbf{x}+\mathbf{r})-\mathbf{x}\right)^{\top}\right\} \\
& =\mathrm{E}\left\{\left(\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{r}\right)\left(\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{r}\right)^{\top}\right\} \\
& =\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathrm{E}\left\{\mathbf{r r}^{\top}\right\} \mathbf{G}\left(\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1}\right)^{\top},
\end{aligned}
$$

and is given by

$$
\begin{equation*}
\operatorname{Cov}\left\{\hat{\mathbf{x}}_{\mathrm{LS}}\right\}=\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{R G}\left(\left(\mathbf{G}^{\top} \mathbf{G}\right)^{-1}\right)^{\top} . \tag{3.32}
\end{equation*}
$$

Example 3.5. Let us recall the autonomous car shown in Figure 1.4 and its model in Equations (1.6) and (1.7). The positions of the landmarks are

$$
\begin{array}{cccc}
s_{1}^{x}=2, & s_{2}^{x}=0, & s_{3}^{x}=10, & s_{4}^{x}=7,  \tag{3.33}\\
s_{1}^{y}=6, & s_{2}^{y}=0, & s_{3}^{y}=2, & s_{4}^{y}=8 .
\end{array}
$$



Figure 3.1. Least squares estimation example.

The joint measurement noise covariance matrix is a diagonal matrix

$$
\begin{equation*}
\mathbf{R}=\operatorname{diag}\left(R_{1}^{x}, R_{1}^{y}, R_{2}^{x}, R_{2}^{y}, R_{3}^{x}, R_{3}^{y}, R_{4}^{x}, R_{4}^{y}\right), \tag{3.34}
\end{equation*}
$$

where the individial variances are given as

$$
\begin{array}{llll}
R_{1}^{x}=1, & R_{2}^{x}=0.1, & R_{3}^{x}=0.8, & R_{4}^{x}=0.5 \\
R_{1}^{y}=0.3, & R_{2}^{y}=0.5, & R_{3}^{y}=2, & R_{4}^{y}=1.5 \tag{3.35}
\end{array}
$$

Furthermore, in order to use the least squares equations in this section we need to de-bias the measurements using (3.17). Figure 3.2 shows the least squares estimate of the autonomous car position (red dot) and its confidence ellipse (red ellipse) when the relative locations of landmarks (orange) where measured with the standard deviations given above. The true position is shown in blue.

### 3.2.3 Weighted Linear Least Squares

As mentioned in Section 2.3, it is sometimes desirable to use a weighted least squares criterion rather than the least squares criterion to take different levels of certainty into account. In this case, the cost function (2.25) with $\mathbf{W}=\mathbf{R}^{-1}$ is used together with the general linear model (3.10), which yields

$$
\begin{equation*}
J_{\mathrm{WLS}}(\mathbf{x})=(\mathbf{y}-\mathbf{G} \mathbf{x})^{\top} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{G} \mathbf{x}) . \tag{3.36}
\end{equation*}
$$

Following the same steps for the derivation as for the least squares case, we obtain the weighted least squares estimator given by

$$
\begin{equation*}
\hat{\mathbf{x}}_{\mathrm{WLS}}=\left(\mathbf{G}^{\boldsymbol{\top}} \mathbf{R}^{-1} \mathbf{G}\right)^{-1} \mathbf{G}^{\boldsymbol{\top}} \mathbf{R}^{-1} \mathbf{y} . \tag{3.37}
\end{equation*}
$$

Comparing the linear least squares estimator in (3.30) and the weighted version in (3.37), we see that the former is a special case of the latter if and only if the covariance matrix is proportional to the identity matrix, that is, $\mathbf{R}=\sigma_{r}^{2} \mathbf{I}$.


Figure 3.2. Weighted least squares estimation example.

Similarly, the mean and covariance of the weighted least squares estimator can be shown to be

$$
\begin{aligned}
\mathrm{E}\left\{\hat{\mathbf{x}}_{\mathrm{WLS}}\right\} & =\mathbf{x}+\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{R}^{-1} \mathrm{E}\{\mathbf{r}\} \\
& =\mathbf{x}
\end{aligned}
$$

and

$$
\begin{align*}
\operatorname{Cov}\left\{\hat{\mathbf{x}}_{\mathrm{WLS}}\right\} & =\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G}\right)^{-1} \mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{R} \mathbf{R}^{-1} \mathbf{G}\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G}\right)^{-1} \\
& =\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G}\right)^{-1} \tag{3.38}
\end{align*}
$$

respectively.
As discussed in Section 2.3, it is possible to make other choices for the weighing matrix $\mathbf{W}$. However, it can be shown that choosing the inverse noise covariance matrix $\mathbf{R}^{-1}$ is the optimal choice for $\mathbf{W}$ in the sense that it minimizes the trace of the covariance of the estimator.

Example 3.6. Figure 3.2 shows the weighed least squares estimate of the autonomous car position using the same data as in Example 3.5. It can be seen that the confidence interval is now smaller and the estimator is closer to the true position.

### 3.3 Regularized Linear Least Squares

So far, we have assumed that the unknowns $\mathbf{x}$ are completely arbitrary and deterministic. However, as discussed in Section 2.3, in many cases, we have some prior information about the parameters, for example we know beforehand that they might be distributed around some mean $\mathrm{E}\{\mathbf{x}\}=\mathbf{m}$ with covariance $\operatorname{Cov}\{\mathbf{x}\}=\mathbf{P}$. This prior information can be incorporated into the estimator by adding a regularization term to the cost function, which leads to the regularized linear least squares estimator.

The cost function then becomes

$$
\begin{equation*}
J_{\operatorname{ReLS}}(\mathbf{x})=(\mathbf{y}-\mathbf{G} \mathbf{x})^{\top} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{G} \mathbf{x})+(\mathbf{x}-\mathbf{m})^{\top} \mathbf{P}^{-1}(\mathbf{x}-\mathbf{m}) \tag{3.39}
\end{equation*}
$$

where the second term $(\mathbf{x}-\mathbf{m})^{\top} \mathbf{P}^{-1}(\mathbf{x}-\mathbf{m})$ is a regularization term that encodes the prior information. The minimum of this cost function can now be derived by setting the gradient to zero. The gradient is given by

$$
\frac{\partial J_{\operatorname{ReLS}}(\mathbf{x})}{\partial \mathbf{x}}=-2 \mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{y}+2 \mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G} \mathbf{x}-2 \mathbf{P}^{-1} \mathbf{m}+2 \mathbf{P}^{-1} \mathbf{x}
$$

which now has some additional terms resulting from the regularization. Setting it to zero and solving for x then yields the regularized linear least squares estimator

$$
\begin{equation*}
\hat{\mathbf{x}}_{\mathrm{ReLS}}=\left(\mathbf{G}^{\boldsymbol{\top}} \mathbf{R}^{-1} \mathbf{G}+\mathbf{P}^{-1}\right)^{-1}\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{y}+\mathbf{P}^{-1} \mathbf{m}\right) \tag{3.40}
\end{equation*}
$$

The estimator in (3.40) shows that when using the regularization approach, the resulting estimate is a weighted average of the data, which enters through the term $\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{y}$ and the prior information through $\mathbf{P}^{-1} \mathbf{m}$. The contributions of each of the terms are determined by two factors. First, there are the weighing matrices $\mathbf{R}^{-1}$ and $\mathbf{P}^{-1}$, which directly weigh the contributions. Second, the number of data points in $\mathbf{y}$ also affect how much emphasis is put on the data, and how much on the prior information.

The covariance of the estimator can then shown to be

$$
\begin{equation*}
\operatorname{Cov}\left\{\hat{\mathbf{x}}_{\text {ReLS }}\right\}=\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G}+\mathbf{P}^{-1}\right)^{-1} \tag{3.41}
\end{equation*}
$$

It can further be shown that the trace of the above covariance (which is the total variance) is strictly smaller than that of the non-regularized version in Equation (3.38). This is because including prior information decreases the variance of the estimator. However, if we compute the expectation of the estimator (3.40), we find that the estimator is biased - it is unbiased if and only if $\mathbf{P}^{-1}=0$. This bias is due to our inclusion of prior information which forces the estimate slightly towards prior. However, if our prior belief is correct, the estimator is supposed to be biased and the bias actually forces the estimate to be closer to the truth than the measurements alone imply.

It is also possible to express the mean and covariance equations in alternative, but equivalent form, which will be especially useful in sequential least squares that we discuss in the next section. The matrix inversion formula, which is also called Woodbury matrix identity, gives

$$
\begin{equation*}
\left(\mathbf{G}^{\top} \mathbf{R}^{-1} \mathbf{G}+\mathbf{P}^{-1}\right)^{-1}=\mathbf{P}-\mathbf{P} \mathbf{G}^{\top}\left(\mathbf{G} \mathbf{P} \mathbf{G}^{\top}+\mathbf{R}\right)^{-1} \mathbf{G} \mathbf{P} \tag{3.42}
\end{equation*}
$$

If we further substitute this into the estimator equation and simplify, we can express the estimator (3.40) and its covariance (3.41) as

$$
\begin{equation*}
\hat{\mathbf{x}}_{\mathrm{ReLS}}=\mathbf{m}+\mathbf{K}(\mathbf{y}-\mathbf{G m}) \tag{3.43}
\end{equation*}
$$



Figure 3.3. Regularized least squares estimation example.
and

$$
\begin{equation*}
\operatorname{Cov}\left\{\hat{\mathbf{x}}_{\mathrm{ReLS}}\right\}=\mathbf{P}-\mathbf{K}\left(\mathbf{G} \mathbf{P} \mathbf{G}^{\top}+\mathbf{R}\right) \mathbf{K}^{\top}, \tag{3.44}
\end{equation*}
$$

where $\mathbf{K}$ is a gain given by

$$
\begin{equation*}
\mathbf{K}=\mathbf{P G}^{\top}\left(\mathbf{G P G} \mathbf{G}^{\top}+\mathbf{R}\right)^{-1} \tag{3.45}
\end{equation*}
$$

In this form, a second interpretation of the regularized linear least squares estimator is possible. Recall that $\mathbf{m}$ was the prior mean. Then, the term $\mathbf{y}-\mathbf{G m}$ is the error between the output and the output predicted by the prior mean. The final estimate is then the prior mean that is corrected by the error term through the weight $\mathbf{K}$.

Finally, we also point out that the regularized least squares solution can also be derived as non-regularized (weighted) least squares solution where the regularization term is interpreted as additional set of measurements. Because ( $\mathbf{x}-$ $\mathbf{m})^{\top} \mathbf{P}^{-1}(\mathbf{x}-\mathbf{m})=(\mathbf{m}-\mathbf{x})^{\top} \mathbf{P}^{-1}(\mathbf{m}-\mathbf{x})$, we can rewrite the cost function (3.39) as

$$
\begin{align*}
J_{\operatorname{ReLS}}(\mathbf{x}) & =(\mathbf{y}-\mathbf{G} \mathbf{x})^{\top} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{G} \mathbf{x})+(\mathbf{m}-\mathbf{x})^{\top} \mathbf{P}^{-1}(\mathbf{m}-\mathbf{x}) \\
& =\left(\left[\begin{array}{c}
\mathbf{y} \\
\mathbf{m}
\end{array}\right]-\left[\begin{array}{c}
\mathbf{G} \\
\mathbf{I}
\end{array}\right] \mathbf{x}\right)^{\top}\left[\begin{array}{cc}
\mathbf{R}^{-1} & 0 \\
0 & \mathbf{P}^{-1}
\end{array}\right]\left(\left[\begin{array}{c}
\mathbf{y} \\
\mathbf{m}
\end{array}\right]-\left[\begin{array}{c}
\mathbf{G} \\
\mathbf{I}
\end{array}\right] \mathbf{x}\right), \tag{3.46}
\end{align*}
$$

where the last form is a non-regularized (weighted) least squares problem where the measurement contains $\mathbf{y}$ and $\mathbf{m}$, the measurement model matrix contains an additional identity matrix, and the covariance of the measurement is a block matrix containing $\mathbf{R}^{-1}$ and $\mathbf{P}^{-1}$ on the diagonal.

By using Equations (3.37) and (3.38) we can now derive the regularized estimator and its covariance. This form will be used in nonlinear least squares problems as it allows us to derive algorithms for plain weighted least squares problems and then the solutions to regularized least squares problems can be solved with the same algorithms by using the trick above.

Example 3.7. Figure 3.3 illustrates the effect of regularization in least squares estimate of the autonomous car position using the same data as in Examples 3.5
and 3.6. On the left we have prior information where the mean is close to the truth, but the prior variance is quite large. In this case the estimate is slightly shifted towards the true position. On the right, the prior mean is in the wrong place with relatively small variance, which in this case shifts the estimate to the wrong place. If the prior information were correct, the estimate would be even better.

### 3.4 Sequential Linear Least Squares

In many cases, the sensor data arrives sequentially at the estimator. Assume that we have calculated the least squares estimate $\hat{\mathbf{x}}_{n-1}$ according to (3.37) with covariance $\mathbf{P}_{n-1}=\operatorname{Cov}\left\{\hat{\mathbf{x}}_{n-1}\right\}$ as in (3.38) for the dataset $\mathbf{y}_{1: n-1}=\left\{\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{n-1}\right\}$. When a new measurement $\mathbf{y}_{n}$ arrives, the weighted least squares cost function can be written as

$$
\begin{align*}
J_{\mathrm{SLS}}(\mathbf{x}) & =\left(\mathbf{y}_{1: n-1}-\mathbf{G}_{1: n-1} \mathbf{x}\right)^{\top} \mathbf{R}_{1: n-1}^{-1}\left(\mathbf{y}_{1: n-1}-\mathbf{G}_{1: n-1} \mathbf{x}\right) \\
& +\left(\mathbf{y}_{n}-\mathbf{G}_{n} \mathbf{x}\right)^{\top} \mathbf{R}_{n}^{-1}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \mathbf{x}\right), \tag{3.47}
\end{align*}
$$

where the first part is the cost function that was minimized when $n-1$ measurements were available, and the second part is the additional cost for the $n$th sample.

Minimizing (3.47) is achieved in the same way as for the general linear model in the previous section, except for that there is an additional term. The gradient is given by

$$
\begin{align*}
\frac{\partial J_{\mathrm{SLS}}(\mathbf{x})}{\partial \mathbf{x}} & =-2 \mathbf{G}_{1: n-1}^{\top} \mathbf{R}_{1: n-1}^{-1} \mathbf{y}_{1: n-1}+2 \mathbf{G}_{1: n-1}^{\top} \mathbf{R}_{1: n-1}^{-1} \mathbf{G}_{1: n-1} \mathbf{x} \\
& -2 \mathbf{G}_{n}^{\top} \mathbf{R}_{n}^{-1} \mathbf{y}_{n}+2 \mathbf{G}_{n}^{\top} \mathbf{R}_{n}^{-1} \mathbf{G}_{n} \mathbf{x} \tag{3.48}
\end{align*}
$$

Setting the gradient to zero and solving for $\mathbf{x}$ then yields the updated estimate

$$
\begin{aligned}
& \hat{\mathbf{x}}_{n}=\left(\mathbf{G}_{1: n-1}^{\top} \mathbf{R}_{1: n-1}^{-1} \mathbf{G}_{1: n-1}+\mathbf{G}_{n}^{\top} \mathbf{R}_{n}^{-1} \mathbf{G}_{n}\right)^{-1} \\
& \times\left(\mathbf{G}_{1: n-1}^{\top} \mathbf{R}_{1: n-1}^{-1} \mathbf{y}_{1: n-1}+\mathbf{G}_{n}^{\top} \mathbf{R}_{n}^{-1} \mathbf{y}_{n}\right) \\
&=\left(\mathbf{P}_{n-1}^{-1}+\mathbf{G}_{n}^{\top} \mathbf{R}_{n}^{-1} \mathbf{G}_{n}\right)^{-1}\left(\mathbf{G}_{1: n-1}^{\top} \mathbf{R}_{1: n-1}^{-1} \mathbf{y}_{1: n-1}+\mathbf{G}_{n}^{\top} \mathbf{R}_{n}^{-1} \mathbf{y}_{n}\right) .
\end{aligned}
$$

Using the matrix inversion formula as in the previous section and some further simplifications, we can rewrite this as

$$
\begin{equation*}
\hat{\mathbf{x}}_{n}=\hat{\mathbf{x}}_{n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right) \tag{3.49}
\end{equation*}
$$

with the gain

$$
\begin{equation*}
\mathbf{K}_{n}=\mathbf{P}_{n-1} \mathbf{G}_{n}^{\top}\left(\mathbf{G}_{n} \mathbf{P}_{n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right)^{-1} \tag{3.50}
\end{equation*}
$$

Knowing that $\mathrm{E}\left\{\hat{\mathbf{x}}_{n-1}\right\}=\mathbf{x}$ (which follows from Section 3.2), the mean of the updated estimate is given by

$$
\begin{align*}
\mathrm{E}\left\{\hat{\mathbf{x}}_{n}\right\} & =\mathrm{E}\left\{\hat{\mathbf{x}}_{n-1}\right\}+\mathbf{K}_{n}\left(\mathrm{E}\left\{\mathbf{y}_{n}\right\}-\mathbf{G}_{n} \mathrm{E}\left\{\hat{\mathbf{x}}_{n-1}\right\}\right)  \tag{3.51}\\
& =\mathbf{x}
\end{align*}
$$



Figure 3.4. Sequential least squares estimation example.
that is, the updated estimate is unbiased provided that the previous one was. Furthermore, the covariance of the updated estimate is

$$
\begin{align*}
& \operatorname{Cov}\left\{\hat{\mathbf{x}}_{n}\right\} \\
& =\mathrm{E}\left\{\left(\hat{\mathbf{x}}_{n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right)-\mathbf{x}\right)\left(\hat{\mathbf{x}}_{n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right)-\mathbf{x}\right)^{\top}\right\} \\
& =\mathrm{E}\left\{\left(\hat{\mathbf{x}}_{n-1}-\mathbf{x}\right)\left(\hat{\mathbf{x}}_{n-1}-\mathbf{x}\right)^{\top}\right\}+\mathrm{E}\left\{\left(\hat{\mathbf{x}}_{n-1}-\mathbf{x}\right)\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right)^{\top} \mathbf{K}_{n}^{\top}\right\} \\
& \quad+\mathrm{E}\left\{\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right)\left(\hat{\mathbf{x}}_{n-1}-\mathbf{x}\right)^{\top}\right\} \\
& \quad+\mathrm{E}\left\{\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right)\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n-1}\right)^{\top} \mathbf{K}_{n}^{\top}\right\} \\
& =\mathbf{P}_{n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{n} \mathbf{P}_{n-1} \mathbf{G}^{\top}+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\top} . \tag{3.52}
\end{align*}
$$

Note how the covariance update in (3.52) illustrates how adding new measurements reduces the uncertainty of the estimate.

It turns out that regularization can be easily incorporated into sequential estimation. All we need to do is to set $\hat{\mathbf{x}}_{0}=\mathbf{m}$ and $\mathbf{P}_{0}=\mathbf{P}$ from the regularization term in (3.39) and then proceed with the above update equations.

Example 3.8. Figure 3.4 illustrates the sequential least squares estimation of the autonomous car position using the same data as in Examples 3.5, 3.6, and 3.7. On the left top we have used only one measurement, on the right top two measurements, on the bottow left three, and in the bottom right all the four measurements. It
can be seen how each measurement gives more information and after all four measurements the result exactly coincides with the weighted least squares estimate.

## Chapter 4

## Static Nonlinear Models

In Chapter 3, we discussed problems where the sensor model is linear in the parameters of interest. For these model types, closed-form estimators exist and furthermore, we can analytically determine the estimators' properties such as biasedness or covariance. However, despite these desirable characteristics, linear models can be limiting and a nonlinear model may be more accurate in describing the relationship between the sensors' measurements and the parameters of interest.

### 4.1 Nonlinear Model

In this section, we develop estimation algorithms for general models of the form

$$
\begin{equation*}
\mathbf{y}=\mathbf{g}(\mathbf{x})+\mathbf{r} \tag{4.1}
\end{equation*}
$$

where $\mathbf{g}(\mathbf{x})$ may be an arbitrary nonlinear (vector-valued) function. As for the linear models, they aim will be to minimize the weighted least squares cost function

$$
\begin{equation*}
J_{\mathrm{WLS}}(\mathbf{x})=(\mathbf{y}-\mathbf{g}(\mathbf{x}))^{\mathrm{T}} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{g}(\mathbf{x})) \tag{4.2}
\end{equation*}
$$

with respect to $\mathbf{x}$.
Although the above cost function does not contain a regularization term, we can recall from Section 3.3 that a regularized linear squares problem can be reformulated as a non-regularized least squares problem. The same trick also works for nonlinear problems and hence the algorithms presented in this section will also work for regularized linear squares problems. We come back to this in Section 4.6.

In some special cases, it is still possible to find analytical expressions at least for the (W)LS estimator and possibly even for its properties. Hence, it is always worth trying to derive it. However, for the vast majority of models, this is not possible and we have to resort to iterative, numerical optimization methods to minimize the cost function (4.2). For this purpose, we discuss the following methods here:

- Gradient descent,


Figure 4.1. Illustration of the cost function for a nonlinear problem with two unknowns $\mathbf{x}=\left[\begin{array}{ll}x_{1} & x_{2}\end{array}\right]^{\top}$. (a) Contours of the cost function, and (b) vector field defined by the negative gradient.

- the Gauss-Newton algorithm, and
- the Levenberg-Marquardt algorithm.

Note that all of these algorithms typically only are able to find local minima, unless there is only a global minimum. Hence, it is important to have good initial guesses of the parameters.

### 4.2 Gradient Descent

The first approach, gradient descent is a method based on the following strategy. Assume that we are given the current parameter estimate $\hat{\mathbf{x}}^{(i)}$ at the algorithm's $i$ th iteration. Then, the gradient of the objective function at $\hat{\mathbf{x}}^{(i)}$ indicates the direction of the function input $\mathbf{x}$ in which the function value increases. Hence, taking steps from $\hat{\mathbf{x}}^{(i)}$ in the direction proportional to the negative gradient, the function value should decrease (see Figure 4.1).

