Chapter 5

An Accurate Sensor Model Representation for Radio Frequency Sensors

5.1 Introduction

The previous chapters argued that the sensor models used to date are not sufficient for high-accuracy localisation and that an improved description of the Radio Frequency (RF) sensor behaviour (or the propagation behaviour) is needed to achieve this goal.
The main ideas presented in this chapter are summarised as:

- An alternative representation for RF sensors is shown, which allows for a detailed description of RF propagation. The model is compared to the nth power law.
- An implementation of Bayesian Decision Theory (BDT) is used to select the correct sensor model out of many available models. This has practical relevance in applications where the observation model to be used may change over time and automated model switching has to be performed.
- Multi-sensor localisation using range-only sensors is introduced to demonstrate how localisation and tracking can be performed with range-only sensors. The basic principle is shown, and a discussion is provided for issues such as the effects of ranging errors. Additionally, the use of external information to aid the localisation process is shown as an extension to the Close Proximity System (CPS) setup as presented in chapter 2.

5.2 A more accurate sensor model representation for Radio Frequency sensors

This section of the thesis presents and develops an alternative closed form description of an RF sensor model that can be used for probabilistic localisation and tracking. The main advantage of this model is that it accurately describes the RF signal mean. Using this sensor model, it is possible to track the position of agents in 2D in a single-sensor, or multi-sensor setup [50, 52]. The CPS shown in subsection 2.6.1 is an example of such a setup. The two-ray model is used as the basis for the theoretical development of the new model. This model illustrates the complexity of RF signal propagation, and that in comparison, the path loss is only very roughly described using the nth power law.

The implementation of the model in a real-world application will be based on geo-referenced measurements, which provide data to calculate certain parameters of the model. These parameters are partially obtained through a regression and partially through interpolation. In this section, the task of modelling the RF signal behaviour is divided into the following four components:

- Modelling of the free space loss and the antenna influence.
- Modelling of the residual components representing reflections, refraction and diffraction.
- Incorporation of statistical behaviour.
- Modelling of the angular dependency on the azimuth angle.

This procedure is shown graphically in figure 5.1, and serves as a guideline for the subsections to follow. This explains the entire procedure of modelling the RF signal behaviour.

5.2.1 Alternative Radio Frequency sensor model

Modelling of the free space loss and the antenna influence

The first component of the modelling process is to represent the free space loss, and the influence of the antenna mounting height on the propagation pattern as shown in figure 5.1. These two components are shown in isolation from other propagation effects in figure 5.2 (a). Here, the receiver was mounted at 2 m in height, while the transmitter was mounted at 2 m, 3 m and 4 m in height. The antenna radiation pattern considered is the previously described pattern of a $\lambda/4$ antenna. Alternative representations must be able to closely represent these curves obtained from physical considerations.

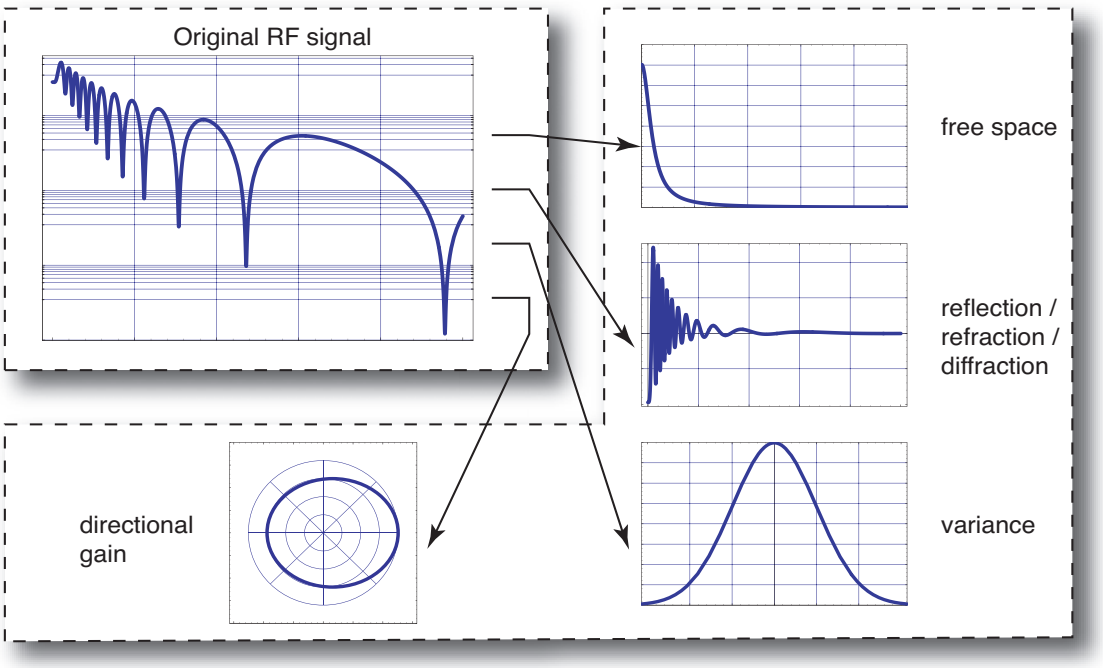

Figure 5.1: RF sensor modelling diagram — The task of modelling the RF sensor is split into four sub-tasks.

The alternative description for the free space loss with antenna influence, as just presented, is given by[1]

$$I_{FS,approx} = f(a,b,c,d,x) = a\exp(bx) + c\exp(dx) \tag{5.1}$$

This representation is an approximation, and apart from the sender-receiver distance x, also contains the four parameters a, b, c and d. These parameters allow for the adjustment of the RF-signal model to describe the real occurring signal drop obtained from experimental data. The parameters are chosen to minimise the mean squared error between the given data and equation (5.1), where the data can be obtained from experimental data, or from theoretical formulations.

It can be seen that the free space loss given by nth power model can be sufficiently modelled by equation (5.1). This result is also seen in the case using a short dipole antenna. Figure 5.2(b) shows this for the free space loss, while (c) and (d) represent the cases where the influence of the antenna becomes apparent.

[1]The notation I is used in the following sections to describe signal intensity in general, whether this be power, energy or Received Signal Strength Indicator (RSSI).