This leads to the update rule given by

$$
\begin{equation*}
\hat{\mathbf{x}}^{(i+1)}=\hat{\mathbf{x}}^{(i)}-\left.\gamma \frac{\partial J_{\mathrm{WLS}}(\mathbf{x})}{\partial \mathbf{x}}\right|_{\mathbf{x}=\hat{\mathbf{x}}^{(i)}} \tag{4.3}
\end{equation*}
$$

where $\partial J_{\mathrm{WLS}}(\mathbf{x}) / \partial \mathbf{x}$ denotes the gradient of $J_{\mathrm{WLS}}(\mathbf{x})$ with respect to $\mathbf{x}$ and $\gamma>0$ is a positive constant defining the step length. For simplicity, we start our derivation assuming scalar measurements $y_{n}$ together with the (non-weighted) least squares cost function and later generalize it to the general model (4.1)-(4.2). In this case, we have that

$$
J_{\mathrm{LS}}(\mathbf{x})=\sum_{n=1}^{N}\left(y_{n}-g_{n}(\mathbf{x})\right)^{2}
$$

and thus

$$
\begin{aligned}
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial \mathbf{x}} & =\frac{\partial}{\partial \mathbf{x}} \sum_{n=1}^{N}\left(y_{n}-g_{n}(\mathbf{x})\right)^{2} \\
& =\sum_{n=1}^{N}-2\left(y_{n}-g_{n}(\mathbf{x})\right) \frac{\partial g_{n}(\mathbf{x})}{\partial \mathbf{x}} .
\end{aligned}
$$

The sum can be rewritten in vector form as

$$
\begin{aligned}
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial \mathbf{x}} & \left.=-2\left[\begin{array}{lllc}
\frac{\partial g_{1}(\mathbf{x})}{\partial \mathbf{x}} & \frac{\partial g_{2}(\mathbf{x})}{\partial \mathbf{x}} & \ldots & \frac{\partial g_{N}(\mathbf{x})}{\partial \mathbf{x}}
\end{array}\right]\left[\begin{array}{c}
y_{1} \\
y_{2} \\
\vdots \\
y_{N}
\end{array}\right]-\left[\begin{array}{c}
g_{1}(\mathbf{x}) \\
g_{2}(\mathbf{x}) \\
\vdots \\
g_{N}(\mathbf{x})
\end{array}\right]\right) \\
& =-2\left[\begin{array}{cccc}
\frac{\partial g_{1}(\mathbf{x})}{\partial x_{1}} & \frac{\partial g_{2}(\mathbf{x})}{\partial x_{1}} & \ldots & \frac{\partial g_{N}(\mathbf{x})}{\partial x_{1}} \\
\frac{\partial g_{1}(\mathbf{x})}{\partial x_{2}} & \frac{\partial g_{2}(\mathbf{x})}{\partial x_{2}} & & \vdots \\
\vdots & & \ddots & \frac{\partial g_{N}(\mathbf{x})}{\partial x_{K-1}} \\
\frac{\partial g_{1}(\mathbf{x})}{\partial x_{K}} & \ldots & \frac{\partial g_{N-1}(\mathbf{x})}{\partial x_{K}} & \frac{\partial g_{N}(\mathbf{x})}{\partial x_{K}}
\end{array}\right](\mathbf{y}-\mathbf{g}(\mathbf{x})),
\end{aligned}
$$

and the matrix of derivatives can be identified as the transpose of the Jacobian matrix given by

$$
\mathbf{G}_{\mathbf{x}}(\mathbf{x})=\left[\begin{array}{cccc}
\frac{\partial g_{1}(\mathbf{x})}{\partial x_{1}} & \frac{\partial g_{1}(\mathbf{x})}{\partial x_{2}} & \cdots & \frac{\partial g_{1}(\mathbf{x})}{\partial x_{K}}  \tag{4.4}\\
\frac{\partial g_{2}(\mathbf{x})}{\partial x_{1}} & \frac{\partial g_{2}(\mathbf{x})}{\partial x_{2}} & & \vdots \\
\vdots & & \ddots & \frac{\partial g_{N-1}(\mathbf{x})}{\partial x_{K}} \\
\frac{\partial g_{N}(\mathbf{x})}{\partial x_{1}} & \cdots & \frac{\partial g_{N}(\mathbf{x})}{\partial x_{K-1}} & \frac{\partial g_{N}(\mathbf{x})}{\partial x_{K}}
\end{array}\right]
$$

which is a matrix containing the gradients of the elements of the vector valued $\mathbf{g}(\mathbf{x})$ at each row. Hence, we have that

$$
\begin{equation*}
\frac{\partial J_{\mathrm{LS}}(\mathbf{x})}{\partial \mathbf{x}}=-2 \mathbf{G}_{\mathbf{x}}^{\mathrm{T}}(\mathbf{x})(\mathbf{y}-\mathbf{g}(\mathbf{x})) \tag{4.5}
\end{equation*}
$$

We can now directly generalize (4.5) to the general cost (4.2), which yields

$$
\begin{aligned}
\frac{\partial J_{\mathrm{WLS}}(\mathbf{x})}{\partial \mathbf{x}} & =\frac{\partial}{\partial \mathbf{x}}(\mathbf{y}-\mathbf{g}(\mathbf{x}))^{\top} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{g}(\mathbf{x})) \\
& =\frac{\partial}{\partial \mathbf{x}}\left[\mathbf{y}^{\top} \mathbf{R}^{-1} \mathbf{y}-\mathbf{y}^{\top} \mathbf{R}^{-1} \mathbf{g}(\mathbf{x})-\mathbf{g}(\mathbf{x})^{\top} \mathbf{R}^{-1} \mathbf{y}+\mathbf{g}(\mathbf{x}) \mathbf{R}^{-1} \mathbf{g}(\mathbf{x})\right] \\
& =-2 \mathbf{G}_{\mathbf{x}}^{\top}(\mathbf{x}) \mathbf{R}^{-1}(\mathbf{y}-\mathbf{g}(\mathbf{x}))
\end{aligned}
$$

Note that the direction of the parameter update is given by $-\mathbf{G}_{\mathbf{x}}^{\mathbf{T}}(\mathbf{x}) \mathbf{R}^{-1}(\mathbf{y}-$ $\mathbf{g}(\mathbf{x}))$ which we will multiply with the step size $\gamma$. We can now absorb the factor 2 into the step size which yields the gradient descent update (4.3) given by

$$
\begin{equation*}
\hat{\mathbf{x}}^{(i+1)}=\hat{\mathbf{x}}^{(i)}+\gamma \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right) . \tag{4.6}
\end{equation*}
$$

```
Algorithm 4.1 Gradient Descent
Input: Initial parameter guess \(\hat{\mathbf{x}}^{(0)}\), data \(\mathbf{y}\), function \(\mathbf{g}(\mathbf{x})\), Jacobian \(\mathbf{G}_{\mathbf{x}}(\mathbf{x})\)
Output: Parameter estimate \(\hat{\mathbf{x}}_{\text {WLS }}\)
    Set \(i \leftarrow 0\)
    repeat
        Calculate the update direction
\[
\Delta \mathbf{x}^{(i+1)}=\mathbf{G}_{\mathbf{x}}^{\mathbf{T}}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)
\]
        Select a suitable \(\gamma^{(i+1)}\)
        Calculate
                        \(\hat{\mathbf{x}}^{(i+1)}=\hat{\mathbf{x}}^{(i)}+\gamma^{(i+1)} \Delta \mathbf{x}^{(i+1)}\)
        Set \(i \leftarrow i+1\)
    until Converged
```




Figure 4.2. Gradient descent optimization example. The left figure shows the evolution of the estimates from start to the end, and right figure is a zoom to the end of the optimization.

It remains to determine how long the step length $\gamma$ should be chosen. Choosing $\gamma$ too large may cause the cost function to increase, whereas too small steps might cause unnecessarily slow convergence. A typical strategy is to simply choose it small enough so that the cost decreases at every step. However, quite often it is advisable to change the step length during the iterations in one way or another. One approach is to use a line search for selecting it, but we will postpone the discussion on line-search methods to Section 4.4, where we discuss it in the context of GaussNewton methods. A general gradient descent algorithm is shown as Algorithm 4.1.

Example 4.1. Let us consider a similar car localization problem as in Example 3.5 which was depicted in Figure 1.4. However, let us assume that we are only able to measure the ranges to the landmarks as specified in the model in Equations (1.8). In this case the model becomes non-linear. We assume that the variances of the range measurements are $R_{1}^{R}=1, R_{2}^{R}=0.1, R_{3}^{R}=0.8$, and $R_{4}^{R}=0.5$. The result of applying the gradient descent algorithm to the corresponding weighted least squares problem is shown in Figure 4.2. It can be seen that gradient descent
indeed finds the minimum, but the route that it finds is far from a direct route and when we get closer to the minimum, the steps become smaller and smaller although we have kept $\gamma$ constant.

While intuitive, the gradient descent approach suffers from an important problem. In areas where the cost function is flat, the gradient is very small. Hence, the algorithm can only take small steps, which leads to the problem that a large number of iterations is needed to cross such areas. Thus in general, the pure gradient descent method is a quite inefficient algorithm and if, for example, Gauss-Newton, Levenberg-Marquardt, or Quasi-Newton methods can be used, they should be preferred. However, recently, gradient descent methods and in particular their stochastic versions, stochastic gradient decent methods, have gained some popularity due to its computational efficiency in models with high number of parameters such as large neural networks.

### 4.3 Gauss-Newton Algorithm

The second method, the Gauss-Newton algorithm, is based on a different approach. Given the estimate $\hat{\mathbf{x}}^{(i)}$ of $\mathbf{x}$ at the $i$ th iteration, the nonlinear function $\mathbf{g}(\mathbf{x})$ can locally be approximated by using a first order Taylor series expansion. This approximation is given by

$$
\begin{equation*}
\mathbf{g}(\mathbf{x}) \approx \mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)+\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right) \tag{4.7}
\end{equation*}
$$

where $\mathbf{G}_{\mathbf{x}}$ is the Jacobian matrix of $g(\mathbf{x})$ given by (4.4). Using this local linear approximation of the nonlinear function, the weighted least squares cost can be approximated by

$$
\begin{align*}
J_{\mathrm{WLS}}(\mathbf{x}) \approx & \left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right)^{\top} \mathbf{R}^{-1} \\
& \times\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right) \\
= & \left(\mathbf{e}^{(i)}-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right)^{\top} \mathbf{R}^{-1}\left(\mathbf{e}^{(i)}-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right) \tag{4.8}
\end{align*}
$$

were we have made use of $\mathbf{e}^{(i)}=\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)$ (which is the error of the $i$ th iteration). The approximated cost function (4.8) is linear in the parameters $\mathbf{x}$ (remember that $\hat{\mathbf{x}}^{(i)}$ is constant and given from the last iteration) and hence, we can derive a closed form solution for the parameter update. This is achieved by setting the approximative cost function's gradient to zero and solving for $\mathbf{x}$ as

$$
\begin{aligned}
\frac{\partial J_{\mathrm{WLS}}(\mathbf{x})}{\partial \mathbf{x}} \approx & \frac{\partial}{\partial \mathbf{x}}\left(\left(\mathbf{e}^{(i)}\right)^{\top} \mathbf{R}^{-1} \mathbf{e}^{(i)}-\left(\mathbf{e}^{(i)}\right)^{\top} \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right. \\
& \quad-\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)^{\top} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{e}^{(i)} \\
& \left.+\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)^{\top} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right) \\
=- & 2 \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{e}^{(i)}+2 \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)=0 .
\end{aligned}
$$

```
Algorithm 4.2 Gauss-Newton Algorithm
Input: Initial parameter guess \(\hat{\mathbf{x}}^{(0)}\), data \(\mathbf{y}\), function \(g(\mathbf{x})\), Jacobian \(\mathbf{G}_{\mathbf{x}}\)
Output: Parameter estimate \(\hat{\mathbf{x}}_{\text {WLS }}\)
    Set \(i \leftarrow 0\)
    repeat
        Calculate the update direction
            \(\Delta \mathbf{x}^{(i+1)}=\left(\mathbf{G}_{\mathbf{x}}^{\mathbf{T}}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\mathbf{T}}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\)
        Calculate
            \(\hat{\mathbf{x}}^{(i+1)}=\hat{\mathbf{x}}^{(i)}+\Delta \mathbf{x}^{(i+1)}\)
        Set \(i \leftarrow i+1\)
    until Converged
    Set \(\hat{\mathbf{x}}_{\text {WLS }}=\hat{\mathbf{x}}^{(i)}\)
```

We then set the next estimate $\hat{\mathbf{x}}^{(i+1)}$ to be equal to the minimum, which gives

$$
\begin{align*}
\hat{\mathbf{x}}^{(i+1)} & =\hat{\mathbf{x}}^{(i)}+\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{e}^{(i)} \\
& =\hat{\mathbf{x}}^{(i)}+\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right) . \tag{4.9}
\end{align*}
$$

The resulting method is given in Algorithm 4.2. The algorithm produces the least squares estimate of the parameters $\hat{\mathbf{x}}_{\text {WLS }}$. Given this estimate, we can also approximate the covariance of the estimate by using the linearized model analogously to the linear case in Equation (3.38), which gives

$$
\begin{equation*}
\operatorname{Cov}\left\{\hat{\mathbf{x}}_{\mathrm{WLS}}\right\} \approx\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}_{\mathrm{WLS}}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}_{\mathrm{WLS}}\right)\right)^{-1} . \tag{4.10}
\end{equation*}
$$

A major disadvantage of the Gauss-Newton approach is that the linearization of the measurement model is local. In highly nonlinear problems, this means that it is dangerous to extrapolate too much and only small steps can be taken at a time to avoid missing local minima. This problem can be diminished by using a line search method to find a suitable step size. We discuss this kind of methods in next section.

Example 4.2. Figure 4.3 shows the result of applying Gauss-Newton algorithm to the problem that we considered in Example 4.1. It can be seen that Gauss-Newton takes a large step already in the beginning and only a few additional iterations are needed to obtain convergence.

Example 4.3. A problem of Gauss-Newton is illustrated in Figure 4.4. In this case we have started the optimization from a relatively far away point and it can be seen that Gauss-Newton's initial step is taken to a wrong direction. As can be seen in the cost function evolution, the initial step actually increases the cost - because the initial step is too large.


Figure 4.3. Comparison of gradient descent and Gauss-Newton in weighted least squared problem. The full evolution of estimates is shown on the left and the end of the estimation is zoomed on the right.


Figure 4.4. Comparison of gradient descent and Gauss-Newton in weighted least squared problem when the starting point is further away. The evolution of estimates is shown on the left and the cost $J(\mathbf{x})$ as function of iteration number is shown on the right.

### 4.4 Gauss-Newton Algorithm with Line Search

In practice, taking a full step according to (4.9) in Gauss-Newton algorithm might be too large with respect to the neighborhood for which the Taylor series approximation (4.7) is valid. To avoid this, a scaled Gauss-Newton step can be done instead, proportional to the direction given by the local approximation. This yields

$$
\begin{align*}
\hat{\mathbf{x}}^{(i+1)} & =\hat{\mathbf{x}}^{(i)}+\gamma \Delta \hat{\mathbf{x}}^{(i+1)}  \tag{4.11a}\\
\Delta \hat{\mathbf{x}}^{(i+1)} & =\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-g\left(\hat{\mathbf{x}}^{(i)}\right)\right), \tag{4.11b}
\end{align*}
$$

where $\gamma$ is the scaling factor.
One way to select the scaling parameter is to use a line search as follows. If we substitute the updated parameter vector to the cost function $J_{\mathrm{WLS}}(\mathbf{x})$, we can

```
Algorithm 4.3 Line Search on Grid
Input: Previous iterate \(\hat{\mathbf{x}}^{(i)}\), the update direction \(\Delta \hat{\mathbf{x}}^{(i+1)}\), the cost function \(J_{\mathrm{WLS}}(\mathbf{x})\), and
    the grid size \(N_{g}\).
Output: Optimal step size \(\gamma^{*}\).
    Set \(\gamma^{*} \leftarrow 0\) and \(J^{*} \leftarrow J_{\text {WLS }}\left(\hat{\mathbf{x}}^{(i)}\right)\)
    for \(j \in\left\{1,2, \ldots, N_{g}\right\}\) do
        Set \(\gamma \leftarrow j / N_{g}\)
        if \(J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}+\gamma \Delta \hat{\mathbf{x}}^{(i+1)}\right)<J^{*}\) then
            Set \(\gamma^{*} \leftarrow \gamma\)
        end if
    end for
```

consider the cost as function of the step size parameter:

$$
\begin{equation*}
J_{\mathrm{WLS}}^{(i)}(\gamma)=J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}+\gamma \Delta \hat{\mathbf{x}}^{(i+1)}\right) . \tag{4.12}
\end{equation*}
$$

Then the idea of line search is that on every step we select the step size by locally optimizing it in the range $[0,1]$.

One way to perform line search is simply by evaluating $J_{\text {WLS }}^{(i)}(\gamma)$ on a grid of values and then selecting the step size that gives the minimum. This is often called exact line search in optimization literature (Nocedal and Wright, 2006). This grid search algorithm is shown as Algorithm 4.3. It would also be possible to use more sophisticated minimum search methods such as bisection algorithm or interpolation methods to find the minimum.

However, it turns out that the line search does not need to be exact as we can guarantee to find the minimum. One simple idea is to use backtracking, where we decrease the parameter $\gamma$ until it provides a sufficient decrease in the cost. One way is to simply halve the step size until it causes a cost decrease. This kind of approach is used, for example, in Gustafsson (2018) and it often works well.

The Armijo backtracking is an inexact line search method where we demand that the cost is decreased at least with an amount that is predicted by a first order Taylor series expansions as follows. Let us expand the right hand side of (4.12) using a first order Taylor series expansion as follows (Nocedal and Wright, 2006):
$J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}+\gamma \Delta \hat{\mathbf{x}}^{(i+1)}\right) \approx J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}\right)-2 \gamma\left[\Delta \hat{\mathbf{x}}^{(i+1)}\right]^{\top} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)$.
Then we demand that the cost decrease should satisfy the Armijo condition

$$
\begin{align*}
& J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}+\gamma \Delta \hat{\mathbf{x}}^{(i+1)}\right)-J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}\right) \\
& \quad \leq-2 \beta \gamma\left[\Delta \hat{\mathbf{x}}^{(i+1)}\right]^{\mathrm{T}} \mathbf{G}_{\mathbf{x}}^{\mathrm{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right) \tag{4.14}
\end{align*}
$$

where $\beta$ is a parameter that we can choose freely on range $[0,1)$, typical parameter value being $\beta=0.1$. In the backtracking we then decrease $\gamma$ by multiplying it with

```
Algorithm 4.4 Line Search with Armijo Backtracking
Input: Previous iterate \(\hat{\mathbf{x}}^{(i)}\), the update direction \(\Delta \hat{\mathbf{x}}^{(i+1)}\), the cost function \(J_{\mathrm{WLS}}(\mathbf{x})\), and
    parameters \(\beta\) and \(\tau\).
Output: Suitable step size \(\gamma\).
    Set \(\gamma \leftarrow 1\) and \(J_{0} \leftarrow J_{\text {WLS }}\left(\hat{\mathbf{x}}^{(i)}\right)\).
    Set \(d \leftarrow-2 \beta\left[\Delta \hat{\mathbf{x}}^{(i+1)}\right]^{\top} \mathbf{G}_{\mathbf{x}}^{\mathbf{\top}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\)
    while \(J_{\text {WLS }}\left(\hat{\mathbf{x}}^{(i)}+\gamma \Delta \hat{\mathbf{x}}^{(i+1)}\right)>J_{0}+\gamma d\) do
        Set \(\gamma \leftarrow \tau \gamma\)
    end while
```

```
Algorithm 4.5 Gauss-Newton Algorithm with Line Search
Input: Initial parameter guess \(\hat{\mathbf{x}}^{(0)}\), data \(\mathbf{y}\), function \(\mathbf{g}(\mathbf{x})\), Jacobian \(\mathbf{G}_{\mathbf{x}}(\mathbf{x})\)
Output: Parameter estimate \(\hat{\mathbf{x}}_{\text {WLS }}\)
    Set \(i \leftarrow 0\)
    repeat
        Calculate the update direction
            \(\Delta \mathbf{x}^{(i+1)}=\left(\mathbf{G}_{\mathbf{x}}^{\mathbf{T}}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\mathbf{T}}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\)
        Compute optimal \(\gamma^{(i+1)}\) with line search (Algorithm 4.3 or 4.4)
        Calculate
            \(\hat{\mathbf{x}}^{(i+1)}=\hat{\mathbf{x}}^{(i)}+\gamma^{(i+1)} \Delta \mathbf{x}^{(i+1)}\)
        Set \(i \leftarrow i+1\)
    until Converged
    Set \(\hat{\mathbf{x}}_{\text {WLS }}=\hat{\mathbf{x}}^{(i)}\)
```

a parameter $\tau$ (e.g., $\tau=0.5$ ) on the range $(0,1)$ until the condition is satisfied. The resulting method is shown as Algorithm 4.4.

A Gauss-Newton algorithm with a generic line search is shown as Algorithm 4.5. The algorithm produces the least squares estimate $\hat{\mathbf{x}}_{\text {WLS }}$ and its covariance can be approximated using (4.10). More advanced line search methods can be found in optimization literature (e.g. Nocedal and Wright, 2006).

Example 4.4. Figure 4.5 shows the result of applying the Gauss-Newton algorithm with line search to the far-away initial-point case considered in Example 4.3. It can be seen that the initial step is shorter than with plain Gauss-Newton and the cost is decreased at every step. However, although the cost is strictly decreasing, the intermediate costs during the iterations are larger than with plain Gauss-Newton. However, the strict decrease in the cost enhances the stability of the algorithm.

### 4.5 Levenberg-Marquardt Algorithm

The Levenberg-Marquardt algorithm (Levenberg, 1944; Marquardt, 1963; Nocedal and Wright, 2006) uses a different strategy from line search to enhance the performance of Gauss-Newton algorithm. The underlying idea is to restrict the step


Figure 4.5. Illustration of Gauss-Newton algorithm with (exact) line search.
size by using regularization in the least squares solution following the linearization made in the Gauss-Newton algorithm.