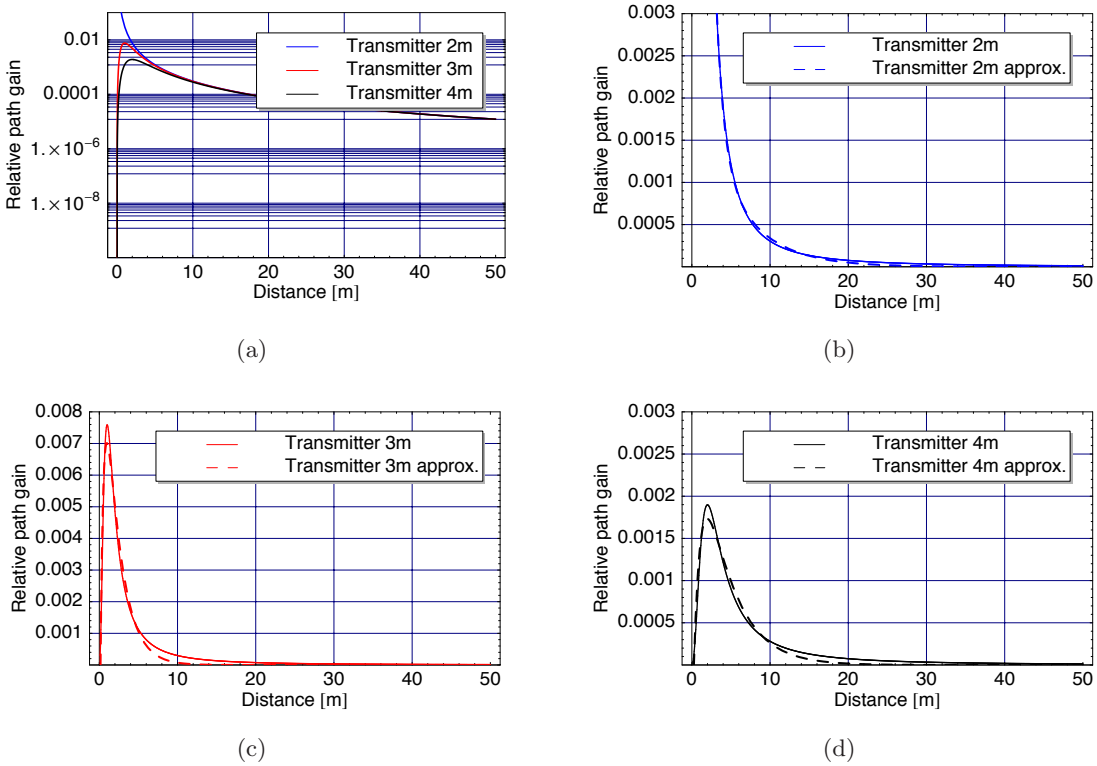

Figure 5.2: Path gain with antenna influence approximation — In all graphs the receiver was mounted at 2 m height and a $\lambda/4$ antenna was considered. (b) corresponds to the free space loss when transmitter and receiver are mounted at the same height and the antenna gain effectively becomes one. (c) and (d) show the effect of the antenna pattern which comes into effect as transmitter and receiver are mounted at different heights. The approximations (dashed) closely follow the given data (solid), which in this case is obtained from the deterministic two-ray model.

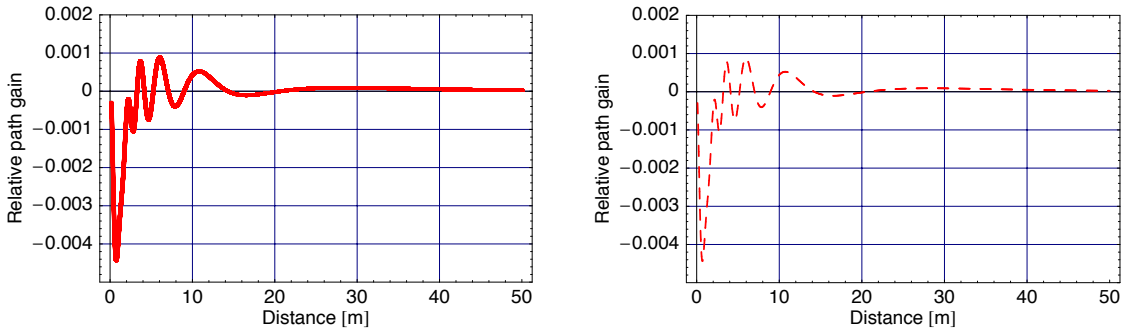

Figure 5.3: Residual components — (a) shows the residual components for the transmitter mounted at 3 m. (b) shows the approximation of the residual components using a FFT based approach.

Approximation of the residual components

The difference between the original signal obtained from the two-ray model and the assumed free-space loss with antenna effect is taken to be the residual components. Figure 5.3(a) shows an example of the case where the signal oscillates in a sinusoidal fashion. The nature of the resulting difference signal motivates the use of a Fourier Series Approximation[2] to model the reflective components. Using a Fast Fourier Transformation (FFT) on the difference signal, the complex Fourier coefficients c_k are obtained. The residual (difference) signal can be alternatively described by

$$I_{RESID,approx}(x) = \sum_{k=1}^{N} c_k e^{ik\omega x} \tag{5.2}$$
$$\omega = 2\pi f \qquad N \ldots \text{Number of sampling points}$$

Figure 5.3 shows the initial difference signal and the approximation $I_{RESID,approx}$ are very close, and consequently it is reasonable to use the FFT based approximation. This result is not surprising as the approximation is actually an interpolant, which means that the curve resulting from the approximation goes through the given data points. This is not the case for the approximation of the free space loss with antenna influence given by equation (5.1), which is a regression curve. It should be noted here that, to be able to obtain the FFT, the points must be sampled at constant spatial frequency, which in a real application is mostly not the case. This can be corrected through a resampling step before the FFT is applied.

[2] Appendix B offers a compact refresher on the topic of Fourier Series.

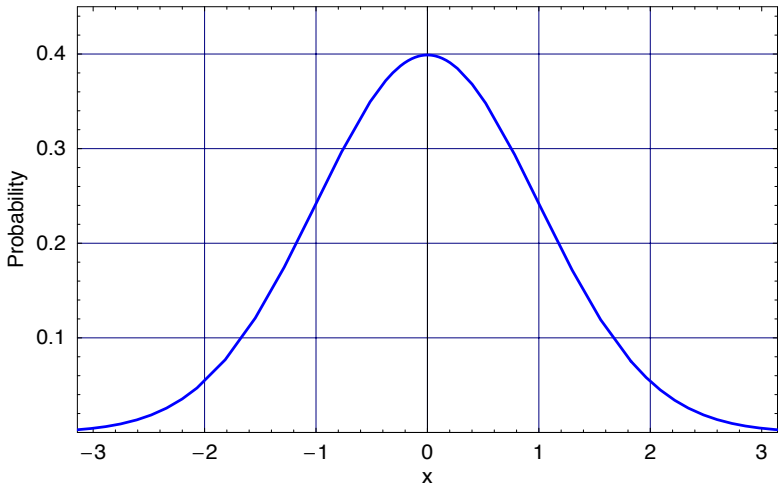

Figure 5.4: Static signal distribution — The signal for a constant distance between transmitter and receiver might vary according to a Gaussian distribution.

Variance of the sensor signal for static transmitter and receiver

In this section, the behaviour of the signal with varying distance and mounting heights of transmitter and receiver has been investigated. The resulting description of these effects occurred in a deterministic manner. The RF signal in reality does not behave so deterministically, there is also a random component. It can be argued that this random component is due to an inability to completely describe the entire RF communication link with its temporally and spatially occurring effects. It is observed, that in general, the signal strength between a transmitter and a receiver varies for a given constant distance. One reason for this is the electronic hardware is subjected to thermal noise, which is certainly a factor for this variation. Nevertheless, a probabilistic model offers the possibility of including uncertainty, for whatever reason this uncertainty might occur. It is therefore of interest to describe and characterise the distribution of the signal. One possibility is that it varies according to a Gaussian distribution as shown in figure 5.4, which is defined as

$$g(z;\mu,\sigma) \triangleq \frac{1}{\sqrt{(2\pi)}\sigma} \exp\left(-\frac{(z-\mu)^2}{2\sigma^2}\right) \tag{5.3}$$

where μ and σ are the mean and variance of the signal respectively, and z is the measured value, for example the RSSI. In the context of modelling the behaviour of an RF sensor the mean of this distribution could be obtained from the two-ray model or any other suitable approximation, such as the model developed in this chapter. To obtain the actual distribution of the signal it is usually necessary to conduct experiments, as it is not feasible to do this based on theoretical considerations alone.

It is noted at this point that the distribution of the RF signal is, for the hardware used, Gaussian if RSSI is measured in dBm. The variance is a function of the mean, shown in equation (5.4). The reader is also reminded that in a similar fashion, the general path loss model also includes a Gaussian random variable to account for the uncertainties.

$$\sigma = f(\mu) \tag{5.4}$$

For the RF sensor under consideration, it is found that the dependency of the standard deviation on the mean can be modelled using a simple linear equation. As the standard deviation increases, the mean of the signal decreases. Since the same signal level can be measured for different distances, it was verified that the standard deviation is independent of the separation distance, and only depends on the measured signal intensity. The general description of the dependency given by equation (5.4), can now be written as

$$\sigma(z) = mz + c \tag{5.5}$$

In this equation $\sigma(z)$ is the signal variance for the measurement z, whereas m is the gradient, and c an added constant.

One explanation for the linear dependence between signal mean and signal standard deviation lies in the hardware used, more specifically in the receiver chain of the radio chip. The hardware used has a built-in *superheterodyne receiver*, which usually has an amplifier (automatic gain control) in the receiver stage to adjust the incoming analog signal amplitude to the correct levels for the conversion to a digital signal [98]. The gain control will also affect the variance or standard deviation of the signal stemming from thermal noise in the receiver. This is because the gain control will increase the gain for low signals or decrease it for strong signals, and at the same time also amplify the noise. The receiver used in the experiments contains a linear amplifier.

As can be seen here, knowledge about the hardware (plant) is very useful in the modelling process. It helps to explain phenomena occurring and ultimately leads to a more accurate sensor model.

Angular dependency of the sensor signal intensity

Antennas often exhibit a directional radiation pattern. This pattern may be distorted on purpose by the design of the antenna, for example to focus the radiated energy in particular directions and thus extend the range in these directions. The distortion might also occur due to the installation environment, for example when the antenna is installed near a large metal structure. Consequently, the received intensity will not only be a function of the distance,

but also a function of both elevation and azimuth angle. The sensor model presented in this section is based on theoretical considerations for a short dipole antenna. This type of antenna has a pattern independent of the azimuth angle. For real applications this might not be the case. It will depend on what type of antenna is chosen, or where the receiver and transmitter are mounted. In this case the influence of the antenna has to be modelled as a function of azimuth α and elevation angle θ.

$$I_{ANG} = f(\alpha, \theta) \tag{5.6}$$

As this thesis deals with localisation in 2 dimensions only, the dependency of the radiation pattern on the elevation angle θ is not considered further.

Figure 5.5 shows an example of a radiation pattern which is not independent of the azimuth angle. If the direction with minimum gain, α, min, is chosen as reference, it is possible to express the increase in gain as a function of the azimuth

$$\Delta I(\alpha) = I(\alpha) - I_{\alpha,min} \tag{5.7}$$

This dependency can again be modelled using a Fourier Series approximation, or in the case of discrete data-points, using a FFT.

$$\begin{aligned} I_{ANG,approx}(\alpha) &= \sum_{k=1}^{N} c_k e^{ik\omega\alpha} \\ & \omega = 2\pi f \qquad N \ldots \text{Number of sampling points} \end{aligned} \tag{5.8}$$

The complete sensor model

At this point it is necessary to combine all of the previously developed approximations into one model $I = f(x, \alpha)$, expressing intensity as a function of the distance x, and the azimuth angle α. Using the sum of equation (5.1), equation (5.3) and equation (5.7), it is possible to approximate the two-ray model as shown in figure 5.6. This sum constitutes the signal mean and is given by

$$\begin{aligned} I_M(x, \alpha) &= I_{FS,approx} + I_{REFL,approx} + I_{ANG,approx} \\ &= a\exp(bx) + c\exp(dx) + \sum_{k=1}^{N} c_k e^{ik\omega x} + \sum_{k=1}^{M} c_k e^{ik\omega\alpha} \\ & \omega = 2\pi f \qquad N, M \ldots \text{Number of Fourier coefficients} \end{aligned} \tag{5.9}$$

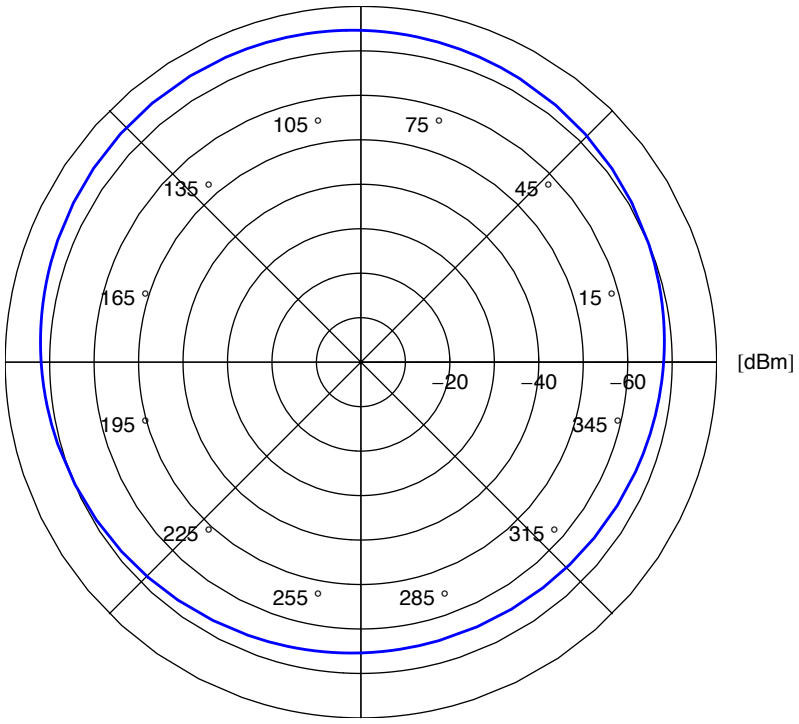

Figure 5.5: Simulation data for rotating receiver — The pattern shows the dependency on the azimuth angle α and exhibits a symmetry about the axis 120-300 degrees azimuth angle. Maximum gain due to the azimuth angle occurs in this figure at about 300 degrees.

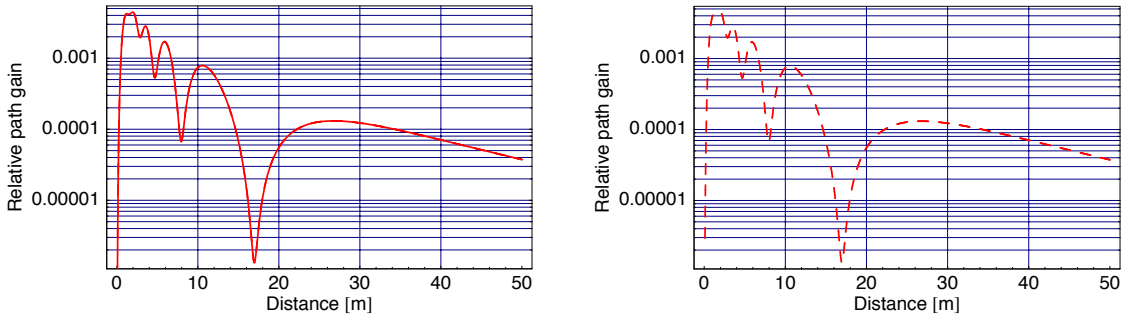

(a) (b)

Figure 5.6: Approximation of the two-ray model — (a) shows the two-ray model as obtained from the analytical formulation, while (b) shows the approximation of the mean for a single azimuth angle as obtained using equation (5.9).