One way of deriving the Levenberg-Marquardt algorithm is indeed to see it as a regularized Gauss-Newton algorithm. The Taylor series approximation of the nonlinear function (4.7) is only valid in the local neighborhood of the current parameter estimate $\hat{\mathbf{x}}^{(i)}$. Hence, we can use this prior information to regularize the approximated weighted least squares cost function using a regularization term as introduced in Section 3.3. This yields the cost function approximation

$$
\begin{aligned}
J_{\operatorname{ReLS}}(\mathbf{x}) \approx & \left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right)^{\top} \\
& \times \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right) \\
+ & \lambda\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)^{\top}\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right) \\
= & \left(\mathbf{e}^{(i)}-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right)^{\top} \mathbf{R}^{-1}\left(\mathbf{e}^{(i)}-\mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)\right) \\
+ & \lambda\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)^{\top}\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)
\end{aligned}
$$

where the second term is the regularization around the last iteration's estimate with weight or damping factor $\lambda$. The gradient of the cost function approximation now becomes

$$
\begin{align*}
\frac{\partial J_{\operatorname{ReLS}}(\mathbf{x})}{\partial \mathbf{x}} & \approx-2 \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{e}^{(i)}+2 \mathbf{G}_{\mathbf{x}}^{\top} \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)+2 \lambda\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right) \\
& =-2 \mathbf{G}_{\mathbf{x}}^{\top} \mathbf{R}^{-1} \mathbf{e}^{(i)}+2\left(\mathbf{G}_{\mathbf{x}}^{\top} \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)+\lambda \mathbf{I}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right) \tag{4.15}
\end{align*}
$$

Setting (4.15) to zero and rearranging yields

$$
\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)+\lambda \mathbf{I}\right)\left(\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right)=\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{e}^{(i)}
$$

and thus,

$$
\begin{align*}
\hat{\mathbf{x}}^{(i+1)} & =\hat{\mathbf{x}}^{(i)}+\Delta \hat{\mathbf{x}}^{(i+1)}  \tag{4.16a}\\
\Delta \hat{\mathbf{x}}^{(i+1)} & =\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)+\lambda \mathbf{I}\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right) \tag{4.16b}
\end{align*}
$$

Equation (4.16b) shows that when $\lambda$ approaches zero, the algorithm converges to the Gauss-Newton algorithm. On the other hand, for large values of $\lambda$, the term $\lambda \mathbf{I}$ will dominate. In this case, we have that

$$
\Delta \hat{\mathbf{x}}^{(i+1)} \approx \frac{1}{\lambda} \mathbf{G}_{\mathbf{x}}^{\mathbf{\top}} \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)
$$

which is a scaled gradient descent step. Thus, an important question is how to choose $\lambda$. Ideally, in flat regions of the cost function where the linear approximation is good, $\lambda$ should be chosen small, whereas in steep regions, it should be chosen large. Unfortunately, there is no definite method on how to chose and adapt $\lambda$ and several different approaches have been proposed.

A simple strategy which is often used, and which is a simplified version of the method proposed by Marquardt (1963) and used, for example, in Pujol (2007) is to start from some damping value, for example, $\lambda^{(0)}$ (e.g. $\lambda^{(0)}=10^{-2}$ ) and select a fixed factor $\nu$ (e.g. $\nu=10$ ). Then at each step we do the following:

- First compute using Equation (4.16) a candidate $\hat{\mathbf{x}}^{(i+1)}$ using the previous parameter value $\lambda^{(i-1)}$. Then proceed as follows:
- If $J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i+1)}\right)<J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}\right)$ then accept $\hat{\mathbf{x}}^{(i+1)}$ and decrease the damping parameter by $\lambda^{(i)}=\lambda^{(i-1)} / \nu$.
- Otherwise continue with $\hat{\mathbf{x}}^{(i+1)}=\hat{\mathbf{x}}^{(i)}$ and increase the damping parameter by $\lambda^{(i)}=\nu \lambda^{(i-1)}$.

Furthermore, the regularization procedure in Equation (4.16) has the disadvantage that it is not scale invariant. To make the regularization scale-invariant, Marquardt (1963) suggested to normalize the regularized cost function approximation by using the diagonal values of $\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)$. This procedure is equivalent to replacing the regularization term $\lambda \mathbf{I}$ with $\lambda \operatorname{diag}\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right.$ ), which yields to the following scaled parameter update:

$$
\begin{align*}
\hat{\mathbf{x}}^{(i+1)}= & \hat{\mathbf{x}}^{(i)}+\Delta \hat{\mathbf{x}}^{(i+1)},  \tag{4.17a}\\
\Delta \hat{\mathbf{x}}^{(i+1)}=( & \left.\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)+\lambda \operatorname{diag}\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\right)^{-1} \\
& \times \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right) . \tag{4.17b}
\end{align*}
$$

The resulting Levenberg-Marquardt algorithm, including both the non-scaled and scaled versions, is shown in Algorithm 4.6.

```
Algorithm 4.6 Levenberg-Marquardt Algorithm with simple adaptation
Input: Initial parameter guess \(\hat{\mathbf{x}}^{(0)}\), data \(\mathbf{y}\), function \(\mathbf{g}(\mathbf{x})\), Jacobian \(\mathbf{G}_{\mathbf{x}}\), initial damping
    \(\lambda^{(0)}\), and parameter \(\nu\).
Output: Parameter estimate \(\hat{\mathbf{x}}_{\text {WLS }}\)
    Set \(i \leftarrow 0\) and \(\lambda \leftarrow \lambda^{(0)}\).
    repeat
        Compute the (candidate) parameter update:
        if using scaled version then
            \(\begin{aligned} \Delta \hat{\mathbf{x}}^{(i+1)}=( & \left.\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)+\lambda \operatorname{diag}\left(\mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\right)^{-1} \\ & \times \mathbf{G}_{\mathbf{x}}^{\top}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\end{aligned}\)
            else
                    \(\Delta \hat{\mathbf{x}}^{(i+1)}=\left(\mathbf{G}_{\mathbf{x}}^{\boldsymbol{\top}}\left(\hat{\mathbf{x}}^{(i)}\right) \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}\left(\hat{\mathbf{x}}^{(i)}\right)+\lambda \mathbf{I}\right)^{-1} \mathbf{G}_{\mathbf{x}}^{\boldsymbol{\top}} \mathbf{R}^{-1}\left(\mathbf{y}-\mathbf{g}\left(\hat{\mathbf{x}}^{(i)}\right)\right)\)
        end if
        if \(J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}+\Delta \hat{\mathbf{x}}^{(i+1)}\right)<J_{\mathrm{WLS}}\left(\hat{\mathbf{x}}^{(i)}\right)\) then
            Accept the candidate and decrease \(\lambda\) :
                \(\begin{aligned} \hat{\mathbf{x}}^{(i+1)} & =\hat{\mathbf{x}}^{(i)}+\Delta \hat{\mathbf{x}}^{(i+1)} \\ \lambda & \leftarrow \lambda / \nu\end{aligned}\)
            Set \(i \leftarrow i+1\)
        else
            Reject the candidate and increase \(\lambda\) :
                                    \(\lambda \leftarrow \nu \lambda\)
        end if
    until Converged
    Set \(\hat{\mathbf{x}}_{\text {WLS }}=\hat{\mathbf{x}}^{(i)}\)
```

Example 4.5. Figure 4.6 illustrates the operation of the Levenberg-Marquardt algorithm in comparison to Gauss-Newton method with line search. It can be seen that the performance of the algorithm is similar to the Gauss-Newton although the path that the optimization takes is quite much different. Furthermore, Figure 4.7 shows the evolution of errors in all the considered optimization methods. It can be seen that the performance of the methods is similar in this case although exact path that the cost function takes differs between the methods.

### 4.6 Regularized Non-Linear Models

Although the optimization algorithms in this chapter have been formulated to compute solutions to non-regularized linear squares problems, they can also be used to compute solutions to regularized problems. Similarly to Section 3.3, we can rewrite


Figure 4.6. Illustration of the operation of the Levenberg-Marquardt algorithm.
the regularized cost as follows:

$$
\begin{align*}
J_{\operatorname{ReLS}}(\mathbf{x}) & =(\mathbf{y}-\mathbf{g}(\mathbf{x}))^{\top} \mathbf{R}^{-1}(\mathbf{y}-\mathbf{g}(\mathbf{x}))+(\mathbf{m}-\mathbf{x})^{\top} \mathbf{P}^{-1}(\mathbf{m}-\mathbf{x}) \\
& =\left(\left[\begin{array}{c}
\mathbf{y} \\
\mathbf{m}
\end{array}\right]-\left[\begin{array}{c}
\mathbf{g}(\mathbf{x}) \\
\mathbf{x}
\end{array}\right]\right)^{\top}\left[\begin{array}{cc}
\mathbf{R}^{-1} & 0 \\
0 & \mathbf{P}^{-1}
\end{array}\right]\left(\left[\begin{array}{c}
\mathbf{y} \\
\mathbf{m}
\end{array}\right]-\left[\begin{array}{c}
\mathbf{g}(\mathbf{x}) \\
\mathbf{x}
\end{array}\right]\right) \tag{4.18}
\end{align*}
$$

which has the form of a non-regularized cost and hence all the algorithms presented in this chapter are again applicable.

### 4.7 Quasi-Newton Methods

Although here we have only concentrated on least squares problems and restricted our consideration to gradient descent, Gauss-Newton, and Levenberg-Marguardt algorithms, there exist other algorithms that can be used for solving these optimization problems. One class of algorithms is based on the classical Newton's method as follows.

Let us consider a generic cost function $J(\mathbf{x})$ which we wish to minimize. The so-called Newton's method is based on the following iterative procedure. Assume that our current guess for the minimum is $\mathbf{x}^{(i)}$. We can now Taylor expand the cost


Figure 4.7. Evolution of cost in all the considered optimization methods.
function as follows:

$$
\begin{align*}
J(\mathbf{x}) \approx & J\left(\mathbf{x}^{(i)}\right)+\left.\left[\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}}\right]^{\top}\right|_{\mathbf{x}=\mathbf{x}^{(i)}}\left(\mathbf{x}-\mathbf{x}^{(i)}\right) \\
& +\left.\frac{1}{2}\left(\mathbf{x}-\mathbf{x}^{(i)}\right)^{\top}\left[\frac{\partial^{2} J(\mathbf{x})}{\partial \mathbf{x}^{2}}\right]\right|_{\mathbf{x}=\mathbf{x}^{(i)}}\left(\mathbf{x}-\mathbf{x}^{(i)}\right), \tag{4.19}
\end{align*}
$$

where $\partial^{2} J(\mathbf{x}) / \partial \mathbf{x}^{2}$ denotes the Hessian matrix of $J(\mathbf{x})$. The strategy is now to minimize the right hand side with respect to x and use the result as the next guess. The resulting Newton's method takes the following form:

$$
\begin{equation*}
\mathbf{x}^{(i+1)}=\mathbf{x}^{(i)}-\left.\left[\frac{\partial^{2} J(\mathbf{x})}{\partial \mathbf{x}^{2}}\right]^{-1} \frac{\partial J(\mathbf{x})}{\partial \mathbf{x}}\right|_{\mathbf{x}=\mathbf{x}^{(i)}} . \tag{4.20}
\end{equation*}
$$

Unfortunately, computation of the Hessian is often not desirable and therefore, in so called quasi-Newton methods the Hessian is approximated. This approximation can be formed in various ways which leads to multiple classes of quasi-Newton methods (Nocedal and Wright, 2006) and among them the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm has turned out to be particularly successful. In fact, the Gauss-Newton method can also be seen as a quasi-Newton method where we approximate the Hessian by $2 \mathbf{G}_{\mathbf{x}}^{\top} \mathbf{R}^{-1} \mathbf{G}_{\mathbf{x}}$. The line search procedure is also typically an essential part of quasi-Newton methods.

### 4.8 Convergence Criteria

The algorithms introduced in the previous sections are all based on an iterative scheme, where the parameter estimates are improved iteratively, based on the last estimate. An important question therefore is, when to terminate the search. Three simple and easy to check criteria are:

- The absolute or relative change in the parameter estimate falls below a threshold. We can, for example, stop iterations when the vector norm of the update direction is below a threshold: $\left\|\Delta \mathbf{x}^{(i)}\right\|<\epsilon_{p}$.
- The absolute or relative change in the cost falls below a certain threshold. For example, a criterion could be $\left|\left(J\left(\mathbf{x}^{i}\right)-J\left(\mathbf{x}^{(i+1)}\right)\right) / J\left(\mathbf{x}^{i}\right)\right|<\epsilon_{c}$.
- A maximum number of iterations is reached.

It is common practice to monitor all or at least a subset of these criteria (plus possibly others, too) and terminate the search when at least one of the criteria is fulfilled. In order to avoid infinite loops, the last criterion (a maximum number of iterations) should always be checked for, possibly issuing a warning when this is the reason for termination of the algorithm.

## Chapter 5

## Dynamic Models

Most sensor fusion problems involve dynamically changing variables. For example, in autonomous driving, other road users continuously appear, disappear, and move in the traffic scene. In this case, we are interested in estimating dynamically varying parameters. We refer to these parameters as states or the state, because they describe the state a system is in (e.g., the location and velocity of a car, or the position, attitude, and velocity of a unmanned aerial vehicle, etc.).

In such dynamic scenarios, we have to employ sensors that provide repeated measurements in time, that is, they sample periodically. In between those samples, the state of the system evolves according to some process and thus, the states at different times are different from each other. However, since the evolution of the state follows some dynamic process, they are related and thus, we can relate the states at two times to each other. To do that, we need a principled way of describing how the state evolves between samples. We find a suitable approach in differential equations (for continuous-time processes) and difference equations (for discrete-time processes), which can be used to describe dynamic processes. Indeed, we will make use of differential equations to derive state-space models that turn an $L$ th order differential (or difference) equation into a first order vectorvalued differential (or difference) equation.

### 5.1 Continuous-Time State-Space Models

### 5.1.1 Deterministic Linear State-Space Models

We start the derivation of the state-space formulation based on an example. Figure 5.1 shows a spring-damper system, a typical mechanical system. This system is governed by Newton's second law of motion, from which we can find the following inhomogeneous ordinary differential equation (ODE):

$$
\begin{equation*}
m a(t)=-k p(t)-\eta v(t)+u(t) \tag{5.1}
\end{equation*}
$$



Figure 5.1. Example of a mechanical dynamic system: A spring-damper system with forcing function $u(t)$. The positive direction of the motion $p$ is defined to the right.
where $p(t), v(t)$, and $a(t)$ denote the position, velocity, and acceleration of the mass $m$, respectively, $k$ is the spring constant, and $\eta$ is the damping constant. Furthermore, the forcing function $u(t)$ is a deterministic input to the system.

By introducing a second equation, $v(t)=v(t)$, which always holds, and dividing (5.1) by $m$, we obtain the equation system

$$
\begin{align*}
& v(t)=v(t)  \tag{5.2a}\\
& a(t)=-\frac{k}{m} p(t)-\frac{\eta}{m} v(t)+\frac{1}{m} u(t) \tag{5.2b}
\end{align*}
$$

This can be rewritten in matrix form, such that we obtain

$$
\left[\begin{array}{c}
v(t)  \tag{5.3}\\
a(t)
\end{array}\right]=\left[\begin{array}{cc}
0 & 1 \\
-\frac{k}{m} & -\frac{\eta}{m}
\end{array}\right]\left[\begin{array}{c}
p(t) \\
v(t)
\end{array}\right]+\left[\begin{array}{c}
0 \\
\frac{1}{m}
\end{array}\right] u(t)
$$

Observe that in (5.3), the vector on the left hand side of the equation is the derivative of the vector on the right hand side. Hence, we can define the vector

$$
\mathbf{x}(t)=\left[\begin{array}{l}
p(t) \\
v(t)
\end{array}\right]
$$

and rewrite (5.3) as

$$
\dot{\mathbf{x}}(t)=\left[\begin{array}{cc}
0 & 1  \tag{5.4}\\
-\frac{k}{m} & -\frac{\eta}{m}
\end{array}\right] \mathbf{x}(t)+\left[\begin{array}{c}
0 \\
\frac{1}{m}
\end{array}\right] u(t)
$$

where

$$
\dot{\mathbf{x}}(t)=\frac{\mathrm{d} \mathbf{x}(t)}{\mathrm{d} t}
$$

denotes the time-derivative.
The form (5.4) is the state-space representation of the dynamic system in Figure 5.1 and the vector $\mathbf{x}(t)$ is called the state vector (or simply state). The state vector now encodes all the information about the state the system is in (position, speed, etc.). Therefore, solving this first-order vector valued ODE representation rather than the original differential equation in (5.1), in some sense, provides richer information about the system.

We can now generalize this approach. Consider the $L$ th order ODE of the form

$$
\begin{equation*}
\frac{\mathrm{d}^{L} x(t)}{\mathrm{d} t^{L}}=a_{0} x(t)+a_{1} \frac{\mathrm{~d} x(t)}{\mathrm{d} t}+\cdots+a_{L-1} \frac{\mathrm{~d}^{L-1} x(t)}{\mathrm{d} t^{L-1}}+b_{1} u(t) \tag{5.5}
\end{equation*}
$$

Then we can introduce $L-1$ equations of the form $\mathrm{d}^{l} x(t) / \mathrm{d} t^{l}=\mathrm{d}^{l} x(t) / \mathrm{d} t^{l}$ to obtain the equation system

$$
\begin{align*}
\frac{\mathrm{d} x(t)}{\mathrm{d} t} & =\frac{\mathrm{d} x(t)}{\mathrm{d} t}  \tag{5.6a}\\
\frac{\mathrm{~d}^{2} x(t)}{\mathrm{d} t^{2}} & =\frac{\mathrm{d}^{2} x(t)}{\mathrm{d} t^{2}} \tag{5.6b}
\end{align*}
$$

$$
\begin{equation*}
\frac{\mathrm{d}^{L} x(t)}{\mathrm{d} t^{L}}=a_{0} x(t)+a_{1} \frac{\mathrm{~d} x(t)}{\mathrm{d} t}+\cdots+a_{L-1} \frac{\mathrm{~d}^{L-1} x(t)}{\mathrm{d} t^{L-1}}+b_{1} u(t) \tag{5.6c}
\end{equation*}
$$

Rewriting (5.6) in vector form yields

$$
\left[\begin{array}{c}
\frac{\mathrm{d} x(t)}{\mathrm{d} t} \\
\frac{\mathrm{~d}^{2} x(t)}{\mathrm{d} t^{2}} \\
\vdots \\
\frac{\mathrm{~d}^{L} x(t)}{\mathrm{d} t^{L}}
\end{array}\right]=\left[\begin{array}{ccccc}
0 & 1 & 0 & \ldots & 0 \\
0 & 0 & 1 & & \vdots \\
\vdots & & & \ddots & 0 \\
a_{0} & a_{1} & & \ldots & a_{L-1}
\end{array}\right]\left[\begin{array}{c}
x(t) \\
\frac{\mathrm{d} x(t)}{\mathrm{d} t} \\
\vdots \\
\frac{\mathrm{~d}^{L-1} x(t)}{\mathrm{d} t^{L-1}}
\end{array}\right]+\left[\begin{array}{c}
0 \\
0 \\
\vdots \\
b_{1}
\end{array}\right] u(t) .
$$

This in turn, can be written compactly as the dynamic model

$$
\begin{equation*}
\dot{\mathbf{x}}(t)=\mathbf{A} \mathbf{x}(t)+\mathbf{B}_{u} \mathbf{u}(t) \tag{5.7}
\end{equation*}
$$

where
$\mathbf{x}(t)=\left[\begin{array}{c}x(t) \\ \frac{\mathrm{d} x(t)}{\mathrm{d} t} \\ \vdots \\ \frac{\mathrm{~d}^{L-1} x(t)}{\mathrm{d} t^{L-1}}\end{array}\right], \mathbf{A}=\left[\begin{array}{ccccc}0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 1 & & \vdots \\ \vdots & & & \ddots & 0 \\ a_{0} & a_{1} & & \ldots & a_{L-1}\end{array}\right], \mathbf{B}_{u}=\left[\begin{array}{c}0 \\ 0 \\ \vdots \\ b_{1}\end{array}\right], \mathbf{u}(t)=u(t)$
While quite general, Equation (5.7) is a state-space representation for models of the type in (5.5). However, we can also derive the same type of representation for other dynamic models. For example, consider a freshly brewed hot cup of coffee. Newton's law of cooling then states that the change in temperature of the coffee is proportional to the temperature difference between the coffee and its surrounding (assuming that the surroundings are much larger than the coffee cup). This can be formulated as a differential equation of the form

$$
\begin{equation*}
\frac{\mathrm{d} T_{\mathrm{c}}(t)}{\mathrm{d} t}=-k_{1}\left(T_{\mathrm{c}}(t)-T_{\mathrm{r}}(t)\right) \tag{5.9}
\end{equation*}
$$

where $T_{\mathrm{c}}(t)$ denotes the coffee's temperature, $T_{\mathrm{r}}(t)$ the room temperature, and $k_{1}$ is a constant. The change in room temperature on the other hand depends on the
heat that is produced inside the room (e.g., due to heating) as well as the losses to the outside. Again assuming that Newton's law of cooling applies, we can write the differential equation as

$$
\begin{equation*}
\frac{\mathrm{d} T_{\mathrm{r}}(t)}{\mathrm{d} t}=-k_{2}\left(T_{\mathrm{r}}(t)-T_{\mathrm{a}}(t)\right)+h(t) \tag{5.10}
\end{equation*}
$$

where $k_{2}$ is a constant, $T_{\mathrm{a}}(t)$ denotes the outside temperature and $h(t)$ is the heating input. Observe that (5.9)-(5.10) is a coupled system that can be written as the equation system

$$
\begin{align*}
& \frac{\mathrm{d} T_{\mathrm{r}}(t)}{\mathrm{d} t}=-k_{2}\left(T_{\mathrm{r}}(t)-T_{\mathrm{a}}(t)\right)+h(t)  \tag{5.11a}\\
& \frac{\mathrm{d} T_{\mathrm{c}}(t)}{\mathrm{d} t}=-k_{1}\left(T_{\mathrm{c}}(t)-T_{\mathrm{r}}(t)\right) \tag{5.11b}
\end{align*}
$$

or equivalently on matrix form as

$$
\left[\begin{array}{c}
\frac{\mathrm{d} T_{\mathrm{r}}(t)}{\mathrm{d} t}  \tag{5.12}\\
\frac{\mathrm{~d} T_{\mathrm{c}}(t)}{\mathrm{d} t}
\end{array}\right]=\left[\begin{array}{cc}
-k_{2} & 0 \\
k_{1} & -k_{1}
\end{array}\right]\left[\begin{array}{c}
T_{\mathrm{r}}(t) \\
T_{\mathrm{c}}(t)
\end{array}\right]+\left[\begin{array}{cc}
k_{2} & 1 \\
0 & 0
\end{array}\right]\left[\begin{array}{c}
T_{\mathrm{a}}(t) \\
h(t)
\end{array}\right]
$$

This has again the form of (5.7), where

$$
\mathbf{x}(t)=\left[\begin{array}{c}
T_{\mathrm{r}}(t)  \tag{5.13}\\
T_{\mathrm{c}}(t)
\end{array}\right], \mathbf{A}=\left[\begin{array}{cc}
-k_{2} & 0 \\
k_{1} & -k_{1}
\end{array}\right], \mathbf{B}_{u}=\left[\begin{array}{cc}
k_{2} & 1 \\
0 & 0
\end{array}\right], \mathbf{u}(t)=\left[\begin{array}{c}
T_{\mathrm{a}}(t) \\
h(t)
\end{array}\right]
$$

Hence, the model (5.7) is a quite general dynamic model of the time-varying behavior for many different processes. Recall that our original aim was to estimate the dynamically changing state $\mathbf{x}(t)$. This requires that we measure $\mathbf{x}(t)$ in some way, that is, we need to combine the dynamic model with a measurement model. The simplest case is that of a linear measurement model of the form (3.10) where we replace the static $\mathbf{x}$ with time varying state $\mathbf{x}(t)$. Since the measurements are obtained at discrete time instances $t_{n}$ (i.e., at $t_{1}, t_{2}, \ldots$ ), the measurement model becomes

$$
\begin{equation*}
\mathbf{y}_{n}=\mathbf{G} \mathbf{x}\left(t_{n}\right)+\mathbf{r}_{n}, \tag{5.14}
\end{equation*}
$$

where $\mathbf{r}_{n}$ denotes the measurement noise with covariance matrix $\operatorname{Cov}\left\{\mathbf{r}_{n}\right\}=\mathbf{R}_{n}$.
Defining $\mathbf{x}_{n} \triangleq \mathbf{x}\left(t_{n}\right)$ and combining the dynamic model (5.7) and the measurement model (5.14) then yields the deterministic linear state space model

$$
\begin{align*}
\dot{\mathbf{x}}(t) & =\mathbf{A} \mathbf{x}(t)+\mathbf{B}_{u} \mathbf{u}(t)  \tag{5.15a}\\
\mathbf{y}_{n} & =\mathbf{G} \mathbf{x}_{n}+\mathbf{r}_{n} \tag{5.15b}
\end{align*}
$$

### 5.1.2 Stochastic Linear State-Space Models

Often, the behavior of dynamic systems is not completely deterministic, or the input $u(t)$ is not known. For example, when tracking an airplane using radar, we


Figure 5.2. Example of a white noise process $w(t)$ with $\sigma_{w}^{2}=2$. (a) One realization of the process, and (b) its power spectral density $S_{w}(\omega)$ averaged over 100 realizations.
do not know the inputs made by the pilots. Additionally, variables such as wind or pressure fields are not known either and it is difficult to model them accurately. Instead, such random effects can be modeled by including a stochastic process as an input to the differential equation (and hence to its state-space representation), which turns an ODE into a stochastic differential equation (SDE; Øksendal, 2010; Särkkä and Solin, 2019).