With the addition of the signal variance (equation (5.3)), and the possibility that the variance is a function of the mean (equation (5.4)), the complete model can be written as

$$\lambda(x, z, \alpha) = \frac{1}{\sqrt{(2\pi)}\sigma(I_M(x,\alpha))} \exp\left(-\frac{(z - I_M(x,\alpha))^2}{2\sigma(I_M(x,\alpha))^2}\right) \tag{5.10}$$

This model is a joint distribution of x, z and α. For a given azimuth angle α, the function describing the behaviour of the sensor becomes a joint function of the distance x and the measurement z. It describes a surface in two dimensions, where at each point in the $x - z$ space the probability for any particular combination of distance and measurement pair is given by its function value. Although such functions are very often termed sensor Probability Density Function (PDF), they are in fact not, since the integral is not equal to one. More correctly, these functions are known as conditional sensor likelihood functions. An example of such a likelihood function is shown in figure 5.7.

The process of obtaining this function for a single azimuth angle is depicted in figure 5.8, where a selection of the Gaussian distributions obtained from the two-ray model, and centred at the mean, are shown. The two dimensional surface is created in the limit when the Gaussians for all distances are obtained.

Application of the model in localisation and tracking

For tracking applications, it is of interest to find the conditional likelihood function for the RF sensor. This function describes the likelihood of being at a certain distance and angle to the sensor given a measurement. The likelihood function is obtained from the model by instantiating the measurement obtained into the unconditional sensor likelihood function, and is given by

$$\lambda(x, \alpha \mid z = z_m) = \frac{1}{\sqrt{(2\pi)}\sigma(I_M(x,\alpha))} \exp\left(-\frac{(z_m - I_M(x,\alpha))^2}{2\sigma(I_M(x,\alpha))^2}\right) \tag{5.11}$$

The likelihood functions obtained through instantiation of a measurement z_m are now a function of the distance x and the azimuth angle α. They describe the likelihood of being at a particular distance and angle, when a particular measurement was made.

Figure 5.9 shows examples of likelihood functions as they are obtained for different measurements. The examples are based on the two-ray model and include the directional behaviour as previously presented. From these pictures it is also apparent that a lower measurement value does not automatically correspond to a larger distance. This can be observed when comparing the figures for the measurement of -35 dBm and -25 dBm. Although the measurement of -25 dBm is higher, there exists a likelihood for certain angles to be at a larger distance than for the measurement of -35 dBm.

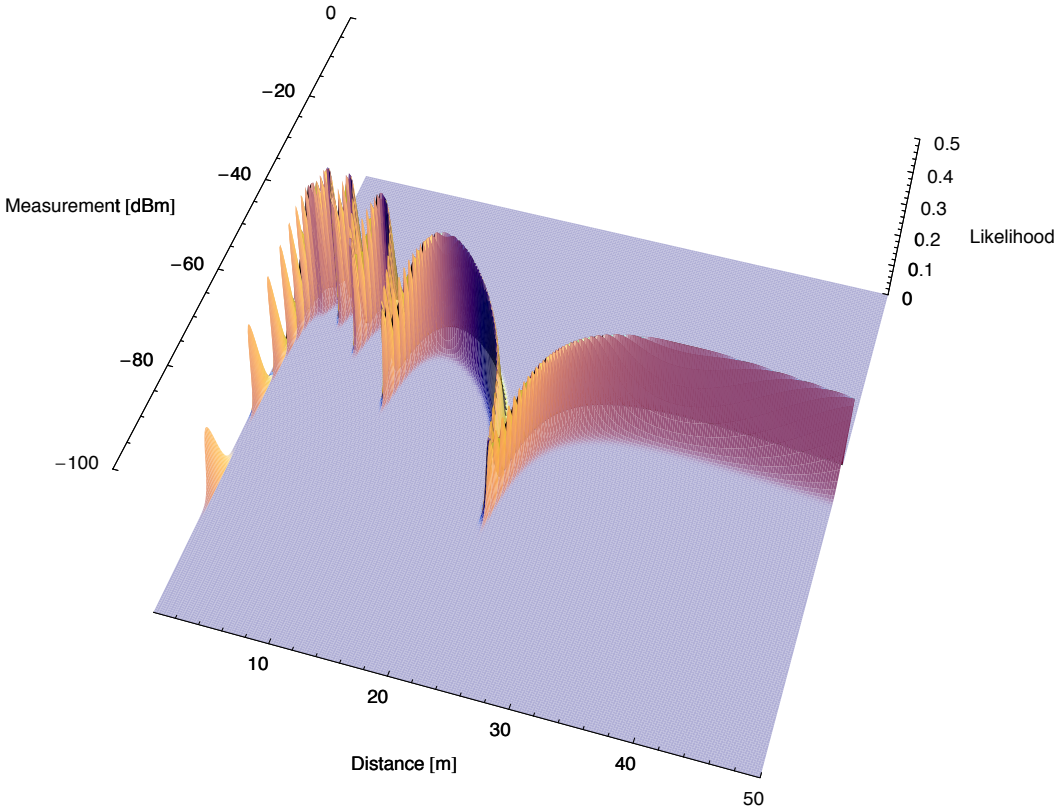

Figure 5.7: Sensor Likelihood Function (theoretical) — The sensor model based on the two-ray model is shown here for a given single azimuth angle for the frequency of 434 MHz, transmitter height of 3 m and receiver height of 2 m.

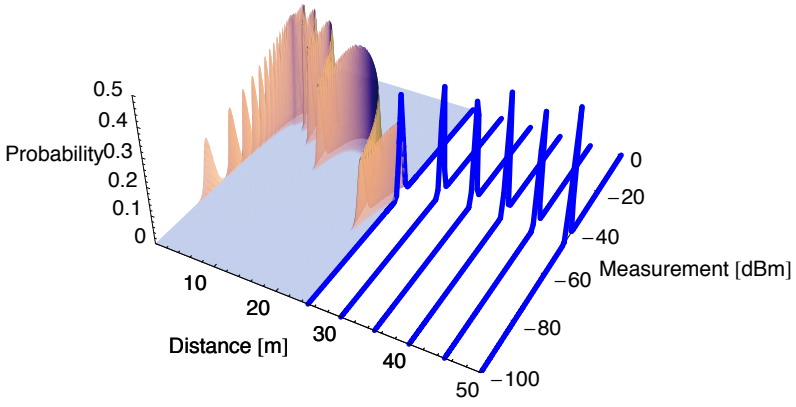

Figure 5.8: Sensor Likelihood Function (Conditional Probabilities) — Some of the Gaussians describing the uncertainty in the received signal for a given distance are shown in blue. When repeating the process of obtaining these Gaussians for all possible distances a two dimensional surface is created, which is shown for separation distances between 0 m and 25 m.

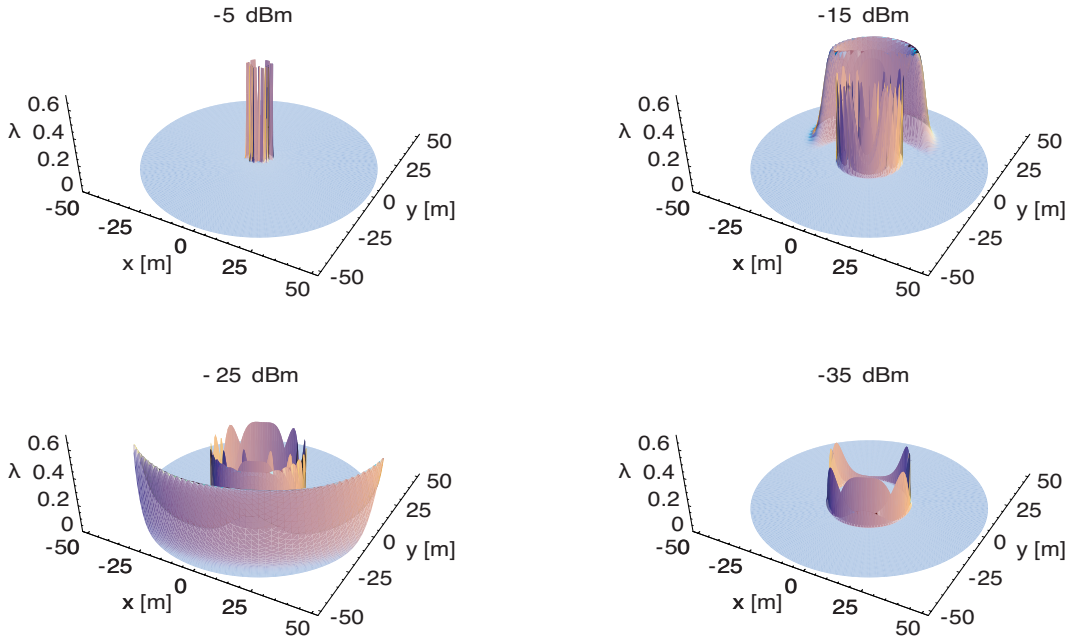

Figure 5.9: Conditional likelihood for a measurement considering distance and azimuth angle — Shown is the likelihood of being at a certain distance and angle given different measurements.The plots are in Cartesian $x-y$, rather than polar coordinates. The receiver is located at (0, 0).

Figure 5.10 shows graphically that assuming a fixed (given) single azimuth angle, these likelihood functions can be considered as cuts parallel to the distance axis through the sensor likelihood function plotted in $x-z$, i.e. the distance-measurement space. Depending on the underlying function I_M (equation (5.9)), the individual likelihood functions obtained might be non-Gaussian and multi-modal.

The likelihood functions shown here are the measurement inputs for the probabilistic filtering techniques described earlier. It is clear that the distributions are not appropriate for Kalman Filters (KFs) since they cannot be approximated by a single Gaussian.

Compression of the alternative sensor representation

The use of a Fourier Series based approximation to model the residual components, and the angular dependency, offers the possibility for compression of this model. The new sensor model described in this thesis can be thought of as a Fourier density [17, 18]. A real Fourier density consists only of Fourier components, does not have the additional exponential components used for the description of the free space loss, and satisfies the requirements for a probability density, i.e. integrates to one. The sensor description does not, though. Nevertheless, the sensor model has the potential to be compressed.

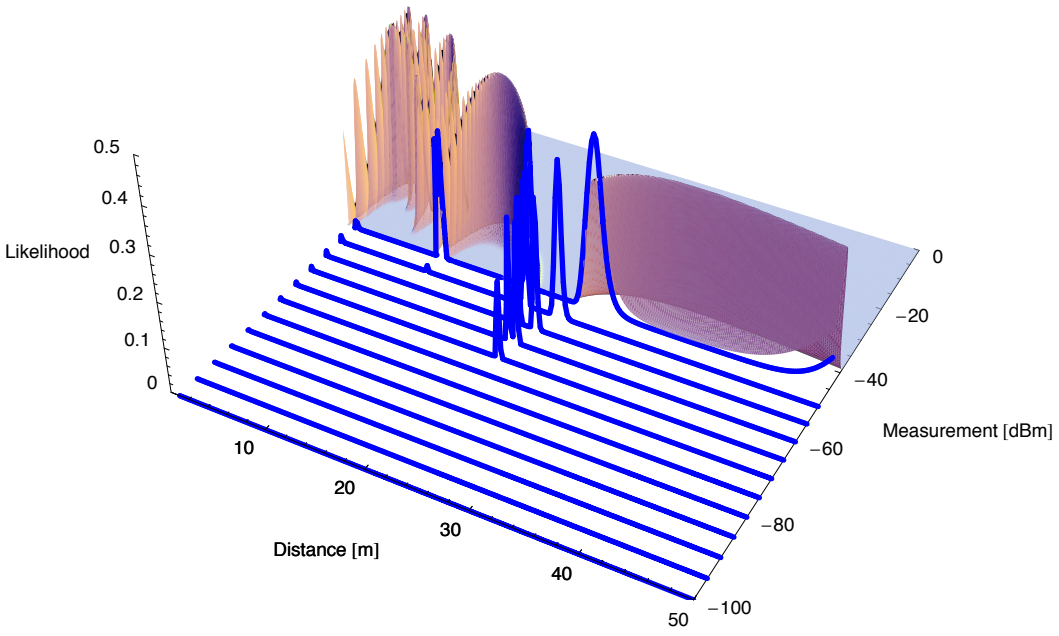

Figure 5.10: Sensor Likelihood Function (Conditional Likelihoods) — For a given measurement and azimuth angle, a likelihood function as function of the distance is obtained (blue).

The complex Fourier coefficients of the model, for both the distance and the angular dependency, are obtained through a FFT of given real data. This data should be sampled at a high spatial frequency to capture the fading characteristics of the environment. As a consequence a high number of coefficients may be obtained. For storage, computational and communication bandwidth reasons (when transmitting the sensor model or measurements) it may be desirable to reduce the number of coefficients and thus *compress* the model. It should also be noted that the sensor model obtained is a parametric model with a fixed number of parameters. These are the four parameters in the exponential approximation, and the Fourier coefficients in the distance and angular approximation.

It may seem natural to reduce the number of coefficients n_c based on their magnitude, i.e. to remove the components with smallest magnitude first. Distance measures like the Kullback-Leibler (KL) distance may provide a helpful tool for monitoring how the sensor density with reduced coefficients resembles the density with all coefficients. However, it is not guaranteed that a sensor model with a reduced set of coefficients will behave in the same way as the model with a full set of coefficients, even though the KL distance is small. A good sense of engineering is necessary, as well as practical tests and simulations to ensure that no unwanted effects occur.

For example, the KL distance between a sensor model with all coefficients and a set of sensor models with reduced number of coefficients is calculated. Models are constructed in such a way that the first model has the single coefficient with highest magnitude, the second model has the two coefficients with highest magnitude and so on. This results in $N-1$ models being created, where N is the maximum number of coefficients available. The resulting set of models is ordered based on the magnitude of the Fourier coefficients. Unfortunately the corresponding KL distances are not monotonically decreasing. This means that there will be situations where, although the number of Fourier coefficients is increased, the KL distance will also increase. Therefore, a simple thresholding solution will be difficult.

Apart from this behaviour, a question that remains to be answered is: how does the KL distance relate to what the sensor model physically represents? The KL distance does not relate at all to the physical sensor representation. It is just a measure of the difference between two densities. Nothing can be assumed about the physical implications of the difference between the two densities that are compared. Therefore, care must be taken to verify that a compressed density based on a reduced number of coefficients is still useful.