A simple class of stochastic processes are stationary stochastic processes which can be characterized by their autocorrelation functions. The autocorrelation function of a stationary stochastic process $w(t)$ is

$$
R_{w w}(\tau)=\mathrm{E}\{w(t+\tau) w(t)\}
$$

An equivalent characterization is the power spectral density (which is the Fourier transform of the autocorrelation function; Papoulis (1984)). A suitable assumption in many models is that the stochastic process is a white noise process. In this case, the autocorrelation function is

$$
\begin{equation*}
R_{w w}(\tau)=\sigma_{w}^{2} \delta(\tau) \tag{5.16}
\end{equation*}
$$

where $\delta(\tau)$ denotes the Dirac delta function. Hence, the power spectral density is a constant, that is,

$$
\begin{equation*}
S_{w w}=\sigma_{w}^{2} \tag{5.17}
\end{equation*}
$$

In other words, the white noise process contains equal contributions from each frequency. One realization of such a process, together with its power spectral density are shown in Figure 5.2.

The derivation of the state-space representation of the resulting SDE follows the same steps as for the deterministic case, where the random process takes the role of the input. Consider the $L$ th order SDE given by

$$
\begin{equation*}
\frac{\mathrm{d}^{L} x(t)}{\mathrm{d} t^{L}}=a_{0} x(t)+a_{1} \frac{\mathrm{~d} x(t)}{\mathrm{d} t}+\cdots+a_{L-1} \frac{\mathrm{~d}^{L-1} x(t)}{\mathrm{d} t^{L-1}}+b_{1} w(t) \tag{5.18}
\end{equation*}
$$



Figure 5.3. Illustration of a spacecraft orbiting the earth.
which is of the same form as (5.5) but with the stochastic process $w(t)$ as the input. Rewriting (5.18) into matrix form yields

$$
\left[\begin{array}{c}
\frac{\mathrm{d} x(t)}{\mathrm{d} t} \\
\frac{\mathrm{~d}^{2} x(t)}{\mathrm{d} t^{2}} \\
\vdots \\
\frac{\mathrm{~d}^{L} x(t)}{\mathrm{d} t^{L}}
\end{array}\right]=\left[\begin{array}{ccccc}
0 & 1 & 0 & \ldots & 0 \\
0 & 0 & 1 & & \vdots \\
\vdots & & & \ddots & 0 \\
a_{0} & a_{1} & & \ldots & a_{L-1}
\end{array}\right]\left[\begin{array}{c}
x(t) \\
\frac{\mathrm{d} x(t)}{\mathrm{d} t} \\
\vdots \\
\frac{\mathrm{~d}^{L-1} x(t)}{\mathrm{d} t^{L-1}}
\end{array}\right]+\left[\begin{array}{c}
0 \\
0 \\
\vdots \\
b_{1}
\end{array}\right] w(t) .
$$

Hence, this can also be written compactly in the form of a first order vector valued SDE system

$$
\dot{\mathbf{x}}(t)=\mathbf{A} \mathbf{x}(t)+\mathbf{B}_{w} \mathbf{w}(t)
$$

Naturally, also coupled models (like the coffee-cup example) can be written in this form. Together with a linear measurement equation, the stochastic linear statespace model becomes

$$
\begin{align*}
\dot{\mathbf{x}}(t) & =\mathbf{A} \mathbf{x}(t)+\mathbf{B}_{w} \mathbf{w}(t)  \tag{5.19a}\\
\mathbf{y}_{n} & =\mathbf{G} \mathbf{x}_{n}+\mathbf{r}_{n} \tag{5.19b}
\end{align*}
$$

Note that some dynamic models may contain both deterministic inputs $\mathbf{u}(t)$ and stochastic inputs $\mathbf{w}(t)$. Naturally, both of these can be incorporated in the model as well.

### 5.1.3 Nonlinear State-Space Models

So far, we have only discussed the state-space form of linear ODEs and SDEs. However, many dynamic systems actually behave nonlinearly. Fortunately, we can handle nonlinear models with a similar formalism.

We start our discussion again based on an example. Figure 5.3 shows an illustration of a spacecraft orbiting the earth. The forces acting upon the craft are the gravitational pull of the earth $F_{g}$ and the propulsion force $F_{p}$. The position $\mathbf{p}(t)$ of the craft is defined according to the coordinate system shown in Figure 5.3, with the origin at the center of the earth. The gravitational acceleration of an object at a distance $|\mathbf{p}(t)|$ from the earth's center is approximately

$$
g \approx g_{0}\left(\frac{r_{e}}{|\mathbf{p}(t)|}\right)^{2}
$$

where $g_{0}=9.81 \mathrm{~m} / \mathrm{s}^{2}$ is the gravitational acceleration on the earth's surface and $r_{e}=6371 \mathrm{~km}$ is the mean earth radius. The gravitational pull is in the opposite direction of $\mathbf{p}(t)$ and hence, we can write it on vector form as

$$
\begin{aligned}
\mathbf{F}_{g} & =-m g_{0}\left(\frac{r_{e}}{|\mathbf{p}(t)|}\right)^{2} \frac{\mathbf{p}(t)}{|\mathbf{p}(t)|} \\
& =-m g_{0} r_{e}^{2} \frac{\mathbf{p}(t)}{|\mathbf{p}(t)|^{3}}
\end{aligned}
$$

The propulsion force is perpendicular to the direction of the position vector $\mathbf{p}(t)$. Hence, we can write

$$
\mathbf{F}_{p}=F_{p} \frac{1}{|\mathbf{p}(t)|}\left[\begin{array}{c}
-p^{y}(t) \\
p^{x}(t)
\end{array}\right]
$$

When tracking the spacecraft from the ground, for example using radar, the magnitude of the propulsion force $F_{p}$ is unknown. But since we know that the engines are only used to make small flight path corrections to conserve fuel, we can model this as a stochastic process, that is, we can assume that $F_{p}=w(t)$ with some spectral density $\sigma_{w}^{2}$.

Using Newton's second law it follows that the motion of the spacecraft is governed by the vector-valued differential equation

$$
m \mathbf{a}(t)=-m g_{0} r_{e}^{2} \frac{\mathbf{p}(t)}{|\mathbf{p}(t)|^{3}}+\frac{1}{|\mathbf{p}(t)|}\left[\begin{array}{c}
-p^{y}(t)  \tag{5.20}\\
p^{x}(t)
\end{array}\right] w(t)
$$

The highest order of the derivative here is the acceleration on the left hand side. Hence, a suitable state vector includes the position (zeroth derivative) and the speed (first derivative), that is,

$$
\mathbf{x}(t)=\left[\begin{array}{llll}
p^{x}(t) & p^{y}(t) & v^{x}(t) & v^{y}(t)
\end{array}\right]^{\top}
$$

However, when trying to rewrite (5.20) into the form of a state-space representation (5.19), we notice that this is not possible since the right hand side of (5.20) is not linear in $\mathbf{p}(t)$. Nevertheless, we can still write it as a nonlinear equation system
in terms of $\mathbf{x}(t)$ as

$$
\begin{align*}
{\left[\begin{array}{c}
v^{x}(t) \\
v^{y}(t) \\
a^{x}(t) \\
a^{y}(t)
\end{array}\right] } & =\left[\begin{array}{c}
v^{x}(t) \\
v^{y}(t) \\
-g_{0} r_{e}^{2} \frac{p^{x}(t)}{|\mathbf{p}(t)|^{3}} \\
-g_{0} r_{e}^{2} \frac{p^{y}(t)}{|\mathbf{p}(t)|^{3}}
\end{array}\right]+\left[\begin{array}{c}
0 \\
0 \\
-\frac{p^{y}(t)}{m|\mathbf{p}(t)|} \\
\frac{p^{x}}{m|\mathbf{p}(t)|}
\end{array}\right] w(t)  \tag{5.21}\\
& =\left[\begin{array}{c}
f_{1}(\mathbf{x}(t)) \\
f_{2}(\mathbf{x}(t)) \\
f_{3}(\mathbf{x}(t)) \\
f_{4}(\mathbf{x}(t))
\end{array}\right]+\left[\begin{array}{c}
0 \\
0 \\
-\frac{p^{y}(t)}{m|\mathbf{p}(t)|} \\
\frac{p^{x}(t) \mid}{m|\mathbf{p}(t)|}
\end{array}\right] w(t)
\end{align*}
$$

where the $f_{i}(\mathbf{x}(t))$ are nonlinear functions of $\mathbf{x}(t)$.
This idea can be generalized to any arbitrary nonlinear dynamic system that is described by a nonlinear ODE or SDE, including coupled systems. Given the state vector $\mathbf{x}(t)=\left[\begin{array}{llll}x_{1}(t) & x_{2}(t) & \ldots & x_{d_{x}}(t)\end{array}\right]^{\top}$ of dimension $d_{x}$, we can write the ODE/SDE as a vector-valued, nonlinear equation system of the form

$$
\left[\begin{array}{c}
\dot{x}_{1}(t)  \tag{5.22}\\
\dot{x}_{2}(t) \\
\vdots \\
\dot{x}_{d_{x}}(t)
\end{array}\right]=\left[\begin{array}{c}
f_{1}(\mathbf{x}(t)) \\
f_{2}(\mathbf{x}(t)) \\
\vdots \\
f_{d_{x}}(\mathbf{x}(t))
\end{array}\right]+\left[\begin{array}{ccc}
b_{11}(\mathbf{x}(t)) & \ldots & b_{1 d_{w}}(\mathbf{x}(t)) \\
b_{21}(\mathbf{x}(t)) & & \vdots \\
\vdots & \ddots & \\
b_{d_{x} 1}(\mathbf{x}(t)) & \ldots & b_{d_{x} d_{w}}(\mathbf{x}(t))
\end{array}\right] \mathbf{w}(t)
$$

More compactly, this can be written as

$$
\begin{equation*}
\dot{\mathbf{x}}(t)=\mathbf{f}(\mathbf{x}(t))+\mathbf{B}_{w}(\mathbf{x}(t)) \mathbf{w}(t) \tag{5.23}
\end{equation*}
$$

where $\mathbf{f}(\mathbf{x}(t))$ is a vector-valued nonlinear function of the state, and $\mathbf{B}_{w}(\mathbf{x}(t))$ is a matrix that depends on the state $\mathbf{x}(t)$.

In order to estimated the state $\mathbf{x}(t)$ based on observations $\mathbf{y}_{n}$, we again need to relate the measurements to the state. This is achieved by combining the dynamic model (5.23) with a measurement model as introduced in Chapters 3-4. The most general model arises if we chose a nonlinear observation model of the form (4.1) where the state $\mathbf{x}_{n} \triangleq \mathbf{x}\left(t_{n}\right)$ at time $t_{n}$ takes the place of the unknown $\mathbf{x}$. This yields the general nonlinear state-space model

$$
\begin{align*}
\dot{\mathbf{x}}(t) & =\mathbf{f}(\mathbf{x}(t))+\mathbf{B}_{w}(\mathbf{x}(t)) \mathbf{w}(t)  \tag{5.24a}\\
\mathbf{y}_{n} & =\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n}, \tag{5.24b}
\end{align*}
$$

where $\mathbf{f}(\mathbf{x}(t))$ is the nonlinear state transition function, $\mathbf{w}(t)$ is the driving stochastic process, $\mathbf{g}\left(\mathbf{x}_{n}\right)$ is the measurement model, and $\mathbf{r}_{n}$ is the measurement noise.

### 5.2 Discrete-Time State-Space Models

### 5.2.1 Deterministic Linear State-Space Models

As discussed in the previous section, ODEs and SDEs are a natural way for modeling the behavior of dynamic systems. Transforming them to state-space form furthermore provides a way of describing how the state of a system evolves over time.

However, not all dynamic systems can be modeled by using differential equations. For some systems, the state is only defined at some discrete points in time $t_{1}, t_{2}, \ldots$, that is, only in discrete-time. In this case, another approach than differential equations is needed to model the dynamic behavior. The discrete-time equivalent to differential equations are difference equations that describe the relationship between the value of a variable at time $t_{n}$ to its previous values at times $t_{n-1}, t_{n-2}, \ldots$

The process of obtaining a state-space representation from difference equations is closely related to the approach used for differential equations. Consider a $d_{x}$ th order equation system with discrete-time, deterministic inputs $u_{1, n} \triangleq u_{1}\left(t_{n}\right), u_{2, n} \triangleq u_{2}\left(t_{n}\right), \ldots, u_{d_{u}, n} \triangleq u_{d_{u}}\left(t_{n}\right)$ of the form

$$
\begin{aligned}
& x_{1, n}=a_{11} x_{1, n-1}+\cdots+a_{1 d_{x}} x_{d_{x}, n-1}+b_{11} u_{1, n}+\cdots+b_{1 d_{u}} u_{d_{u}, n} \\
& x_{2, n}=a_{21} x_{1, n-1}+\cdots+a_{2 d_{x}} x_{d_{x}, n-1}+b_{21} u_{1, n}+\cdots+b_{2 d_{u}} u_{d_{u}, n} \\
& \quad \vdots \\
& x_{d_{x}, n}=a_{d_{x} 1} x_{1, n-1}+\cdots+a_{d_{x} d_{x}} x_{d_{x}, n-1}+b_{d_{x} 1} u_{1, n}+\cdots+b_{d_{x} d_{u}} u_{d_{u}, n}
\end{aligned}
$$

where $x_{i, n} \triangleq x_{i}\left(t_{n}\right)$ is the $i$ th variable at time $t_{n}$. This can be rewritten in matrix form according to

$$
\left[\begin{array}{c}
x_{1, n} \\
\vdots \\
x_{d_{x}, n}
\end{array}\right]=\left[\begin{array}{ccc}
a_{11} & \ldots & a_{1 d_{x}} \\
\vdots & \ddots & \vdots \\
a_{d_{x} 1} & \ldots & a_{d_{x} d_{x}}
\end{array}\right]\left[\begin{array}{c}
x_{1, n-1} \\
\vdots \\
x_{d_{x}, n-1}
\end{array}\right]+\left[\begin{array}{ccc}
b_{11} & \ldots & b_{1 d_{u}} \\
\vdots & \ddots & \vdots \\
b_{d_{x} 1} & \ldots & b_{d_{x} d_{u}}
\end{array}\right]\left[\begin{array}{c}
u_{1, n} \\
\vdots \\
u_{d_{u}, n}
\end{array}\right]
$$

More compactly, this can also be written as

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{F} \mathbf{x}_{n-1}+\mathbf{B}_{u} \mathbf{u}_{n} \tag{5.25}
\end{equation*}
$$

with
$\mathbf{x}_{n}=\left[\begin{array}{c}x_{1, n} \\ \vdots \\ x_{d_{x}, n}\end{array}\right], \mathbf{F}=\left[\begin{array}{ccc}a_{11} & \ldots & a_{1 d_{x}} \\ \vdots & \ddots & \vdots \\ a_{d_{x} 1} & \ldots & a_{d_{x} d_{x}}\end{array}\right], \mathbf{B}_{u}=\left[\begin{array}{ccc}b_{11} & \ldots & b_{1 d_{u}} \\ \vdots & \ddots & \vdots \\ b_{d_{x} 1} & \ldots & b_{d_{x} d_{u}}\end{array}\right], \mathbf{u}_{n}=\left[\begin{array}{c}u_{1, n} \\ \vdots \\ u_{d_{u}, n}\end{array}\right]$.
Equation (5.25) is the general dynamic model for linear discrete-time systems with deterministic inputs $\mathbf{u}_{n}$.

To convert the $L$ th order difference equation (with single input $u_{n}$ )

$$
\begin{equation*}
z_{n}=c_{1} z_{n-1}+c_{2} z_{n-2}+\cdots+c_{L} z_{n-L}+d_{1} u_{n} \tag{5.26}
\end{equation*}
$$

to the form (5.25), we proceed as follows. First, we have to choose the state variables $\mathbf{x}_{n}$. This is easiest done at time $n-1$ and for the model (5.26), we can choose

$$
x_{1, n-1}=z_{n-1}, x_{2, n-1}=z_{n-2}, \ldots, x_{d_{x}, n-1}=z_{n-L}
$$

Then we obtain that

$$
\begin{aligned}
x_{1, n} & =z_{n}=c_{1} z_{n-1}+c_{2} z_{n-2}+\cdots+c_{L} z_{n-L}+d_{1} u_{n} \\
& =c_{1} x_{1, n-1}+c_{2} x_{2, n-1}+\cdots+c_{L} x_{d_{x}, n-1}+d_{1} u_{n} \\
x_{2, n} & =z_{n-1} \\
& =x_{1, n-1} \\
& \vdots \\
x_{d_{x}, n} & =z_{n-L+1} \\
& =x_{d_{x}-1, n-1}
\end{aligned}
$$

where we have expressed the state at time $n$ solely in terms of the state at previous times as well as the input $u_{n}$. Rewriting this on matrix form finally yields

$$
\left[\begin{array}{c}
x_{1, n} \\
x_{2, n} \\
\vdots \\
x_{d_{x}, n}
\end{array}\right]=\left[\begin{array}{cccc}
c_{1} & c_{2} & \ldots & c_{L} \\
1 & 0 & & \vdots \\
\vdots & \ddots & & \\
0 & \ldots & 1 & 0
\end{array}\right]\left[\begin{array}{c}
x_{1, n-1} \\
x_{2, n-1} \\
\vdots \\
x_{d_{x}, n-1}
\end{array}\right]+\left[\begin{array}{c}
d_{1} \\
0 \\
\vdots \\
0
\end{array}\right] u_{n}
$$

which is of the form (5.25).
Equation (5.25) only models the dynamic behavior of the state. As for the continuous case, only noisy measurements of the dynamic process are available for estimation. Combining the dynamic model with a measurement model thus again yields the complete discrete-time state-space model. Assuming a linear measurement model of the form (3.10) then yields the linear discrete-time model

$$
\begin{align*}
\mathbf{x}_{n} & =\mathbf{F} \mathbf{x}_{n-1}+\mathbf{B}_{u} \mathbf{u}_{n}  \tag{5.27a}\\
\mathbf{y}_{n} & =\mathbf{G} \mathbf{x}_{n}+\mathbf{r}_{n} \tag{5.27b}
\end{align*}
$$

### 5.2.2 Stochastic Linear Dynamic Models

As for the continuous case, the deterministic discrete-time dynamic model can not take into account random effects, uncertainty, and unknown inputs. To account for such effects, we need an approach that is similar to the random process input
in the continuous case. For the discrete-time case, we can model this using a random variable as the input to the dynamic model (compared to a stochastic process for the continuous case). We denote this random input, commonly called the process noise, as $\mathbf{q}_{n}$. Then, replacing the deterministic input $\mathbf{u}_{n}$ with $\mathbf{q}_{n}$ yields the stochastic discrete-time dynamic model

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{F} \mathbf{x}_{n-1}+\mathbf{B}_{q} \mathbf{q}_{n} \tag{5.28}
\end{equation*}
$$

The random variable $\mathbf{q}_{n}$ follows some probability density function such that

$$
\mathbf{q}_{n} \sim p\left(\mathbf{q}_{n}\right) .
$$

Here, we assume that $\mathbf{q}_{n}$ is zero-mean, that is $\mathrm{E}\left\{\mathbf{q}_{n}\right\}=0$, and that it has covariance $\operatorname{Cov}\left\{\mathbf{q}_{n}\right\}=\mathbf{Q}_{n}$. Furthermore, we also assume that the cross-covariance between two random inputs $\mathbf{q}_{m}$ and $\mathbf{q}_{n}$ is $\operatorname{Cov}\left\{\mathbf{q}_{m}, \mathbf{q}_{n}\right\}=0$ for $m \neq n$.