Other measures to compare densities, or how to order the Fourier coefficients may be better suited, but these are outside the scope of this thesis.

5.2.2 Comparison of the new Radio Frequency sensor model with the nth power law

This subsection compares the expected performance of the sensor model developed in this thesis to the commonly used nth power model. This comparison is based on the assumption that the new sensor description reflects the true signal propagation more accurately.

Inferring wrong distances with the nth power model

Neglecting the probabilistic aspects of both models, a comparison is made using the signal mean, which can be described deterministically.

When making a measurement, a sensor reading z is obtained, corresponding to the true distance x, based on the true sensor likelihood function. The true function is most likely not available, but for the purposes of demonstration it is represented by the two-ray model. The measurement in turn is interpreted using the nth power model, which leads us to an inferred distance $\hat{x}$.

An observation occurring at the true position x_1, and where the measurement value on the nth power model is lower than on the true model, will lead to an inferred distance $\hat{x}_1$ which is shorter than the true distance x_1. This situation is shown in figure 5.11, where the numbers $1-4$ indicate the order.

If, in contrast to the previous situation, the measurement value z_2 on the nth power model is higher than on the true model, then the inferred distance $\hat{x}_2$ will be larger than the true distance x_2. This situation is illustrated in figure 5.12.

It can now be seen that, except for the situations where the new model and the nth power model coincide, the measured distance will always differ from the true distance.

Model differences

When the probabilistic aspects are included and the variations in the signals are considered, it can often be found that the nth power model is applied with a signal variance that is too optimistic. The nth power model is generally obtained through a regression of data points. For given distances, multiple measurements are obtained. As the distribution of the measurements is Gaussian, the resulting fitted curve for the nth power model should be a curve going through the mean of the data points (or at least close to the mean). The Gaussian distribution depicting the variance should also be centred on this particular mean. This condition is necessary to accurately reflect the assumed *Gaussianity* of the signal. As the nth power model is found through a regression, the previously stated condition is not guaranteed and is in practical applications often violated. In contrast, the alternative representation guarantees that the conditional Gaussian distributions are centred on the mean of

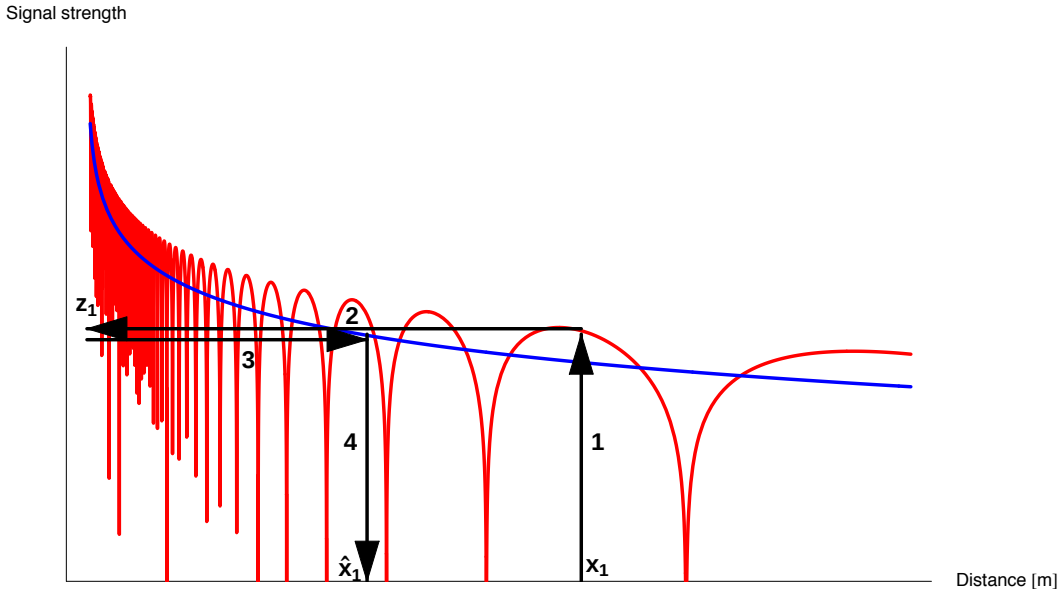

Figure 5.11: nth power model vs. new developed model — Inferring shorter distances: The true sensor likelihood is shown in red, while the nth power model is depicted in blue.

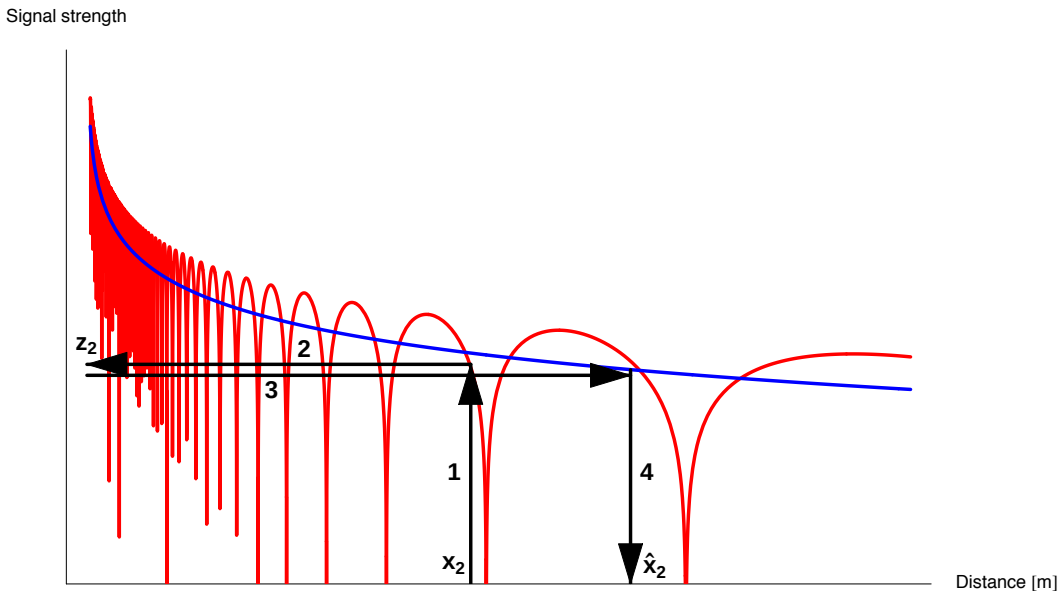

Figure 5.12: nth power model vs. new developed model — Inferring larger distances: The true sensor likelihood is shown in red, while the nth power model is depicted in blue.