Together with a linear sensor model, we then obtain the stochastic linear discrete-time state-space model given by

$$
\begin{align*}
\mathbf{x}_{n} & =\mathbf{F} \mathbf{x}_{n-1}+\mathbf{B}_{q} \mathbf{q}_{n}  \tag{5.29a}\\
\mathbf{y}_{n} & =\mathbf{G} \mathbf{x}_{n}+\mathbf{r}_{n} \tag{5.29b}
\end{align*}
$$

### 5.2.3 Nonlinear Dynamic Model

Similar to differential equations, difference equations may be nonlinear. In this case, we start from the nonlinear equation system with the random variables $q_{1, n}, q_{2, n}, \ldots, q_{d_{q}, n}$ as the inputs given by

$$
\begin{aligned}
& x_{1, n}=f_{1}\left(x_{1, n-1}, x_{2, n-1}, \ldots, x_{d_{x}, n-1}\right)+b_{11} q_{1, n}+\cdots+b_{1 d_{q}} q_{d_{q}, n} \\
& x_{2, n}=f_{2}\left(x_{1, n-1}, x_{2, n-1}, \ldots, x_{d_{x}, n-1}\right)+b_{21} q_{1, n}+\cdots+b_{2 d_{q}} q_{d_{q}, n} \\
& \quad \vdots \\
& x_{d_{x}, n}=f_{d_{x}}\left(x_{1, n-1}, x_{2, n-1}, \ldots, x_{d_{x}, n-1}\right)+b_{d_{x} 1} q_{1, n}+\cdots+b_{d_{x} d_{q}} q_{d_{q}, n} .
\end{aligned}
$$

This can directly be rewritten into vector form, which yields

$$
\left[\begin{array}{c}
x_{1, n} \\
\vdots \\
x_{d_{x}, n}
\end{array}\right]=\left[\begin{array}{c}
f_{1}\left(x_{1, n-1}, x_{2, n-1}, \ldots, x_{d_{x}, n-1}\right) \\
\vdots \\
f_{d_{x}}\left(x_{1, n-1}, x_{2, n-1}, \ldots, x_{d_{x}, n-1}\right)
\end{array}\right]+\left[\begin{array}{ccc}
b_{11} & \ldots & b_{1 d_{q}} \\
\vdots & \ddots & \vdots \\
b_{d_{x} 1} & \cdots & b_{d_{x} d_{q}}
\end{array}\right]\left[\begin{array}{c}
q_{1, n} \\
\vdots \\
q_{d_{q}, n}
\end{array}\right],
$$

or more compactly

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{B}_{q} \mathbf{q}_{n} \tag{5.30}
\end{equation*}
$$

The nonlinear dynamic model (5.30) together with the general nonlinear measurement model (4.1) yields the nonlinear discrete-time state-space model

$$
\begin{align*}
\mathbf{x}_{n} & =\mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{B}_{q} \mathbf{q}_{n},  \tag{5.31a}\\
\mathbf{y}_{n} & =\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n}, \tag{5.31b}
\end{align*}
$$

where $\mathbf{q}_{n} \sim p\left(\mathbf{q}_{n}\right), \mathrm{E}\left\{\mathbf{q}_{n}\right\}=0, \operatorname{Cov}\left\{\mathbf{q}_{n}\right\}=\mathbf{Q}_{n}$, and $\mathbf{r}_{n} \sim p\left(\mathbf{r}_{n}\right), \mathrm{E}\left\{\mathbf{r}_{n}\right\}=0$, $\operatorname{Cov}\left\{\mathbf{r}_{n}\right\}=\mathbf{R}_{n}$.

### 5.3 Discretization of Linear Dynamic Models

In practice, modern sensor fusion systems are implemented in a digital computer. Hence, dealing with continuous dynamic models, that are defined on continuous $t$ is not possible. Instead, we have to discretize the continuous-time model to obtain an (approximately) equivalent discrete-time dynamic model. In other words, we have to solve the differential equation on the interval between times $t_{n-1}$ and $t_{n}$, that is, we have to integrate the differential equation. For linear dynamic models, this can be achieved exactly and we develop the necessary tools in this section. For most nonlinear models, exact discretization is not feasible and we have to resort to approximate methods. This is the subject of Section 5.4.

### 5.3.1 Deterministic Linear Dynamic Models

Scalar Model. To derive the necessary expressions for discretizing the linear dynamic model, we first review how to solve a first order scalar case. Consider the non-homogeneous first order ODE

$$
\begin{equation*}
\dot{x}(t)=a x(t)+b u(t) \tag{5.32}
\end{equation*}
$$

which we want to solve on the interval between the two sampling points $t_{n}$ and $t_{n-1}$. To do this, we introduce the integrating factor $e^{-a t}$ and note that it holds that

$$
\begin{equation*}
\frac{\mathrm{d}}{\mathrm{~d} t} e^{-a t} x(t)=e^{-a t} \dot{x}(t)-e^{-a t} a x(t) \tag{5.33}
\end{equation*}
$$

Rearranging (5.32) by moving the term $a x(t)$ to the left hand side and multiplying by the integrating factor yields

$$
e^{-a t} \dot{x}(t)-e^{-a t} a x(t)=e^{-a t} b u(t)
$$

and thus, using (5.33), we have that

$$
\begin{equation*}
\frac{\mathrm{d}}{\mathrm{~d} t} e^{-a t} x(t)=e^{-a t} b u(t) \tag{5.34}
\end{equation*}
$$

Equation (5.34) is separable in $t$. Hence, the ODE can directly be integrated from $t_{n-1}$ to $t_{n}$ according to

$$
\int_{t_{n-1}}^{t_{n}} \mathrm{~d}\left[e^{-a t} x(t)\right]=\int_{t_{n-1}}^{t_{n}} e^{-a t} b u(t) \mathrm{d} t
$$

which yields

$$
\begin{aligned}
{\left[e^{-a t} x(t)\right]_{t=t_{n-1}}^{t_{n}} } & =\int_{t_{n-1}}^{t_{n}} e^{-a t} b u(t) \mathrm{d} t \\
e^{-a t_{n}} x\left(t_{n}\right)-e^{-a t_{n-1}} x\left(t_{n-1}\right) & =\int_{t_{n-1}}^{t_{n}} e^{-a t} b u(t) \mathrm{d} t
\end{aligned}
$$

Now we can solve for $x\left(t_{n}\right)$ by rearranging and multiplying by the factor $e^{a t_{n}}$, which yields

$$
\begin{aligned}
x\left(t_{n}\right) & =e^{a t_{n}-a t_{n-1}} x\left(t_{n-1}\right)+e^{a t_{n}} \int_{t_{n-1}}^{t_{n}} e^{-a t} b u(t) \mathrm{d} t, \\
& =e^{a\left(t_{n}-t_{n-1}\right)} x\left(t_{n-1}\right)+\int_{t_{n-1}}^{t_{n}} e^{a\left(t_{n}-t\right)} b u(t) \mathrm{d} t .
\end{aligned}
$$

Finally, letting $\Delta t \triangleq t_{n}-t_{n-1}$ and using the notation $x_{n} \triangleq x\left(t_{n}\right)$, we obtain

$$
\begin{equation*}
x_{n}=e^{a \Delta t} x_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{a\left(t_{n}-t\right)} b u(t) \mathrm{d} t . \tag{5.35}
\end{equation*}
$$

Note that the second term on the right hand side of (5.35) is the convolution between the factor $e^{a\left(t_{n}-t\right)}$ and the input $u(t)$ on the interval between $t_{n-1}$ and $t_{n}$.

General Case. Based on the steps used for solving the scalar ODE (5.32), the general case can now be solved. First, recall that the general dynamic model on state-space form was given by the vector valued first order ODE system (5.15), that is,

$$
\begin{equation*}
\dot{\mathbf{x}}(t)=\mathbf{A} \mathbf{x}(t)+\mathbf{B}_{u} \mathbf{u}(t) . \tag{5.36}
\end{equation*}
$$

The key here is to note that (5.36) is a first order vector ODE just like (5.32). Hence, we should be able to solve it using an integrating factor similar to the one for the scalar case. Indeed, such an integrating factor can be found through the matrix exponential. The matrix exponential for a square matrix A can be defined in the same way as the regular exponential, that is, as an infinite sum of the form

$$
e^{\mathbf{A}}=\sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{A}^{k} .
$$

Similar to the (scalar) exponential, it holds that the derivative with respect to a scalar $x$ of $e^{\mathbf{A} x}$ is

$$
\begin{equation*}
\frac{\mathrm{d}}{\mathrm{~d} x} e^{\mathbf{A} x}=e^{\mathbf{A} x} \mathbf{A} \tag{5.37}
\end{equation*}
$$

Another useful property of the matrix exponential is that of its transpose, which is given by

$$
\begin{equation*}
\left(e^{\mathbf{A}}\right)^{\top}=e^{\mathbf{A}^{\top}} . \tag{5.38}
\end{equation*}
$$

With this in mind, we can solve the general case following the same steps as for the scalar case. First, we multiply (5.36) by the integrating factor $e^{-\mathbf{A t}}$ and rearrange it to obtain

$$
e^{-\mathbf{A} t} \dot{\mathbf{x}}(t)-e^{-\mathbf{A} t} \mathbf{A} \mathbf{x}(t)=e^{-\mathbf{A} t} \mathbf{B}_{u} \mathbf{u}(t)
$$

Next, similar to (5.33) (product rule), the left hand side is

$$
\frac{\mathrm{d}}{\mathrm{~d} t} e^{-\mathbf{A} t} \mathbf{x}(t)=e^{-\mathbf{A} t} \dot{\mathbf{x}}(t)-e^{-\mathbf{A} t} \mathbf{A} \mathbf{x}(t)
$$

and thus, we have that

$$
\frac{\mathrm{d}}{\mathrm{~d} t} e^{-\mathbf{A} t} \mathbf{x}(t)=e^{-\mathbf{A} t} \mathbf{B}_{u} \mathbf{u}(t)
$$

This is again separable in $t$. Direct integration from $t_{n-1}$ to $t_{n}$ then leads to

$$
\begin{aligned}
\int_{t_{n-1}}^{t_{n}} \mathrm{~d}\left[e^{-\mathbf{A} t} \mathbf{x}(t)\right] & =\int_{t_{n-1}}^{t_{n}} e^{-\mathbf{A} t} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t \\
{\left[e^{-\mathbf{A} t} \mathbf{x}(t)\right]_{t=t_{n-1}}^{t_{n}} } & =\int_{t_{n-1}}^{t_{n}} e^{-\mathbf{A} t} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t \\
e^{-\mathbf{A} t_{n}} \mathbf{x}\left(t_{n}\right)-e^{-\mathbf{A} t_{n-1}} \mathbf{x}\left(t_{n-1}\right) & =\int_{t_{n-1}}^{t_{n}} e^{-\mathbf{A} t} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t
\end{aligned}
$$

Solving for $\mathbf{x}\left(t_{n}\right)$ by rearranging and multiplying both sides by $e^{\mathbf{A} t_{n}}$ yields

$$
\mathbf{x}\left(t_{n}\right)=e^{\mathbf{A}\left(t_{n}-t_{n-1}\right)} \mathbf{x}\left(t_{n-1}\right)+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t
$$

or

$$
\begin{equation*}
\mathbf{x}_{n}=e^{\mathbf{A} \Delta t} \mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t \tag{5.39}
\end{equation*}
$$

In this case, the factor in front of $x_{n-1}$ (i.e., the state at time $t_{n-1}$ ) is a matrix exponential that again depends on the time difference $\Delta t$ between the two discrete times $t_{n-1}$ and $t_{n}$. Note that $\Delta t$ can be arbitrary, meaning that there is no need for uniform sampling between points $t_{n-2}, t_{n-1}, t_{n}$, and so forth. Furthermore, the second term in (5.39) is again the convolution between the deterministic input $u(t)$ and the matrix exponential on the same interval.

If the input $\mathbf{u}(t)$ is a zeroth-order-hold ( ZOH ) signal, which is common in control or robot navigation, the second term can further be simplified. In that case, $\mathbf{u}(t)$ is constant on the interval between $t_{n-1}$ and $t_{n}$ and we can denote it as $\mathbf{u}_{n-1}$. Thus, we then have that

$$
\begin{aligned}
\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t & =\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}_{n-1} \mathrm{~d} t \\
& =\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathrm{~d} t \mathbf{u}_{n-1}
\end{aligned}
$$

where $\mathbf{B}_{u}$ may or may not depend on $t$.

Defining

$$
\begin{align*}
& \mathbf{F}_{n} \triangleq e^{\mathbf{A}\left(t_{n}-t_{n-1}\right)}  \tag{5.40a}\\
& \mathbf{L}_{n} \triangleq \int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathrm{~d} t \tag{5.40b}
\end{align*}
$$

we can rewrite (5.39) as

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{F}_{n} \mathbf{x}_{n-1}+\mathbf{L}_{n} \mathbf{u}_{n-1} \tag{5.41}
\end{equation*}
$$

Equation (5.41) is an exact solution to (5.36) and thus it is an exact discretization and equivalent representation of the continuous-time model. Furthermore, it also has the same form as the discrete-time state-space model (5.27).

### 5.3.2 Stochastic Linear Dynamic Models

Discretization of the stochastic linear dynamic model in (5.19) and given by

$$
\begin{equation*}
\dot{\mathbf{x}}(t)=\mathbf{A} \mathbf{x}(t)+\mathbf{B}_{w} \mathbf{w}(t) \tag{5.42}
\end{equation*}
$$

with $R_{w w}(\tau)=\mathrm{E}\{\mathbf{w}(t+\tau) \mathbf{w}(t)\}=\boldsymbol{\Sigma}_{w} \delta(\tau)$, follows the same steps as the discretization of the deterministic model. This yields

$$
\mathbf{x}\left(t_{n}\right)=e^{\mathbf{A}\left(t_{n}-t_{n-1}\right)} \mathbf{x}\left(t_{n-1}\right)+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t
$$

Now note that special attention has to be paid to the second term on the right hand side. The noise process $\mathbf{w}(t)$ does not have the properties required for regular integration rules to apply (Øksendal, 2010; Särkkä and Solin, 2019).

However, we can define a random variable $\mathbf{q}_{n}$ as

$$
\mathbf{q}_{n} \triangleq \int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t
$$

Next, we can calculate the properties of $\mathbf{q}_{n}$. The mean is given by

$$
\begin{aligned}
\mathrm{E}\left\{\mathbf{q}_{n}\right\} & =\mathrm{E}\left\{\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t\right\} \\
& =\int_{t_{n-1}}^{t_{n}} \mathrm{E}\left\{e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathbf{w}(t)\right\} \mathrm{d} t \\
& =\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathrm{E}\{\mathbf{w}(t)\} \mathrm{d} t \\
& =0
\end{aligned}
$$

Similarly, we can calculate the covariance matrix

$$
\begin{aligned}
\operatorname{Cov}\left\{\mathbf{q}_{n}\right\} & =\mathrm{E}\left\{\left(\mathbf{q}_{n}-\mathrm{E}\left\{\mathbf{q}_{n}\right\}\right)\left(\mathbf{q}_{n}-\mathrm{E}\left\{\mathbf{q}_{n}\right\}\right)^{\top}\right\} \\
& =\mathrm{E}\left\{\mathbf{q}_{n} \mathbf{q}_{n}^{\mathrm{T}}\right\} \\
& =\mathrm{E}\left\{\left(\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t\right)\left(\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-\tau\right)} \mathbf{B}_{w} \mathbf{w}(\tau) \mathrm{d} \tau\right)^{\top}\right\} \\
& =\int_{t_{n-1}}^{t_{n}} \int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \mathrm{E}\left\{\mathbf{w}(t) \mathbf{w}(\tau)^{\top}\right\} \mathbf{B}_{w}^{\top}\left(e^{\mathbf{A}\left(t_{n}-\tau\right)}\right)^{\top} \mathrm{d} \tau \mathrm{~d} t \\
& =\int_{t_{n-1}}^{t_{n}} \int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} R_{w w}(t-\tau) \mathbf{B}_{w}^{\top}\left(e^{\mathbf{A}\left(t_{n}-\tau\right)}\right)^{\top} \mathrm{d} \tau \mathrm{~d} t \\
& =\int_{t_{n-1}}^{t_{n}} \int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-t\right)} \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \delta(t-\tau) \mathbf{B}_{w}^{\top}\left(e^{\mathbf{A}\left(t_{n}-\tau\right)}\right)^{\top} \mathrm{d} \tau \mathrm{~d} t \\
& =\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-\tau\right)} \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top} e^{\mathbf{A}^{\top}\left(t_{n}-\tau\right)} \mathrm{d} \tau \\
& \triangleq \mathbf{Q}_{n} .
\end{aligned}
$$

Hence, the mean and covariance of the random variable $\mathbf{q}_{n}$ given that $\mathbf{w}(t)$ is a zero-mean white noise process with autocorrelation function $R_{w w}(\tau)=\boldsymbol{\Sigma}_{w} \delta(\tau)$ are given by

$$
\begin{align*}
\mathrm{E}\left\{\mathbf{q}_{n}\right\} & =0  \tag{5.43a}\\
\operatorname{Cov}\left\{\mathbf{q}_{n}\right\} & =\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}\left(t_{n}-\tau\right)} \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top} e^{\mathbf{A}^{\top}\left(t_{n}-\tau\right)} \mathrm{d} \tau \triangleq \mathbf{Q}_{n} \tag{5.43b}
\end{align*}
$$

Furthermore, it turns out that in this case, the process noise $\mathbf{q}_{n}$ is actually a Gaussian random variable, that is, it holds that

$$
\begin{equation*}
\mathbf{q}_{n} \sim \mathcal{N}\left(0, \mathbf{Q}_{n}\right) \tag{5.44}
\end{equation*}
$$

In summary, the discretized stochastic dynamic model (5.42) is

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{F}_{n} \mathbf{x}_{n-1}+\mathbf{q}_{n} \tag{5.45}
\end{equation*}
$$

where

$$
\begin{align*}
\mathbf{F}_{n} & \triangleq e^{\mathbf{A}\left(t_{n}-t_{n-1}\right)}  \tag{5.46a}\\
\mathbf{q}_{n} & \sim \mathcal{N}\left(0, \mathbf{Q}_{n}\right) \tag{5.46b}
\end{align*}
$$

with $\mathbf{Q}_{n}$ as in (5.43). Again, the discretized model (5.45)-(5.46) is an exact discretization of the continuous-time model (5.42) and thus completely equivalent.

### 5.4 Discretization of Nonlinear Dynamic Models

Unlike the linear models discussed in Section 5.3, discretization of nonlinear dynamic models is generally harder. In particular, closed-form solutions to the discretization problem can only be found in a few special cases. In most other cases, approximations have to be used. In this section three approaches for approximate discretization are discussed.

### 5.4.1 Discretization of Linearized Nonlinear Models

The first approach is based on a Taylor series expansion of the nonlinear function. Truncating the Taylor series expansion of the nonlinear function $\mathbf{f}(\mathbf{x}(t))$ around the state $\mathbf{x}_{n-1}$ at the linear term yields the following approximation of the (deterministic) nonlinear dynamic model:

$$
\dot{\mathbf{x}}(t) \approx \mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{A}_{x}\left(\mathbf{x}(t)-\mathbf{x}_{n-1}\right)+\mathbf{B}_{u} \mathbf{u}(t)
$$

where $\mathbf{A}_{x}$ is the Jacobian of $\mathbf{f}(\mathbf{x}(t))$ with respect to $\mathbf{x}(t)$, evaluated at $\mathbf{x}(t)=\mathbf{x}_{n-1}$. Rearranging the terms on the right hand side yields

$$
\dot{\mathbf{x}}(t) \approx \mathbf{A}_{x} \mathbf{x}(t)+\mathbf{f}\left(\mathbf{x}_{n-1}\right)-\mathbf{A}_{x} \mathbf{x}_{n-1}+\mathbf{B}_{u} \mathbf{u}(t)
$$

which is linear in $\mathbf{x}(t)$, and $\mathbf{f}\left(\mathbf{x}_{n-1}\right)$ and $\mathbf{A}_{x} \mathbf{x}_{n-1}$ are constants. Thus, we can use the same approach as for linear models in Section 5.3, that is, we can use the integrating factor $e^{-\mathbf{A}_{x} t}$. This yields

$$
\begin{aligned}
& \mathbf{x}_{n} \approx e^{\mathbf{A}_{x} \Delta t} \mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathrm{d} t f\left(\mathbf{x}_{n-1}\right)-\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathrm{d} t \mathbf{A}_{x} \mathbf{x}_{n-1} \\
&+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t
\end{aligned}
$$

Noting that the third term is the integral of the derivative of the matrix exponential with respect to $t$ (see (5.37)), we have that

$$
\begin{aligned}
\mathbf{x}_{n} \approx & e^{\mathbf{A}_{x} \Delta t} \mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathrm{d} t f\left(\mathbf{x}_{n-1}\right)+\left[e^{\mathbf{A}_{x}\left(t_{n}-t\right)}\right]_{t=t_{n-1}}^{t_{n}} \mathbf{x}_{n-1} \\
& \quad+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t \\
= & e^{\mathbf{A}_{x} \Delta t} \mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathrm{d} t f\left(\mathbf{x}_{n-1}\right)+\left(\mathbf{I}-e^{\mathbf{A}_{x}\left(t_{n}-t_{n-1}\right)}\right) \mathbf{x}_{n-1} \\
& \quad+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t,
\end{aligned}
$$

and finally,

$$
\begin{equation*}
\mathbf{x}_{n} \approx \mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathrm{d} t f\left(\mathbf{x}_{n-1}\right)+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t \tag{5.47}
\end{equation*}
$$

As usual, the drawback of this Taylor-series-based solution is that the linearization is only local and large $\Delta t$ s as well as highly nonlinear functions may be problematic. On the other hand, as it can be seen from (5.47), the approximation only affects the deterministic term of the model, whereas the input is calculated in the same way as for the linear model (but with the Jacobian $\mathbf{A}_{x}$ ). Hence, for the stochastic dynamic model, the result immediately follows from (5.47) to be

$$
\begin{equation*}
\mathbf{x}_{n} \approx \mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-t\right)} \mathrm{d} t \mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{q}_{n} \tag{5.48}
\end{equation*}
$$

with

$$
\begin{aligned}
\mathbf{q}_{n} & \sim \mathcal{N}\left(0, \mathbf{Q}_{n}\right) \\
\mathbf{Q}_{n} & \approx \int_{t_{n-1}}^{t_{n}} e^{\mathbf{A}_{x}\left(t_{n}-\tau\right)} \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top} e^{\mathbf{A}_{x}^{\top}\left(t_{n}-\tau\right)} \mathrm{d} \tau
\end{aligned}
$$

### 5.4.2 Euler Discretization of Deterministic Nonlinear Models

The second approach is based on the so-called Euler approximation. We know that the solution for $\mathbf{x}_{n}$ can be written as

$$
\mathbf{x}_{n}=\mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} \mathbf{f}(\mathbf{x}(t)) \mathrm{d} t+\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{u} \mathbf{u}(t) \mathrm{d} t
$$

and the problem is to calculate the integrals in the second and third terms on the right hand side. However, for sufficiently small $\Delta t=t_{n}-t_{n-1}$, we can approximate the integral by using the rectangle approximation. This means that we can approximate the integral as

$$
\begin{aligned}
\int_{t_{n-1}}^{t_{n}} \mathbf{f}(\mathbf{x}(t)) \mathrm{d} t & \approx \mathbf{f}\left(\mathbf{x}_{n-1}\right)\left(t_{n}-t_{n-1}\right) \\
& =\Delta t \mathbf{f}\left(\mathbf{x}_{n-1}\right)
\end{aligned}
$$

and similar for the second integral with the input $\mathbf{u}(t)$.
This then yields the Euler discretization of the deterministic, nonlinear dynamic model given by

$$
\begin{equation*}
\mathbf{x}_{n} \approx \mathbf{x}_{n-1}+\Delta t \mathbf{f}\left(\mathbf{x}_{n-1}\right)+\Delta t \mathbf{B}_{u} \mathbf{u}_{n-1} \tag{5.49}
\end{equation*}
$$

### 5.4.3 Euler-Maruyama Discretization of Stochastic Nonlinear Models

The Euler discretization is quite simple to implement. However, it does not readily work for stochastic dynamic models. In this case, we are interested in solving the integral equation

$$
\mathbf{x}_{n}=\mathbf{x}_{n-1}+\int_{t_{n-1}}^{t_{n}} \mathbf{f}(\mathbf{x}(t)) \mathrm{d} t+\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t
$$

We can still use the rectangle approximation for the integral involving the deterministic part (second term on the right hand side), but for the stochastic part, we are again lacking the integration rules. Thus, we can not use the rectangle approximation in this case. However, we can again define the resulting random variable as

$$
\mathbf{q}_{n} \triangleq \int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t
$$

and investigate its properties instead.
First, the mean is given by

$$
\begin{aligned}
\mathrm{E}\left\{\mathbf{q}_{n}\right\} & =\mathrm{E}\left\{\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t\right\} \\
& =\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathrm{E}\{\mathbf{w}(t)\} \mathrm{d} t \\
& =0
\end{aligned}
$$

where we have made use of the assumption that $\mathbf{w}(t)$ is a zero-mean process. Second, the covariance can be found similar to the linear case as follows.

$$
\begin{aligned}
\operatorname{Cov}\left\{\mathbf{q}_{n}\right\} & =\mathrm{E}\left\{\left(\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathbf{w}(t) \mathrm{d} t\right)\left(\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathbf{w}(\tau) \mathrm{d} \tau\right)^{\top}\right\} \\
& =\int_{t_{n-1}}^{t_{n}} \int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \mathrm{E}\left\{\mathbf{w}(t) \mathbf{w}(\tau)^{\top}\right\} \mathbf{B}_{w}^{\top} \mathrm{d} \tau \mathrm{~d} t \\
& =\int_{t_{n-1}}^{t_{n}} \int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \delta(t-\tau) \mathbf{B}_{w}^{\top} \mathrm{d} \tau \mathrm{~d} t \\
& =\int_{t_{n-1}}^{t_{n}} \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top} \mathrm{d} \tau .
\end{aligned}
$$

At this point, note that the remaining integral does not involve any random process anymore. Hence, we can again use the rectangle approximation of the integral, which yields

$$
\begin{aligned}
\operatorname{Cov}\left\{\mathbf{q}_{n}\right\} & \approx\left(\mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top}\right)\left(t_{n}-t_{n-1}\right) \\
& =\Delta t \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top} \\
& \triangleq \mathbf{Q}_{n}
\end{aligned}
$$

This yields the Euler-Maruyama discretization of the stochastic nonlinear dynamic model given by

$$
\begin{equation*}
\mathbf{x}_{n}=\mathbf{x}_{n-1}+\Delta t \mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{q}_{n} \tag{5.50}
\end{equation*}
$$

with

$$
\begin{aligned}
\mathbf{q}_{n} & \sim \mathcal{N}\left(0, \mathbf{Q}_{n}\right) \\
\mathbf{Q}_{n} & =\Delta t \mathbf{B}_{w} \boldsymbol{\Sigma}_{w} \mathbf{B}_{w}^{\top}
\end{aligned}
$$

## Chapter 6

## Filtering

In Chapter 5, we developed state-space models that describe the dynamic behavior of time-varying processes and how the sensor measurements relate to the state. Unfortunately, the estimation methods developed earlier do not work for such dynamic problems: These solve the estimation problem in the static case, but not when the parameters (i.e., the states) vary between samples. Instead, new methods are required which are able to combine the prior information from the dynamic model and the measurements at different points in time. This methodology is called filtering and the general approach is discussed in Section 6.1, which is followed by the exact solution of the filtering problem for linear state-space models in Section 6.2. Approximate filtering methods for nonlinear systems are discussed in Section 6.36.6.

### 6.1 The Filtering Approach

The basic discussion of the filtering approach is based on the general discrete-time state-space model of the form

$$
\begin{align*}
& \mathbf{x}_{n}=\mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{q}_{n}  \tag{6.1a}\\
& \mathbf{y}_{n}=\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n} \tag{6.1b}
\end{align*}
$$

where the process and measurement noises are zero-mean $\left(\mathrm{E}\left\{\mathbf{q}_{n}\right\}=\mathrm{E}\left\{\mathbf{r}_{n}\right\}=0\right)$ with covariances $\operatorname{Cov}\left\{\mathbf{q}_{n}\right\}=\mathbf{Q}_{n}$ and $\operatorname{Cov}\left\{\mathbf{r}_{n}\right\}=\mathbf{R}_{n}$, respectively. Note that the dynamic model (6.1a) may be an inherently discrete-time model (Section 5.2) or the result of discretizing a continuous-time dynamic model.

An aspect of the state-space model that has been neglected so far are the initial conditions. Since the dynamic model (6.1a) is based on a differential (difference) equation, it also requires knowledge of the initial conditions of the system. Naturally, in the sensor fusion and estimation context, the initial conditions are unknown (e.g., it is not known where exactly a target is located in the beginning). Instead, it is suitable to specify the initial conditions in a probabilistic way. In other words,
we assume that the state at $t_{0}$, denoted as $\mathbf{x}_{0} \triangleq \mathbf{x}\left(t_{0}\right)$, is a random variable which follows a probability density function $p\left(\mathbf{x}_{0}\right)$, that is,

$$
\mathbf{x}_{0} \sim p\left(\mathbf{x}_{0}\right)
$$

This allows us to specify the prior knowledge about where the initial state may be but also take into account that this is an uncertain guess. In practice, the distribution of the initial state is often assumed to follow a Gaussian distribution, but this does not need to be the case. Here, we only make the assumptions that the mean and covariance of the initial state are given by

$$
\begin{align*}
\mathrm{E}\left\{\mathbf{x}_{0}\right\} & =\mathbf{m}_{0},  \tag{6.2a}\\
\operatorname{Cov}\left\{\mathbf{x}_{0}\right\} & =\mathbf{P}_{0} . \tag{6.2b}
\end{align*}
$$

Given the state-space model defined by (6.1) and (6.2), the filtering approach can now be formulated. In particular, filtering iterates between 1) a prediction step, which predicts the current state using the dynamic model and the previous estimate of the state (also called time update), and 2) a measurement update step that estimates the current state using the prediction and the new measurement. This prediction-update strategy is actually very general, and as it will be shown, all the algorithms discussed in this chapter make use of these two steps.

Prediction. The aim of the prediction step is to estimate the current state $\mathbf{x}_{n}$ at $t_{n}$ given the set of all the previous measurement data $\mathbf{y}_{1: n-1}=\left\{\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{n-1}\right\}$, and hence, given the estimate of the state $\mathbf{x}_{n-1}$ at $t_{n-1}$. Here, the mean of the prediction is denoted as $\hat{\mathbf{x}}_{n \mid n-1}$, which is read as "the estimate of the state at $t_{n}$ given the data up to $t_{n-1}$ ". The hat indicates that it is an estimation, whereas the subscript is closely related to the notation for conditional expectations and densities, that is, it is read as " $n$ given $n-1$ ". Similarly, the covariance of the predicted $\mathbf{x}_{n}$ is denoted as $\mathbf{P}_{n \mid n-1}$.

Measurement Update. In the measurement update, the newly obtained measurement $\mathbf{y}_{n}$ is used together with the prediction (and its uncertainty) to estimate the current value of the state $\mathbf{x}_{n}$. Similar to the prediction, the mean of the updated state estimate at time $t_{n}$ is denoted as $\hat{\mathbf{x}}_{n \mid n}$, which is read as "the estimate of the state at $t_{n}$ given the data up to $t_{n} "$. In other words, this estimate is now based on the set of all measurements including the latest measurement at $t_{n}$ given by $\mathbf{y}_{1: n}=\left\{\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{n}\right\}$. Again, the covariance of the updated estimate is denoted as $\mathbf{P}_{n \mid n}$.

### 6.2 Kalman Filter

In this section, filtering for linear state-space models is considered. Recall that the linear state-space model with continuous-time dynamic model is (see Section 5.1)

$$
\begin{align*}
\dot{\mathbf{x}}(t) & =\mathbf{A} \mathbf{x}(t)+\mathbf{B} \mathbf{w}(t),  \tag{6.3a}\\
\mathbf{y}_{n} & =\mathbf{G}_{n} \mathbf{x}\left(t_{n}\right)+\mathbf{r}_{n} . \tag{6.3b}
\end{align*}
$$

Together with the initial distribution (6.2), the model (6.3) has the discrete-time equivalent

$$
\begin{align*}
& \mathbf{x}_{n}=\mathbf{F}_{n} \mathbf{x}_{n-1}+\mathbf{q}_{n}  \tag{6.4a}\\
& \mathbf{y}_{n}=\mathbf{G}_{n} \mathbf{x}_{n}+\mathbf{r}_{n} \tag{6.4b}
\end{align*}
$$

with

$$
\begin{align*}
\mathrm{E}\left\{\mathbf{x}_{0}\right\} & =\mathbf{m}_{0}, \operatorname{Cov}\left\{\mathbf{x}_{0}\right\}=\mathbf{P}_{0}  \tag{6.5a}\\
\mathrm{E}\left\{\mathbf{q}_{n}\right\} & =0, \operatorname{Cov}\left\{\mathbf{q}_{n}\right\}=\mathbf{Q}_{n}  \tag{6.5b}\\
\mathrm{E}\left\{\mathbf{r}_{n}\right\} & =0, \operatorname{Cov}\left\{\mathbf{r}_{n}\right\}=\mathbf{R}_{n} \tag{6.5c}
\end{align*}
$$

Based on (6.4)-(6.5), the prediction and update steps can now be derived. For ease of presentation, the measurement update step is presented first, followed by the prediction.

Measurement Update. For the measurement update at time $t_{n}$, assume that there exists a prediction of the mean of the state $\hat{\mathbf{x}}_{n \mid n-1}$ and its covariance $\mathbf{P}_{n \mid n-1}$. This can be seen as the prior knowledge of the state at $t_{n}$. Then, the new measurement $\mathbf{y}_{n}$ provides the actual (noisy) information about the true state. Hence, the objective is to minimize the error with respect to the measurement, taking the prior information (prediction) into account.

Recall that we have based our estimators mainly on cost functions in general and quadratic cost functions (least squares) in particular. A cost function that is particularly well suited for this kind of problem is the regularized least squares cost function. Hence, for the filtering problem with the linear measurement model, the cost function becomes

$$
\begin{align*}
& J_{\operatorname{ReLS}}\left(\mathbf{x}_{n}\right)=\left(\mathbf{y}_{n}-\mathbf{G}_{n} \mathbf{x}_{n}\right)^{\top} \mathbf{R}_{n}^{-1}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \mathbf{x}_{n}\right) \\
&+\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mathbf{P}_{n \mid n-1}^{-1}\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right) \tag{6.6}
\end{align*}
$$

where the first term accounts for the measurement $\mathbf{y}_{n}$ and the second term, the regularization term, accounts for the prior information from the prediction.

To estimate the state, we can now solve the regularized least squares problem

$$
\begin{equation*}
\hat{\mathbf{x}}_{n \mid n}=\underset{\mathbf{x}_{n}}{\operatorname{argmin}} J_{\operatorname{ReLS}}\left(\mathbf{x}_{n}\right) \tag{6.7}
\end{equation*}
$$

Since the cost function is completely linear in the unknown $\mathbf{x}_{n}$, solving (6.7) follows the same steps as solving the linear least squares problems in Chapter 3 and we can find a closed-form solution.

Fortunately, we have already solved this problem for the static case. The solution is given by (see Section 3.3)

$$
\begin{equation*}
\hat{\mathbf{x}}_{n \mid n}=\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}\right) \tag{6.8}
\end{equation*}
$$

where $\mathbf{K}_{n}$ is called the Kalman gain that is given by

$$
\begin{equation*}
\mathbf{K}_{n}=\mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}\left(\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right)^{-1} \tag{6.9}
\end{equation*}
$$

Furthermore, we have also shown that the covariance of $\hat{\mathbf{x}}_{n \mid n}$ is

$$
\begin{equation*}
\mathbf{P}_{n \mid n}=\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\top} \tag{6.10}
\end{equation*}
$$

An important property of the measurement update step is that the updated estimate and its covariance actually are the mean and covariance of the state given the set of all measurements $\mathbf{y}_{1: n}=\left\{\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{n}\right\}$ (Särkkä, 2013). In other words, it holds that

$$
\begin{align*}
\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n}\right\} & =\hat{\mathbf{x}}_{n \mid n}  \tag{6.11a}\\
\operatorname{Cov}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n}\right\} & =\mathbf{P}_{n \mid n} \tag{6.11b}
\end{align*}
$$

where $\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n}\right\}$ denotes the conditional expectation of $\mathbf{x}_{n}$ given the data $\mathbf{y}_{1: n}$ (and similar for the covariance).

Prediction. Given the estimated mean $\mathrm{E}\left\{\mathbf{x}_{n-1} \mid \mathbf{y}_{1: n-1}\right\}=\hat{\mathbf{x}}_{n-1 \mid n-1}$ and its covariance $\operatorname{Cov}\left\{\mathbf{x}_{n-1} \mid \mathbf{y}_{1: n-1}\right\}=\mathbf{P}_{n-1 \mid n-1}$ from the previous time step $t_{n-1}$, the predicted mean and covariance can now be calculated. The mean is given by

$$
\hat{\mathbf{x}}_{n \mid n-1}=\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}
$$

Substituting $\mathbf{x}_{n}$ in the expectation using the dynamic model given in (6.4), the mean can be rewritten as

$$
\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}=\mathrm{E}\left\{\mathbf{F}_{n} \mathbf{x}_{n-1}+\mathbf{q}_{n} \mid \mathbf{y}_{1: n-1}\right\}
$$

Distributing the expectation over the sum and noting that $\mathrm{E}\left\{\mathbf{q}_{n} \mid \mathbf{y}_{1: n-1}\right\}=0$ yields

$$
\begin{align*}
\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\} & =\mathbf{F}_{n} \mathrm{E}\left\{\mathbf{x}_{n-1} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}, \tag{6.12}
\end{align*}
$$

where the last equality makes use of (6.11).
Similarly, the covariance is found from

$$
\begin{aligned}
\mathbf{P}_{n \mid n-1} & =\operatorname{Cov}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{x}_{n}-\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)\left(\mathbf{x}_{n}-\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} .
\end{aligned}
$$

```
Algorithm 6.1 Kalman Filter
    Initialize \(\hat{\mathbf{x}}_{0 \mid 0}=\mathbf{m}_{0}, \mathbf{P}_{0 \mid 0}=\mathbf{P}_{0}\)
    for \(n=1,2, \ldots\) do
        Prediction (time update):
\[
\begin{aligned}
\hat{\mathbf{x}}_{n \mid n-1} & =\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1} \\
\mathbf{P}_{n \mid n-1} & =\mathbf{F}_{n} \mathbf{P}_{n-1 \mid n-1} \mathbf{F}_{n}^{\top}+\mathbf{Q}_{n}
\end{aligned}
\]
```

4: Measurement update:

$$
\begin{aligned}
\mathbf{K}_{n} & =\mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}\left(\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right)^{-1} \\
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}\right) \\
\mathbf{P}_{n \mid n} & =\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\top}
\end{aligned}
$$

end for

Using the expression of the dynamic model (6.4) for $\mathbf{x}_{n}$ and (6.12) for $\mathrm{E}\left\{\mathbf{x}_{n} \mid\right.$ $\left.\mathrm{y}_{1: n-1}\right\}$ yields

```
\(\operatorname{Cov}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}\)
\(=\mathrm{E}\left\{\left(\mathbf{F}_{n} \mathbf{x}_{n-1}+\mathbf{q}_{n}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{F}_{n} \mathbf{x}_{n-1}+\mathbf{q}_{n}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\}\).
```

By rewriting and expanding the quadratic term and distributing the expectation we obtain

$$
\begin{aligned}
& \operatorname{Cov}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
&= \mathrm{E}\left\{\left(\mathbf{F}_{n} \mathbf{x}_{n-1}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{F}_{n} \mathbf{x}_{n-1}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)^{\top}\right\} \\
&+\mathrm{E}\left\{\mathbf{q}_{n}\left(\mathbf{F}_{n} \mathbf{x}_{n-1}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)^{\mathrm{T}} \mid \mathbf{y}_{1: n-1}\right\} \\
&+\mathrm{E}\left\{\left(\mathbf{F}_{n} \mathbf{x}_{n-1}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right) \mathbf{q}_{n}^{\top} \mid \mathbf{y}_{1: n-1}\right\}+\mathrm{E}\left\{\mathbf{q}_{n} \mathbf{q}_{n}^{\top} \mid \mathbf{y}_{1: n-1}\right\} .
\end{aligned}
$$

Noting that the middle terms are zero due to the factors being independent and $\mathbf{q}_{n}$ being zero-mean, what remains is

$$
\begin{aligned}
& \operatorname{Cov}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{F}_{n} \mathbf{x}_{n-1}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{F}_{n} \mathbf{x}_{n-1}-\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& \quad+\mathrm{E}\left\{\mathbf{q}_{n} \mathbf{q}_{n}^{\top} \mid \mathbf{y}_{1: n-1}\right\},
\end{aligned}
$$

and finally,

$$
\begin{equation*}
\operatorname{Cov}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}=\mathbf{F}_{n} \mathbf{P}_{n-1 \mid n-1} \mathbf{F}_{n}^{\top}+\mathbf{Q}_{n} . \tag{6.13}
\end{equation*}
$$

Kalman Filter. Gathering the results (6.8)-(6.10) and (6.12)-(6.13), the filtering algorithm for linear state-space models can now be formulated. First, during the prediction (time update) step the predicted mean and covariance are calculated according to

$$
\begin{align*}
\hat{\mathbf{x}}_{n \mid n-1} & =\mathbf{F}_{n} \hat{\mathbf{x}}_{n-1 \mid n-1}  \tag{6.14a}\\
\mathbf{P}_{n \mid n-1} & =\mathbf{F}_{n} \mathbf{P}_{n-1 \mid n-1} \mathbf{F}_{n}^{\top}+\mathbf{Q}_{n} \tag{6.14b}
\end{align*}
$$

Second, in the measurement update step, the new measurement is incorporated and the new state is estimated using

$$
\begin{align*}
\mathbf{K}_{n} & =\mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}\left(\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right)^{-1}  \tag{6.15a}\\
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}\right)  \tag{6.15b}\\
\mathbf{P}_{n \mid n} & =\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\top} \tag{6.15c}
\end{align*}
$$

Furthermore, the recursion is initialized by letting $\hat{\mathbf{x}}_{0 \mid 0}=\mathbf{m}_{0}$ and $\mathbf{P}_{0 \mid 0}=\mathbf{P}_{0}$. This leads to the algorithm shown in Algorithm 6.1, which is called the Kalman filter (KF) (Kalman, 1960).

It is worth pointing out several aspects of the prediction and measurement update (6.14)-(6.15). First, note that in the prediction, the covariance increases due to the scaling factor $\mathbf{F}_{n}$ and the added uncertainty by the process noise. On the other hand, during the measurement update, the covariance is decreased by a factor that scales quadratically with the gain $\mathbf{K}_{n}$. This follows the intuition that a prediction should increase the uncertainty, while incorporating new information should decrease it.

Second, the measurement update is a sum of the prediction $\hat{\mathbf{x}}_{n \mid n-1}$ and the term $\mathbf{y}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}$ scaled by the gain $\mathbf{K}_{n}$. Noting that $\mathbf{G}_{n}$ actually is the matrix relating the state $\mathbf{x}_{n}$ to the measurement $\mathbf{y}_{n}$, the term $\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}$ can be interpreted as a prediction of the output, which it in fact is. To show this, consider the expected value of the output $\mathbf{y}_{n}$ given all the data up to $t_{n-1}$, that is,

$$
\begin{aligned}
\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} & =\mathrm{E}\left\{\mathbf{G}_{n} \mathbf{x}_{n}+\mathbf{r}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathbf{G}_{n} \mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}+\mathrm{E}\left\{\mathbf{r}_{n} \mid \mathbf{y}_{1: n-1}\right\}
\end{aligned}
$$

and thus

$$
\begin{equation*}
\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}=\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1} \tag{6.16}
\end{equation*}
$$

Hence, the difference between the measurement and its prediction gives an indication about how far the predicted state is from the true state: If the difference is large, the predicted state is far and should be corrected a lot, if it is small, it is close and does not need to be corrected much. Hence, this difference is also called the innovation.