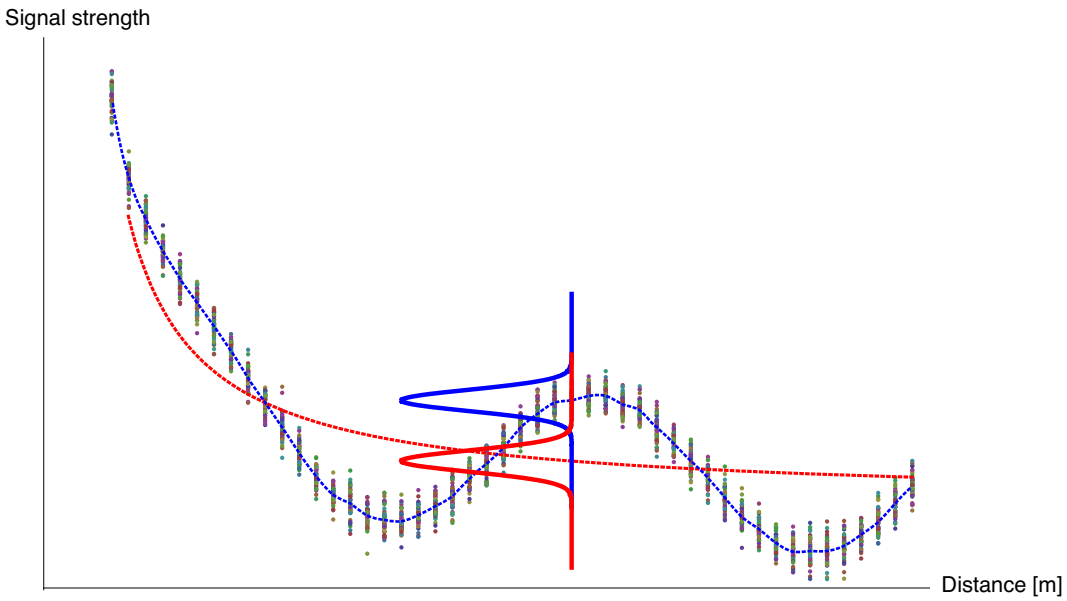

Figure 5.13: nth power model vs. new developed model (Model differences) — For given data points, which in this case here are equi-distant sampled, the distribution is of Gaussian form. The blue Gaussian distribution depicts the underlying sampled points for one distance. The signal mean function for all data points as it can be represented using the new model is shown in blue (dashed) and goes through the mean of the data points for the given sampling distances. For the nth power model this cannot be guaranteed, as this model is obtained through a regression. This behaviour is shown with the red (dashed) curve. Using the nth power model as the basis for modelling, with the same variance (as obtained from the data points) to represent the signal uncertainty, the resulting model differs significantly from the original distribution. This behaviour can be seen when comparing the red Gaussian curve (nth power model) with the blue Gaussian curve (new model). In this example the two distributions have almost zero overlap. Based on the assumption that the recorded data points represent the true position (which they should), this is an indication that the nth power model approximation is not good.

the data, at least for the case where all Fourier coefficients are used to represent the signal mean function. This situation is depicted in figure 5.13, and clearly underpins the importance of correctly representing the signal mean. Depending on the variance obtained from the data points, the overlap between the two Gaussian distributions as shown previously may vary. Clearly, when using the variance as obtained from the data points and *shifting* the Gaussian distribution, as the nth power model does in this case, an error is introduced. To some extent this can be compensated for by artificially increasing the variance for the nth power model so that it includes the data points. This increase still does not solve the problem that the data points are not correctly represented, but should alleviate to some extent the issues occurring from a complete mismatch indicated by almost non-overlapping Gaussian distributions.

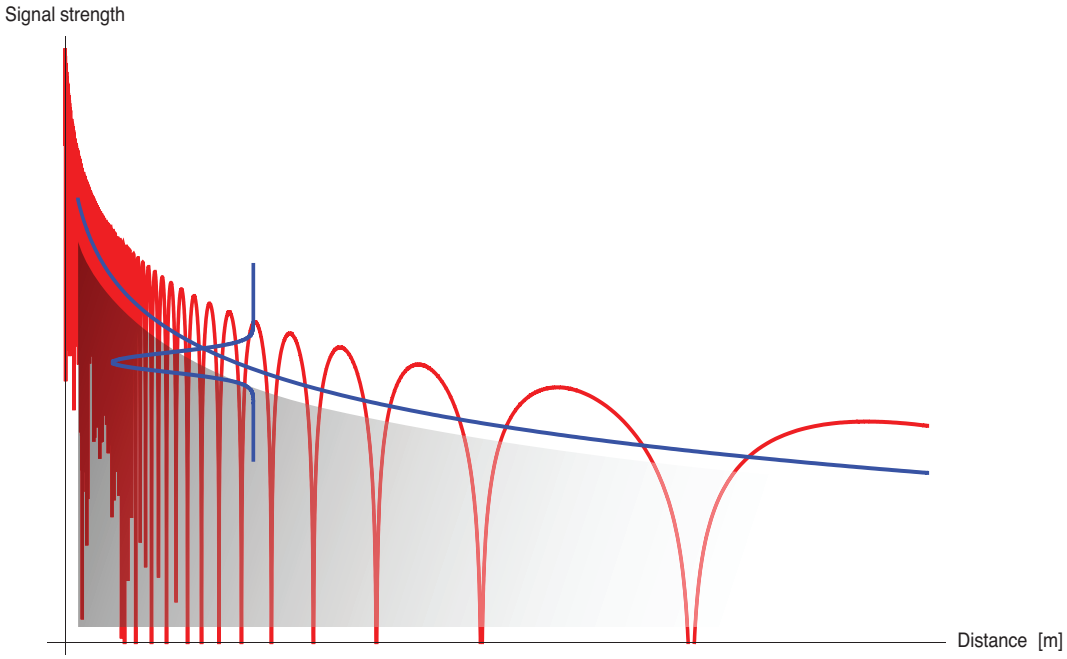

Figure 5.14: nth power model vs. new developed model (Model differences) — The two-ray model shows that for small distances there will be a probability to also measure small signal levels, which is indicated by the gray gradient area. The nth power model in contrast does not reflect this behaviour as indicated by the blue Gaussian curve.

It is important to note that the two-ray model, and the measurements clearly show that there is a possibility of having a very small signal value for close distances. This is not correctly modelled by the nth power model as there is almost no probability mass in this area due to an often too small chosen signal variance. Figure 5.14 illustrates this situation.

From the previous examples it can be seen that the nth power model will lead to incorrectly observed distances. The nth power model, as used in robotics applications, is found through regression. It can be assumed that the true function behaves more like the two-ray model, with oscillations. This in turn means that in many situations the nth power model based distance estimates are either shorter or longer than the true distance. Examples of these behaviours are shown in chapter 6.

5.3 Multi-sensor localisation

One of the goals of the CPS is to provide position information for agents in 2 dimensions. To resolve the bearing uncertainty inherent in range-only sensors it is necessary to have more than one sensor. Including more than one sensor requires sensor fusion, as the information from the different sensors has to be fused in a consistent manner. Data association must

also be solved, i.e. a certain measurement must be assigned to a certain agent (target) before it can be used to update this agent's position. This section addresses these issues, and also introduces a centralised tracking architecture suitable for the CPS. A qualitative analysis is conducted into the effects that occur if the distance to an agent is measured incorrectly by one or multiple sensors, and how this affects the overall localisation is shown. Finally, the inclusion of external range information to enhance the localisation process is discussed.

5.3.1 Range-only localisation in 2D

Range-only localisation attempts to estimate the position of a target by taking only range measurements to the target. Since no bearing to the target is available, the position estimate based on a single measurement has complete uncertainty about the bearing. Successive measurements taken from the same location are not able to resolve this issue. To obtain a bearing estimate and increase the localisation accuracy multiple measurements from spatially different sensor positions have to be taken.

In the context of probabilistic filtering the behaviour of sensors is encoded in sensor-likelihood functions. In the case of a range sensor this function shows, for a particular measurement, the likelihood distribution of the distance. The likelihood-function of a unimodal range measuring sensor has one distinct mode for the most probable distance. An example likelihood-function for a given measurement is shown in figure 5.15. Conversely, range estimates from RF-RSSI sensors are usually multi-modal, which means that for one measurement the sensors can provide more than one very likely distance to the target observed.