Third, the covariance of the predicted output $\mathbf{y}_{n}$ is given by

$$
\begin{aligned}
& \operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{y}_{n}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)\left(\mathbf{y}_{n}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{G}_{n} \mathbf{x}_{n}+\mathbf{r}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{G}_{n} \mathbf{x}_{n}+\mathbf{r}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{G}_{n}\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mathbf{G}_{n}^{\top}+\mathbf{r}_{n} \mathbf{r}_{n}^{\top} \mid \mathbf{y}_{1: n-1}\right\}\right.
\end{aligned}
$$

and finally,

$$
\begin{equation*}
\operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}=\mathbf{G}_{n} \mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top}+\mathbf{R}_{n} \tag{6.17}
\end{equation*}
$$

which is the denominator of the Kalman gain $\mathbf{K}_{n}$. Similarly, the cross-covariance between the predicted state $\mathbf{x}_{n}$ and predicted output $\mathbf{y}_{n}$ is

$$
\begin{aligned}
& \operatorname{Cov}\left\{\mathbf{x}_{n}, \mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{x}_{n}-\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)\left(\mathbf{y}_{n}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{G}_{n} \mathbf{x}_{n}+\mathbf{r}_{n}-\mathbf{G}_{n} \hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mathbf{G}_{n}^{\top} \mid \mathbf{y}_{1: n-1}\right\},
\end{aligned}
$$

which yields

$$
\begin{equation*}
\operatorname{Cov}\left\{\mathbf{x}_{n}, \mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}=\mathbf{P}_{n \mid n-1} \mathbf{G}_{n}^{\top} \tag{6.18}
\end{equation*}
$$

Thus, using (6.16)-(6.18), the measurement update can be written in the alternative form as

$$
\begin{align*}
\mathbf{K}_{n} & =\operatorname{Cov}\left\{\mathbf{x}_{n}, \mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} \operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}^{-1}  \tag{6.19a}\\
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)  \tag{6.19b}\\
\mathbf{P}_{n \mid n} & =\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n} \operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} \mathbf{K}_{n}^{\top} \tag{6.19c}
\end{align*}
$$

From (6.19), we can see that the denominator of the Kalman gain $\mathbf{K}_{n}$ is the covariance of the predicted measurement. Hence, if the uncertainty of the prediction is large (either due to a large uncertainty in the predicted state or due to large measurement noise covariance $\mathbf{R}_{n}$ ), the gain becomes small. This in turn means that the innovation only contributes little information to the updated state. Conversely, if the uncertainty of the prediction is small, the gain becomes large and the innovation contributes a lot.

### 6.3 Extended Kalman Filter

Similar to the static, linear case discussed in Chapter 3, the Kalman filter is the closed-form solution for the state estimation problem in linear state-space models and thus an exact solution. However, if we consider general nonlinear state-space models of the form (6.1)-(6.2), closed-form solutions are normally not available. Similar to the nonlinear static case, we need approximative solutions instead. A
first approximation is found in the extended Kalman filter (EKF), which is based on a linearization using a Taylor series approximation of the nonlinear dynamic and measurement models. Using this linearization, prediction and measurement update steps similar to the Kalman filter are found.

Prediction. For the prediction, assume that the state estimate and its covariance for $t_{n-1}$ are given by $\hat{\mathbf{x}}_{n-1 \mid n-1}$ and $\mathbf{P}_{n-1 \mid n-1}$, respectively. Then, the linear (first order) Taylor series approximation of the dynamic model around the linearization point $\hat{\mathbf{x}}_{n-1 \mid n-1}$ is given by

$$
\begin{align*}
\mathbf{x}_{n} & =\mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{q}_{n}  \tag{6.20}\\
& \approx \mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{q}_{n}
\end{align*}
$$

where $\mathbf{F}_{x}$ is the Jacobian matrix of the vector-value function $\mathbf{f}$. Then, based on the linearized dynamic model (6.20), we can predict the mean and covariance of the state $\mathbf{x}_{n}$ at $t_{n}$ given the set of all the measurements $\mathbf{y}_{1: n-1}=\left\{\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{n-1}\right\}$.

The mean is given by

$$
\begin{aligned}
& \hat{\mathbf{x}}_{n \mid n-1} \\
& =\mathrm{E}\left\{\mathbf{x}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& \approx \mathrm{E}\left\{\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{q}_{n} \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \mathrm{E}\left\{\mathbf{x}_{n-1} \mid \mathbf{y}_{1: n-1}\right\}-\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \hat{\mathbf{x}}_{n-1 \mid n-1} \\
& =\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \hat{\mathbf{x}}_{n-1 \mid n-1}-\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \hat{\mathbf{x}}_{n-1 \mid n-1} \\
& =\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right),
\end{aligned}
$$

and thus

$$
\begin{equation*}
\hat{\mathbf{x}}_{n \mid n-1}=\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \tag{6.21}
\end{equation*}
$$

Similarly, the covariance is

$$
\begin{aligned}
& \mathbf{P}_{n \mid n-1} \\
& =\mathrm{E}\left\{\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& \approx \\
& \mathrm{E}\left\{\left(\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{q}_{n}-\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\right)\right. \\
& \quad \times\left(\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{q}_{n}-\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\right)^{\top} \\
& \left.\quad \mid \mathbf{y}_{1: n-1}\right\} \\
& =\mathrm{E}\left\{\left(\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{q}_{n}\right)\right. \\
& \left.\quad \times\left(\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{q}_{n}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& = \\
& =\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \mathrm{E}\left\{\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)\left(\mathbf{x}_{n-1}-\hat{\mathbf{x}}_{n-1 \mid n-1}\right)^{\top} \mid \mathbf{y}_{1: n-1}\right\} \\
& \\
& \quad \times \mathbf{F}_{x}^{\top}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathrm{E}\left\{\mathbf{q}_{n} \mathbf{q}_{n}^{\top} \mid \mathbf{y}_{1: n-1}\right\}
\end{aligned}
$$

and finally,

$$
\begin{equation*}
\mathbf{P}_{n \mid n-1}=\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \mathbf{P}_{n-1 \mid n-1} \mathbf{F}_{x}^{\top}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{Q}_{n} \tag{6.22}
\end{equation*}
$$

Measurement Update. The derivation of the measurement update follows very similar steps as the prediction in the EKF and the measurement update in the KF. Given the predicted mean $\mathbf{x}_{n \mid n-1}$ and covariance $\mathbf{P}_{n \mid n-1}$, the nonlinear measurement model is linearized using a first order Taylor series expansion around the prediction $\hat{\mathbf{x}}_{n \mid n-1}$. This yields

$$
\begin{align*}
\mathbf{y}_{n} & =\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n}  \tag{6.23}\\
& =\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{r}_{n}
\end{align*}
$$

where $\mathbf{G}_{x}$ denotes the Jacobian matrix of $\mathbf{g}$.
Next, introducing the regularized least squares criterion for the linearized model (6.23) yields

$$
\begin{align*}
J_{\operatorname{ReLS}}\left(\mathbf{x}_{n}\right)=( & \left.\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)-\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)\right)^{\top} \mathbf{R}_{n}^{-1} \\
& \times\left(\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)-\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)\right)  \tag{6.24}\\
& +\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mathbf{P}_{n \mid n-1}^{-1}\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right) .
\end{align*}
$$

By introducing a change of variables with $\mathbf{z}_{n}=\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+$ $\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \hat{\mathbf{x}}_{n \mid n-1}$, (6.24) can be rewritten as

$$
\begin{gather*}
J_{\operatorname{ReLS}}\left(\mathbf{x}_{n}\right)=\left(\mathbf{z}_{n}-\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{x}_{n}\right)^{\top} \mathbf{R}_{n}^{-1}\left(\mathbf{z}_{n}-\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{x}_{n}\right) \\
+\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top} \mathbf{P}_{n \mid n-1}^{-1}\left(\mathbf{x}_{n}-\hat{\mathbf{x}}_{n \mid n-1}\right) \tag{6.25}
\end{gather*}
$$

which is linear in $\mathbf{x}_{n}$ and has the same form as (6.6) but with the Jacobian matrix $\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)$ rather than the measurement matrix $\mathbf{G}_{n}$. Hence, solving

$$
\begin{equation*}
\hat{\mathbf{x}}_{n \mid n}=\underset{\mathbf{x}_{n}}{\operatorname{argmin}} J_{\mathrm{ReLS}}\left(\mathbf{x}_{n}\right) \tag{6.26}
\end{equation*}
$$

yields

$$
\begin{aligned}
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{z}_{n}-\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \hat{\mathbf{x}}_{n \mid n-1}\right) \\
\mathbf{K}_{n} & =\mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\mathrm{T}}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\mathrm{T}}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{R}_{n}\right)^{-1}
\end{aligned}
$$

with covariance

$$
\mathbf{P}_{n \mid n} \approx \mathbf{P}_{n \mid n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\top}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\top}
$$

Finally, changing $\mathbf{z}_{n}$ back to its definition yields

$$
\begin{aligned}
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \hat{\mathbf{x}}_{n \mid n-1}-\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \hat{\mathbf{x}}_{n \mid n-1}\right) \\
& =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\right)
\end{aligned}
$$

for the measurement update.

```
Algorithm 6.2 Extended Kalman Filter
    Initialize \(\hat{\mathbf{x}}_{0 \mid 0}=\mathbf{m}_{0}, \mathbf{P}_{0 \mid 0}=\mathbf{P}_{0}\)
    for \(n=1,2, \ldots\) do
        Prediction (time update):
\[
\begin{aligned}
\hat{\mathbf{x}}_{n \mid n-1} & =\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \\
\mathbf{P}_{n \mid n-1} & =\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \mathbf{P}_{n-1 \mid n-1} \mathbf{F}_{x}^{\top}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{Q}_{n}
\end{aligned}
\]
```

Measurement update:

$$
\begin{aligned}
& \mathbf{K}_{n}=\mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\boldsymbol{\top}}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\boldsymbol{\top}}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{R}_{n}\right)^{-1} \\
& \hat{\mathbf{x}}_{n \mid n}=\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\right) \\
& \mathbf{P}_{n \mid n}=\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\boldsymbol{\top}}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\boldsymbol{\top}} \\
&: \text { end for }
\end{aligned}
$$

Extended Kalman Filter. We can now summarize the prediction and measurement update steps for the EKF. First, during the predictions step, the predicted mean and covariance are calculated according to

$$
\begin{align*}
\hat{\mathbf{x}}_{n \mid n-1} & =\mathbf{f}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right),  \tag{6.27a}\\
\mathbf{P}_{n \mid n-1} & =\mathbf{F}_{x}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right) \mathbf{P}_{n-1 \mid n-1} \mathbf{F}_{x}^{\boldsymbol{\top}}\left(\hat{\mathbf{x}}_{n-1 \mid n-1}\right)+\mathbf{Q}_{n}, \tag{6.27b}
\end{align*}
$$

where $\mathbf{F}_{x}$ is the Jacobian matrix of the dynamic model. Second, the measurement update is

$$
\begin{align*}
\mathbf{K}_{n} & =\mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\top}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\top}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{R}_{n}\right)^{-1},  \tag{6.28a}\\
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathbf{g}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)\right)  \tag{6.28b}\\
\mathbf{P}_{n \mid n} & =\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n}\left(\mathbf{G}_{x}\left(\hat{\mathbf{x}}_{n \mid n-1}\right) \mathbf{P}_{n \mid n-1} \mathbf{G}_{x}^{\top}\left(\hat{\mathbf{x}}_{n \mid n-1}\right)+\mathbf{R}_{n}\right) \mathbf{K}_{n}^{\top} . \tag{6.28c}
\end{align*}
$$

These two steps are then performed iteratively, and the algorithm is initialized with the initial conditions. This yields the complete EKF algorithm shown in Algorithm 6.2. Note that this is essentially the same algorithm as the original KF where the nonlinear function takes the place of the prediction (both in the prediction step and the prediction of the output) and in the covariance updates, the Jacobian matrices take the place of the matrices in the dynamic model and measurement model. Furthermore, if either of the models is linear, the corresponding step (prediction or measurement update) may be replaced by an exact update step from the Kalman filter in Section 6.2. Finally, note that Algorithm 6.2 is an approximation of the filtering problem since the nonlinear function is approximated around the filtered and predicted state, meaning that both the state estimate and its covariance may or may not converge to the true state (similar as for static nonlinear problems).

### 6.4 Unscented Kalman Filtering

One of the major problems of the EKF is that the linearization is local, which often leads to problems such as the covariance of the state being underestimated (which in turn affects the Kalman gain in the next iteration). Another approach to solve the filtering problem within the Kalman filtering framework is the unscented Kalman filter (UKF) (Wan and Van Der Merwe, 2000; Julier and Uhlmann, 2004), which makes use of a so-called unscented transform. Hence, this transform will be discussed first, and then the prediction and measurement update steps for the UKF are derived.

Unscented Transform. The unscented transform is based on calculating the moments of the nonlinear transformation of a finite set of deterministic sampling points as follows. Given a random variable $\mathbf{x}$ with mean $\mathbf{m}$ and covariance $\mathbf{P}$, we can find a set of $J$ points $\left\{\mathbf{x}^{0}, \mathbf{x}^{2}, \ldots, \mathbf{x}^{J-1}\right\}$, so-called sigma-points, such that their weighted sum is equal to the mean and covariance of $\mathbf{x}$, that is, such that

$$
\begin{align*}
\mathbf{m} & =\sum_{j=0}^{J-1} w_{m}^{j} \mathbf{x}^{j}  \tag{6.29a}\\
\mathbf{P} & =\sum_{j=0}^{J-1} w_{P}^{j}\left(\mathbf{x}^{j}-\mathbf{m}\right)\left(\mathbf{x}^{j}-\mathbf{m}\right)^{\top} \tag{6.29b}
\end{align*}
$$

where it must hold that $\sum_{j=0}^{J-1} w_{m}^{j}=1$ since the expected value of the sum should be unbiased.

Then, given the nonlinear function $\mathbf{z}=h(\mathbf{x})$, we can calculate the transformed sigma-points according to

$$
\mathbf{z}^{j}=h\left(\mathbf{x}^{j}\right)
$$

Based on the transformed sigma-points, the mean and covariance of the transformed variable $\mathbf{z}$ as well as the cross-covariance between $\mathbf{x}$ and $\mathbf{z}$ can then be calculated according to

$$
\begin{align*}
\mathrm{E}\{\mathbf{z}\} & \approx \sum_{j=1}^{J} w_{m}^{j} \mathbf{z}^{j},  \tag{6.30a}\\
\operatorname{Cov}\{\mathbf{z}\} & \approx \sum_{j=1}^{J} w_{P}^{j}\left(\mathbf{z}^{j}-\mathrm{E}\{\mathbf{z}\}\right)\left(\mathbf{z}^{j}-\mathrm{E}\{\mathbf{z}\}\right)^{\top},  \tag{6.30b}\\
\operatorname{Cov}\{\mathbf{x}, \mathbf{z}\} & \approx \sum_{j=1}^{J} w_{P}^{j}\left(\mathbf{x}^{j}-\mathbf{m}\right)\left(\mathbf{z}^{j}-\mathrm{E}\{\mathbf{z}\}\right)^{\top} . \tag{6.30c}
\end{align*}
$$

The question then is how to chose the sigma-points $\mathbf{x}^{j}$ and their weights for the mean $w_{m}^{j}$ and covariance $w_{P}^{j}$. The unscented transform is one such approach for
choosing sigma-points that fulfill the conditions (6.29). Here, $J=2 L+1$ (with $L$ being the dimension of the vector $\mathbf{x}$ ) sigma-points are chosen according to

$$
\begin{array}{ll}
\mathbf{x}^{0}=\mathbf{m} \\
\mathbf{x}^{j}=\mathbf{m}+\sqrt{L+\lambda}[\sqrt{\mathbf{P}}]_{j}, & j=1, \ldots, L \\
\mathbf{x}^{j}=\mathbf{m}-\sqrt{L+\lambda}[\sqrt{\mathbf{P}}]_{(j-L)}, & j=L+1, \ldots, 2 L \tag{6.31c}
\end{array}
$$

In (6.31), $\sqrt{\mathbf{P}}$ denotes a matrix square root such that $\mathbf{P}=\sqrt{\mathbf{P}} \sqrt{\mathbf{P}^{\top}}$ (in practice, this can be implemented using the Cholesky factorization) and $[\sqrt{\mathbf{P}}]_{j}$ denotes the $j$ th column of the matrix $\sqrt{\mathbf{P}}$. The weights corresponding to the above sigmapoints are

$$
\begin{align*}
w_{m}^{0} & =\frac{\lambda}{L+\lambda}  \tag{6.32a}\\
w_{P}^{0} & =\frac{\lambda}{L+\lambda}+\left(1-\alpha^{2}+\beta\right),  \tag{6.32b}\\
w_{m}^{j} & =w_{P}^{j}=\frac{1}{2(L+\lambda)}, \quad j=1, \ldots, 2 L \tag{6.32c}
\end{align*}
$$

In both (6.31) and (6.32), $\lambda$ is

$$
\begin{equation*}
\lambda=\alpha^{2}(L+\kappa)-L \tag{6.33}
\end{equation*}
$$

and $\alpha, \beta$, and $\kappa$ are tuning parameters. The parameters $\alpha$ and $\kappa$ determine the scaling of the spread of the sigma-points in the direction of the covariance $\mathbf{P}$ whereas $\beta$ is related the higher order moments of $\mathbf{x}$ and only affects the weight for the central sigma-point in the covariance calculations.

The choice of the tuning parameters may greatly affect the performance of the filter and several default choices have been suggested. The value for $\kappa$ is most often chosen to be zero (i.e., $\kappa=0$ ), which leaves the choices for $\alpha$ and $\beta$. First, note that with $\kappa=0$, the weight of the central sigma-point becomes

$$
\begin{equation*}
w_{m}^{0}=\frac{\lambda}{L+\lambda}=\frac{L \alpha^{2}-L}{L+L \alpha^{2}-L}=\frac{\alpha^{2}-1}{\alpha^{2}} \tag{6.34}
\end{equation*}
$$

and the scaling factor of the root of the covariance matrix becomes

$$
\begin{equation*}
\sqrt{L+\lambda}=\sqrt{L+\alpha^{2} L-L}=\alpha \sqrt{L} \tag{6.35}
\end{equation*}
$$

One suggestion is to choose $\alpha=1 \times 10^{-3}$ (Wan and Van Der Merwe, 2000). This gives a quite large and negative weight $w_{m}^{0}$ [see (6.34)] and a small scaling factor [see (6.35)], which makes the sigma-points lie close to the mean. To avoid this, it is sometimes more intuitive to instead start from the weight of the central point $w_{m}^{0}$ and then determine $\alpha$ based on reformulating (6.34), which yields

$$
\begin{equation*}
\alpha=\sqrt{\frac{1}{1-w_{m}^{0}}} . \tag{6.36}
\end{equation*}
$$

Furthermore, since all the weights for $j>0$ are the same [see (6.32)], it must also hold that

$$
w_{m}^{j}=\frac{1-w_{m}^{0}}{2 L}
$$

Finally, $\beta$ only affects the central weight of the covariance and a good starting point is normally to chose $\beta=2$ (see, for example, Wan and Van Der Merwe (2000)).

Prediction. The unscented transform as introduced above can directly be used in the prediction step to calculate the predicted mean and the covariance. In this case, the nonlinear transformation is the function of the dynamic model, that is, $\mathbf{f}\left(\mathbf{x}_{n-1}\right)$. The sigma-points are calculated using the estimated mean and covariance at $t_{n}$, that is,

$$
\begin{array}{rlr}
\mathbf{x}_{n-1}^{0}=\hat{\mathbf{x}}_{n-1 \mid n-1} \\
\mathbf{x}_{n-1}^{j}=\hat{\mathbf{x}}_{n-1 \mid n-1}+\sqrt{L+\lambda}\left[\sqrt{\mathbf{P}_{n-1 \mid n-1}}\right]_{j}, & j=1, \ldots, L, \\
\mathbf{x}_{n-1}^{j}=\hat{\mathbf{x}}_{n-1 \mid n-1}-\sqrt{L+\lambda}\left[\sqrt{\mathbf{P}_{n-1 \mid n-1}}\right]_{(j-L)}, & j=L+1, \ldots, 2 L . \tag{6.37c}
\end{array}
$$

The weights of the sigma-points are given by (6.32) and $\lambda$ is as in (6.33).
Then, the moments of the transformed variable, that is, the moments of the prediction become

$$
\begin{align*}
\mathbf{x}_{n}^{j} & =\mathbf{f}\left(\mathbf{x}_{n-1}^{j}\right), \quad j=0, \ldots, 2 L  \tag{6.38a}\\
\hat{\mathbf{x}}_{n \mid n-1} & =\sum_{j=0}^{2 L} w_{m}^{j} \mathbf{x}_{n}^{j}  \tag{6.38b}\\
\mathbf{P}_{n \mid n-1} & =\sum_{j=0}^{2 L} w_{c}^{j}\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n-1}\right)^{\top}+\mathbf{Q}_{n} \tag{6.38c}
\end{align*}
$$

Note that the additional term $\mathbf{Q}_{n}$ in the covariance is due to the fact that the dynamic model also includes the process noise term $\mathbf{q}_{n}$, which increases the uncertainty. In other words, the unscented transform only calculates the covariance of the $\mathbf{x}_{n-1}$ transformed by $\mathbf{f}\left(\mathbf{x}_{n-1}\right)$ and does not take the effect of the noise into account in this form.

Measurement Update. For the measurement update, first recall that it can be written in terms of the predicted output, its covariance, and the covariance between the predicted output and the state, see (6.19). Hence, we can also use the unscented transform to calculate these means and covariances.

```
Algorithm 6.3 Unscented Kalman Filter
    Initialize \(\hat{\mathbf{x}}_{0 \mid 0}=\mathbf{m}_{0}, \mathbf{P}_{0 \mid 0}=\mathbf{P}_{0}\)
    for \(n=1,2, \ldots\) do
        Calculate \(\mathbf{x}_{n-1}^{j}, w_{m}^{j}\), and \(w_{P}^{j}\) using (6.37) and (6.32)
        Calculate \(\hat{\mathbf{x}}_{n \mid n-1}\) and \(\mathbf{P}_{n \mid n-1}\) using (6.38)
        Calculate \(\mathbf{x}_{n}^{j}, w_{m}^{j}\), and \(w_{P}^{j}\) using (6.39) and (6.32)
        Calculate \(\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}, \operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\), and \(\operatorname{Cov}\left\{\mathbf{x}_{n}, \mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\)
    using (6.40)
        Measurement update
```

$$
\begin{aligned}
\mathbf{K}_{n} & =\operatorname{Cov}\left\{\mathbf{x}_{n}, \mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} \operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}^{-1} \\
\hat{\mathbf{x}}_{n \mid n} & =\hat{\mathbf{x}}_{n \mid n-1}+\mathbf{K}_{n}\left(\mathbf{y}_{n}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right) \\
\mathbf{P}_{n \mid n} & =\mathbf{P}_{n \mid n-1}-\mathbf{K}_{n} \operatorname{Cov}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\} \mathbf{K}_{n}^{\top}
\end{aligned}
$$

end for

In this case, the sigma points are calculated from the predicted mean $\hat{\mathbf{x}}_{n \mid n-1}$ and the covariance $\mathbf{P}_{n \mid n-1}$ according to

$$
\begin{array}{ll}
\mathbf{x}_{n}^{0}=\hat{\mathbf{x}}_{n \mid n-1} \\
\mathbf{x}_{n}^{j}=\hat{\mathbf{x}}_{n \mid n-1}+\sqrt{L+\lambda}\left[\sqrt{\mathbf{P}_{n \mid n-1}}\right]_{j}, & j=1, \ldots, L, \\
\mathbf{x}_{n}^{j}=\hat{\mathbf{x}}_{n \mid n-1}-\sqrt{L+\lambda}\left[\sqrt{\mathbf{P}_{n \mid n-1}}\right]_{(j-L)}, & j=L+1, \ldots, 2 L, \tag{6.39c}
\end{array}
$$

with the weights as in (6.32). Then, the necessary moments become

$$
\begin{align*}
\mathbf{y}_{n}^{j} & =\mathbf{g}\left(\mathbf{x}_{n}^{j}\right), \quad j=0, \ldots, 2 L, \\
\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}= & \sum_{j=0}^{2 L} w_{m}^{j} \mathbf{y}_{n}^{j}, \\
\operatorname{Cov}\left\{\mathbf{y} \mid \mathbf{y}_{1: n-1}\right\}= & \sum_{j=0}^{2 L} w_{P}^{j}\left(\mathbf{y}_{n}^{j}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)\left(\mathbf{y}_{n}^{j}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)^{\top} \\
& +\mathbf{R}_{n}, \\
\operatorname{Cov}\left\{\mathbf{x}_{n}, \mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}= & \sum_{j=0}^{2 L} w_{P}^{j}\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n-1}\right)\left(\mathbf{y}_{n}^{j}-\mathrm{E}\left\{\mathbf{y}_{n} \mid \mathbf{y}_{1: n-1}\right\}\right)^{\top},
\end{align*}
$$

and the measurement update given in (6.19) can be performed.

Unscented Kalman Filter. Based on the prediction and measurement update discussed above, the unscented Kalman filter then becomes as shown in Algorithm 6.3. Note that the UKF still performs the two steps very similar to the original Kalman filter for linear systems, that is, it basically estimates the mean and covariance of the state at each time step and propagates this information between time steps, but without analytical linearization. The latter also implies that the UKF does not require any Jacobians to be calculated and is thus somewhat easier to implement.

### 6.5 Iterated Extended and Unscented Kalman Filters

In Section 6.3 we derived the extended Kalman filter (EKF) update by computing the regularized (weighted) least squares solution to a linearized model. However, this is exactly the same procedure that was used in Gauss-Newton and LevenbergMarquardt algorithms except that this linearization was iteratively used to compute the solution to the non-linear regularized least squares solution. This same approach can also be used to improve the EKF - instead of taking just a single step, we solve the non-linear regularized least squares problem using a least squares optimization method.

When the non-linear least squares problem is solved using subsequent linearization - which is equivalent to Gauss-Newton optimization - we obtain the iterated extended Kalman filter (Bell and Cathey, 1993). The iteration helps to solve highly non-linear problems where the single linearization used in EKF which is equivalent to single step of Gauss-Newton algorithm - is not enough. This concept can also be extended to include line search procedures and we can also use Levenberg-Marquardt algorithms to solve the non-linear least squares problem (Skoglund et al., 2015).

Furthermore, the unscented Kalman filter (UKF) introduced in Section 6.4 can also be improved by using iteration. In this case the optimization criterion is no longer the non-linear regularized least squares problem, but instead, we optimize the Kullback-Leibler divergence to the posterior distribution. Nevertheless, the resulting iterated posterior linearization filter (García-Fernández et al., 2015) takes a similar form as the iterated extended Kalman filter, but the Taylor series linearizations are replaced with sigma-point approximations.

### 6.6 Bootstrap Particle Filter

Another approach of solving the estimation problem in dynamic systems is particle filtering, which differs quite a lot from the Kalman filtering approaches discussed in the previous sections. Rather than being based on predicting and updating the mean and covariance from time step to time step, particle filtering propagates a set of weighted random samples (called particles) and hence there is some resemblance to the unscented Kalman filter. The particles can be seen as qualified but random
guesses of the state whereas the weights provide an indication of how good the guess is. Particle filtering is a very general methodology and comprises a large class of filtering algorithms. Here, we focus on the bootstrap particle filter, which has an intuitive explanation but also a rigorous mathematical background (Gordon et al., 1993; Doucet and Johansen, 2011; Särkkä, 2013).

To derive the particle filter, we first take another look at the general discretetime state-space model given by

$$
\begin{aligned}
& \mathbf{x}_{n}=\mathbf{f}\left(\mathbf{x}_{n-1}\right)+\mathbf{q}_{n}, \\
& \mathbf{y}_{n}=\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n},
\end{aligned}
$$

with $\mathbf{q}_{n} \sim p\left(\mathbf{q}_{n}\right)$ and $\mathbf{r}_{n} \sim p\left(\mathbf{r}_{n}\right)$. Recall that the dynamic model is a stochastic process, meaning that every time a system is observed, a new realization of that stochastic process is observed. For example, consider the scalar (one-dimensional) random walk dynamic model

$$
x_{n}=x_{n-1}+q_{n}
$$

with

$$
\begin{aligned}
x_{0} & \sim \mathcal{N}(0,1) \\
q_{n} & \sim \mathcal{N}(0,1)
\end{aligned}
$$

This model essentially states that the initial state is a random variable distributed around $x_{0}=0$ with variance 1 and as time elapses, random steps from that initial point are taken. Then, at any time $t_{n}$, there are infinitely many random steps that can be taken through $q_{n}$ (i.e., $q_{n}$ can take on any value, but with some values more likely than others as specified by the distribution of the random variable). This implies in turn that no two realizations of any random process are the same. As an example, Figure 6.1a illustrates 20 realization of the random walk process. As it can be seen, each realization takes a completely different path, and the realizations diverge as time elapses.

In practice, when estimating a system's state, all the observations (measurements) come from one single realization out of all the infinitely possible realizations of the process. The measurements then give us the information about the particular realization that is observed and the realization and the measurements are linked through the measurement model. For example, Figure 6.1b illustrates one particular realization of the random walk process together with the measurements from the linear measurement model

$$
y_{n}=x_{n}+r_{n}
$$

with $r_{n} \sim \mathcal{N}(0,1)$.
This leads to a new strategy for filtering where we first simulate a set of trajectories and then evaluate how well each of these trajectories explain the measurements that we are observing. More formally, in the framework of the prediction


Figure 6.1. Random walk process example: (a) 20 realizations of the random walk process, and (b) one realization of the process together with its measurements.


Figure 6.2. Random walk example with measurements $y_{n}$ (blue) from one particular realization of the random walk process (red) and 20 other realizations.
and measurement update steps introduced in Section 6.1, this approach alternates between:

1. Prediction: Given a set of simulated states $\mathbf{x}_{n-1}^{j}$ (for $j=1, \ldots, J$ ), simulate from $t_{n-1}$ to $t_{n}$ to obtain a set of simulated states $\mathbf{x}_{n}^{j}(j=1, \ldots, J)$;
2. Measurement update: Evaluate how well the simulated states $\mathbf{x}_{n}^{j}$ explain the observed measurement $\mathbf{y}_{n}$.

For example, Figure 6.2 combines the different realizations of the state trajectories in Figure 6.1a (gray) with the measurements (blue dots) in Figure 6.1b (together with the true trajectory that generated the measurements in red). As it can be seen from the figure, several of the gray realizations could be the trajectories that generated the measurements in blue, at least at some points in time, whereas other realizations are very unlikely to be the trajectories that generated the measurements.

To develop a filtering algorithm based on this reasoning, we need to find a strategy to a) simulate samples $\mathbf{x}_{n}^{j}$ given the samples $\mathbf{x}_{n-1}^{j}$, and b) evaluate how well a particular sample $\mathbf{x}_{n}^{j}$ of the state explains the current measurement $\mathbf{y}_{n}$.

An intuitive way to generate new samples is to use the dynamic model to simulate one time step. In other words, we can generate a realization of the random variable $\mathbf{q}_{n}$ and use the nonlinear function $\mathbf{f}\left(\mathbf{x}_{n-1}\right)$ to pass the sample from $t_{n-1}$ to $t_{n}$. This means performing the following two steps for each sample $j=1, \ldots, J$ :

1. Sample $\mathbf{q}_{n}^{j} \sim p\left(\mathbf{q}_{n}\right)$,
2. Calculate $\mathbf{x}_{n}^{j}=\mathbf{f}\left(\mathbf{x}_{n-1}^{j}\right)+\mathbf{q}_{n}^{j}$.

In practice, the process noise $\mathbf{q}_{n}$ is often Gaussian, that is, $\mathbf{q}_{n}$ is a Gaussian random variable (e.g., when the discrete-time model comes from the discretization of a continuous-time model with a white noise process as the input). In that case, sampling $\mathbf{q}_{n}$ amounts to sampling from the normal distribution $\mathcal{N}\left(0, \mathbf{Q}_{n}\right)$ where $\mathbf{Q}_{n}$ is the covariance matrix of $\mathbf{q}_{n}$.

To evaluate the importance of each sample with respect to the current measurement $\mathbf{y}_{n}$, we assign a weight $w_{n}^{j}$ (called the importance weight) to each sample $\mathbf{x}_{n}^{j}$. Intuitively, the weight should represent how well each sample explains the measurement, and thus, the closer the sample is to the true state, the higher the weight should be. If we ensure that the weights sum to one, that is, if

$$
\sum_{j=1}^{J} w_{n}^{j}=1
$$

we can also loosely interpret the weight $w_{n}^{j}$ as the probability of the $j$ th sample representing the correct state. The question then is how the weights should be calculated, given the sample $\mathbf{x}_{n}^{j}$ and the measurement $\mathbf{y}_{n}$. So far, cost functions have been used to evaluate whether a state explains the measurements well. However, in this approach the cost should be low for a state that explains the measurements well, whereas the importance weight should be high and hence, we can not use cost functions ${ }^{1}$.

Instead, we have another look at the measurement model. Recall that we have assumed that the measurement model is of the form

$$
\mathbf{y}_{n}=\mathbf{g}\left(\mathbf{x}_{n}\right)+\mathbf{r}_{n}
$$

where $\mathbf{r}_{n} \sim p\left(\mathbf{r}_{n}\right)$ is a random variable with probability density function $p\left(\mathbf{r}_{n}\right)$. Hence, $\mathbf{y}_{n}$ is a random variable too. Assuming that $\mathbf{x}_{n}$ is given, $\mathbf{y}_{n}$ follows the same distribution as $\mathbf{r}_{n}$ but offset by the constant term $\mathbf{g}\left(\mathbf{x}_{n}\right)$. We can express this probability density function for $\mathbf{y}_{n}$ as $p\left(\mathbf{y}_{n} \mid \mathbf{x}_{n}\right)$ which is read as "the probability density function of $\mathbf{y}_{n}$ given that $\mathbf{x}_{n}$ is known". The probability density function $p\left(\mathbf{y}_{n} \mid \mathbf{x}_{n}\right)$ is commonly known as the likelihood because it indicates how likely a certain state $\mathbf{x}_{n}$ is when observing the measurement $\mathbf{y}_{n}$. Hence, the likelihood is a

[^2]suitable measure for how well any given sample $\mathbf{x}_{n}^{j}$ explains the measurement $\mathbf{y}_{n}$, and we can define the non-normalized weight $\tilde{w}_{n}^{j}$ as
$$
\tilde{w}_{n}^{j}=p\left(\mathbf{y}_{n} \mid \mathbf{x}_{n}^{j}\right)
$$
$\tilde{w}_{n}^{j}$ is non-normalized because the sum over all weights does generally not fulfill the requirement that the weights should sum to one. This can be ensured by simply dividing each non-normalized weight by the sum of all the non-normalized weights, that is,
$$
w_{n}^{j}=\frac{\tilde{w}_{n}^{j}}{\sum_{i=1}^{J} \tilde{w}_{n}^{i}}
$$

In practice, the measurement noise is often assumed to be a zero-mean Gaussian random variable with covariance matrix $\mathbf{R}_{n}$, that is, $p\left(\mathbf{r}_{n}\right)=\mathcal{N}\left(\mathbf{r}_{n} ; 0, \mathbf{R}_{n}\right)$. In this case, the likelihood is also a Gaussian random variable with mean $\mathbf{g}\left(\mathbf{x}_{n}\right)$ (due to the offset) and covariance $\mathbf{R}_{n}$ (due to the uncertainty introduced by $\mathbf{r}_{n}$ ). Hence, the likelihood becomes

$$
p\left(\mathbf{y}_{n} \mid \mathbf{x}_{n}\right)=\mathcal{N}\left(\mathbf{y}_{n} ; \mathbf{g}\left(\mathbf{x}_{n}\right), \mathbf{R}_{n}\right)
$$

Once the importance weights have been calculated, a point estimate of the current state and its covariance can be obtained. Similar to the unscented Kalman filter, these are given by the weighted sum of the individual samples $\mathbf{x}_{n}^{j}$, weighed by the importance weights $w_{n}^{j}$, that is,

$$
\begin{align*}
\hat{\mathbf{x}}_{n \mid n} & =\sum_{j=1}^{J} w_{n}^{j} \mathbf{x}_{n}^{j}  \tag{6.41a}\\
\mathbf{P}_{n \mid n} & =\sum_{j=1}^{J} w_{n}^{j}\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n}\right)\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n}\right)^{\top} . \tag{6.41b}
\end{align*}
$$

Observe that the covariances for the process noise $\left(\mathbf{Q}_{n}\right)$ and measurement noise ( $\mathbf{R}_{n}$ ) do not enter these equations (contrary to the Kalman-filter-type algorithms). This is due to the fact that this uncertainty is accounted for when sampling and calculating the weights.

It would appear that this now yields a working algorithm. However, there is one important problem, illustrated for the random walk model in Figure 6.3: Assuming that the red trajectory represents the true realization of our state trajectory, the other simulated trajectories start to diverge from that trajectory as time elapses. Consequently, toward the end of the time scale, there are only very few trajectories in the neighborhood of the actual trajectory. In practice, this means that the weights for almost all but a few trajectories would become zero and eventually, after some more time has elapsed, there would be no trajectory close to the true realization anymore and the filter would break down completely.


Figure 6.3. Illustration of the divergence problem when simulating trajectories.

To address this problem, one additional step called resampling has to be introduced. The basic idea of resampling is to make sure that samples with low weights, that is, samples that are unlikely to be close to the true state, are discarded and replaced with copies of the samples with high weight. Hence, resampling essentially regenerates the sample $\mathbf{x}_{n}^{j}$ such that there are a total of approximately $\left\lfloor w_{n}^{j} J\right\rceil$ copies of $\mathbf{x}_{n}^{j}$ in the resampled set of particles. Thus, if $\tilde{\mathbf{x}}_{n}^{i}($ for $i=1, \ldots, J)$ denotes the resampled particle, that particle is equal to particle $\mathbf{x}_{n}^{j}$ with probability $w_{n}^{j}$, that is, we have that

$$
\operatorname{Pr}\left\{\tilde{\mathbf{x}}_{n}^{i}=\mathbf{x}_{n}^{j}\right\}=w_{n}^{j},
$$

where $\operatorname{Pr}\{\cdot\}$ denotes the probability. This ensures that trajectories with low weight are discarded and that the samples remain close to the true state.

The final bootstrap particle filtering algorithm then consists of the following three steps:

1. Propagate the particles using the dynamic model,
2. calculate the importance weights using the measurement model, and
3. resample the particles according to their weights.

These three steps are then repeated for the complete dataset and the recursion is initialized by sampling from the initial distribution of the state $p\left(\mathbf{x}_{0}\right)$. Assuming Gaussian distributions for the initial state, process noise, and measurement noise, that is,

$$
\begin{aligned}
\mathbf{x}_{0} & \sim \mathcal{N}\left(\mathbf{m}_{0}, \mathbf{P}_{0}\right), \\
\mathbf{q}_{n} & \sim \mathcal{N}\left(0, \mathbf{Q}_{n}\right), \\
\mathbf{r}_{n} & \sim \mathcal{N}\left(0, \mathbf{R}_{n}\right),
\end{aligned}
$$

the resulting algorithm is as summarized in Algorithm 6.4.

```
\(\overline{\text { Algorithm 6.4 Bootstrap Particle Filter (Gaussian Process and Measurement }}\)
Noises)
    1: Initialization: Sample \((j=1, \ldots, J)\)
\[
\mathbf{x}_{0}^{j} \sim \mathcal{N}\left(\mathbf{m}_{0}, \mathbf{P}_{0}\right)
\]
\[
\text { for } n=1,2, \ldots \text { do }
\]
\[
\text { for } j=1,2, \ldots, J \text { do }
\]
Sample
\[
\mathbf{q}_{n}^{j} \sim \mathcal{N}(0, \mathbf{Q})
\]
5: \(\quad\) Propagate the state
\[
\mathbf{x}_{n}^{j}=\mathbf{f}\left(\mathbf{x}_{n-1}^{j}\right)+\mathbf{q}_{n}^{j}
\]
```

6: $\quad$ Calculate the importance weights

$$
\tilde{w}_{n}^{j}=\mathcal{N}\left(\mathbf{y}_{n} ; \mathbf{g}\left(\mathbf{x}_{n}^{j}\right), \mathbf{R}_{n}\right)
$$

end for
8: $\quad$ Normalize the importance weights $(j=1, \ldots, J)$

$$
w_{n}^{j}=\frac{\tilde{w}_{n}^{j}}{\sum_{i=1}^{J} \tilde{w}_{n}^{i}}
$$

9: Calculate the mean and covariance

$$
\begin{aligned}
\hat{\mathbf{x}}_{n \mid n} & =\sum_{j=1}^{J} w_{n}^{j} \mathbf{x}_{n}^{j} \\
\mathbf{P}_{n \mid n} & =\sum_{j=1}^{J} w_{n}^{j}\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n}\right)\left(\mathbf{x}_{n}^{j}-\hat{\mathbf{x}}_{n \mid n}\right)^{\top}
\end{aligned}
$$

10: Resample such that

$$
\operatorname{Pr}\left\{\tilde{\mathbf{x}}_{n}^{i}=\mathbf{x}_{n}^{j}\right\}=w_{n}^{j}
$$

11: end for

## Bibliography

Bell, B. M. and Cathey, F. W. (1993). The iterated Kalman filter update as a Gauss-Newton method. IEEE Transactions on Automatic Control, 38(2):294297. (Cited on page 87.)

Doucet, A. and Johansen, A. M. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In Crisan, D. and Rozovskii, B., editors, Handbook of Nonlinear Filtering, volume 12 of Oxford Handbooks, pages 656-704. Oxford University Press, Oxford, UK. (Cited on page 88.)

García-Fernández, Á. F., Svensson, L., Morelande, M. R., and Särkkä, S. (2015). Posterior linearization filter: principles and implementation using sigma points. IEEE Transactions on Signal Processing, 63(20):5561-5573. (Cited on page 87.)

Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. Radar and Signal Processing, IEE Proceedings F, 140(2):107-113. (Cited on page 88.)

Gustafsson, F. (2018). Statistical Sensor Fusion. Studentlitteratur, 3rd edition. (Cited on pages v and 44.)

Hartley, R. and Zisserman, A. (2003). Multiple view geometry in computer vision. Cambridge University Press. (Cited on page 4.)

Julier, S. J. and Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. Proceedings of the IEEE, 92(3):401-422. (Cited on page 83.)

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Transactions of the ASME, Journal of Basic Engineering, 82(1):35-45. (Cited on page 78.)

Kay, S. M. (1993). Fundamentals of Statistical Signal Processing: Estimation Theory. Prentice Hall, Upper Saddle River, NJ, USA. (Cited on page v.)

Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. Quarterly of Applied Mathematics, 2(2):164-168. (Cited on page 45. )

Marquardt, D. (1963). An algorithm for least-squares estimation of nonlinear parameters. Journal of the Society for Industrial and Applied Mathematics, 11(2):431-441. (Cited on pages 45 and 47.)

Nocedal, J. and Wright, S. J. (2006). Numerical Optimization. Springer, 2nd edition. (Cited on pages v, 44, 45, and 50.)

Papoulis, A. (1984). Probability, Random Variables, and Stochastic Processes. McGraw-Hill. (Cited on page 57.)

Petersen, K. B. and Pedersen, M. S. (2012). The matrix cookbook. Technical report, Technical University of Denmark. (Cited on page 27.)

Pujol, J. (2007). The solution of nonlinear inverse problems and the LevenbergMarquardt method. Geophysics, 72(4):W1-W16. (Cited on page 47.)

Richards, M. A., Scheer, J., and Holm, W. A. (2010). Principles of Modern Radar. SciTech Publishing. (Cited on page 4.)

Skoglund, M., Hendeby, G., and Axehill, D. (2015). Extended Kalman filter modifications based on an optimization view point. In 18th International Conference on Information Fusion (Fusion), pages 1856-1861. (Cited on page 87.)

Särkkä, S. (2013). Bayesian Filtering and Smoothing. Cambridge University Press. (Cited on pages v, 76, and 88.)

Särkkä, S. and Solin, A. (2019). Applied Stochastic Differential Equations. Cambridge University Press. (Cited on pages v, 57, and 67.)

Wan, E. A. and Van Der Merwe, R. (2000). The unscented Kalman filter for nonlinear estimation. In Adaptive Systems for Signal Processing, Communications, and Control Symposium (AS-SPCC), pages 153-158. (Cited on pages 83, 84, and 85.)

Wilson, J. (2005). Sensor Technology Handbook. Elsevier. (Cited on page 11.)
Øksendal, B. (2010). Stochastic Differential Equations: An Introduction with Applications. Springer, 6th edition. (Cited on pages 57 and 67.)


[^0]:    ${ }^{1}$ There are more general formulations than (2.1) that do not share this trait, but we will not make use of these in these notes.

[^1]:    ${ }^{2}$ That is, a symmetric matrix $\mathbf{M}$ for which it holds that $\mathbf{x}^{\top} \mathbf{M x}>0$ for any arbitrary non-zero vectors $\mathbf{x}$. The eigenvalues of this kind of matrix are all positive.

[^2]:    ${ }^{1}$ In practice, cost functions are closely related to the way the importance weights are calculated, but this is beyond the scope of this course.