To obtain a position estimate with a single range-only sensor the sensor location must be varied for successive measurements, or multiple measurements from spatially distributed sensors have to be fused. In doing so, the spatial overlap of the obtained likelihood functions will most of the time resolve the initial ambiguities and refine the position estimate. This process is shown in figure 5.16, with the result plotted in figure 5.17. Sometimes, the ambiguities cannot be completely resolved as the result using only two combined measurements shows (figure 5.17). The probabilistic approach presented here can be seen as analog to deterministic triangulation.

5.3.2 Data association

One of the first steps for a tracking algorithm is, upon reception of a measurement, to assign this measurement to a target. This process is called *Data Association*. The correct match of the observation to the target generating the observation is of fundamental importance for

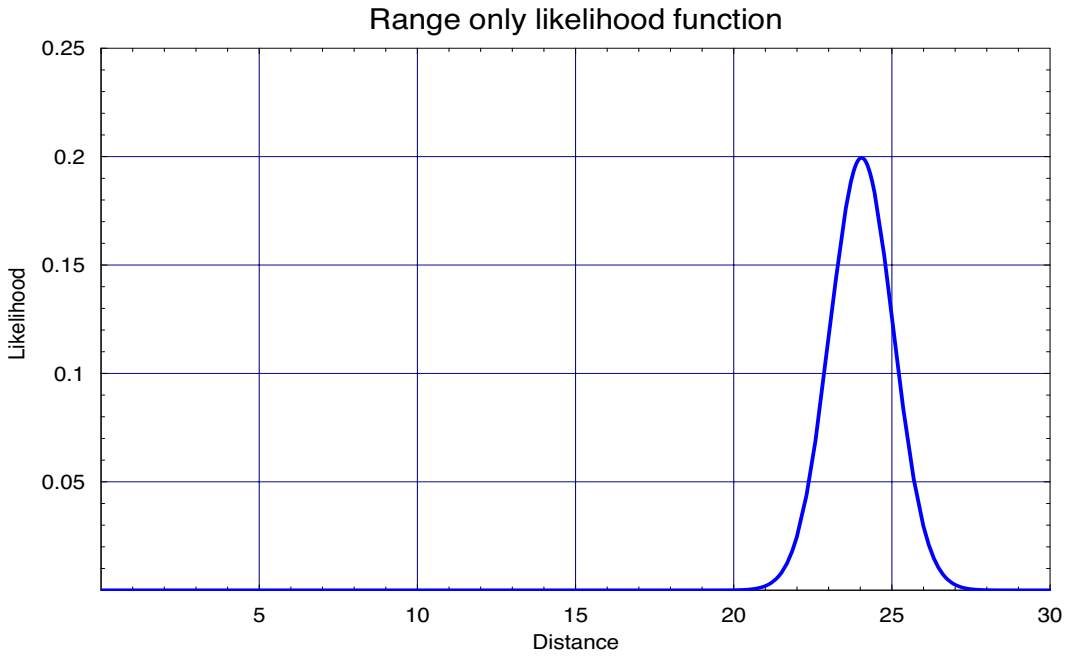

Figure 5.15: Uni-modal range-only sensor likelihood function — The (conditional) likelihood function of a uni-modal range-only sensor usually has one distinct mode denoting the most probable distance for a given observation. Note that although the shape of the likelihood function shown here is Gaussian, this is an example only. The shape for individual types of sensors may vary from the one shown here. Also note that the area under a likelihood-function does not have to integrate to one.

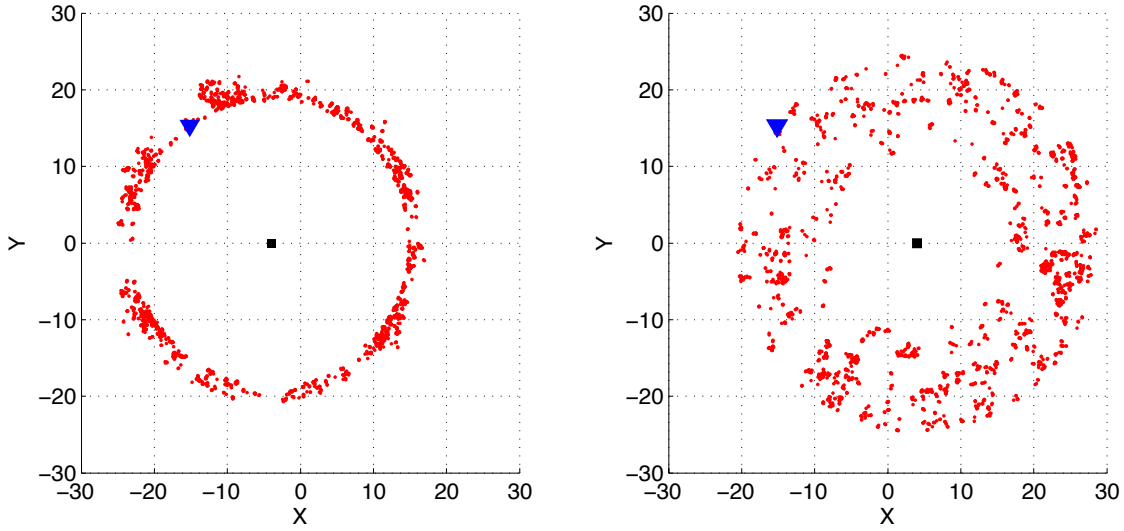

Figure 5.16: Uni-modal range-only sensor — (a) A range-only sensor located at (-4, 0) observes the distance to a target located at (-15, 15). As the sensor does not provide bearing, the possible locations form a circle around the sensors position with the observed distance as the radius. (b) A second range-only sensor located at (4, 0) observes the distance to the same target. Once again, as this sensor does not provide bearing, the possible locations form a circle around the sensors position with the observed distance as the radius. Note how the particles are more spread, as the second sensor is more uncertain about the measured distance when compared to the first sensor.

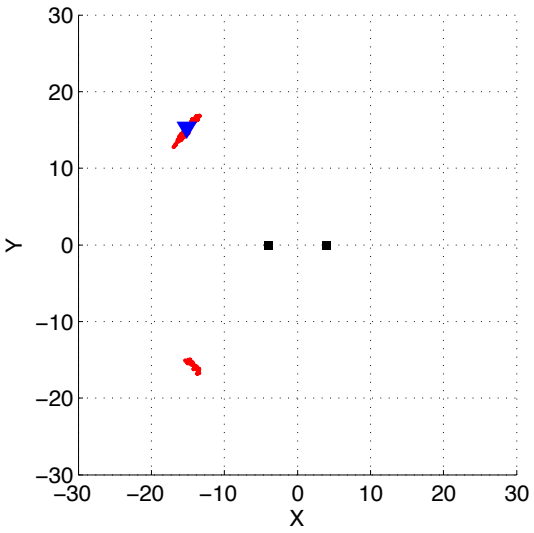

Figure 5.17: Uni-modal range-only sensors fused — The combination of both observations as shown in figure 5.16(a) and (b) eliminates the majority of the sensor-individual hypotheses. As only two sensors are used a mirrored cloud of particles forms at around (-15,-15) in addition to the true location at (-15, 15). The use of a third sensor, not co-linearly aligned with the other ones would most likely eliminate the second hypothesis.

any tracking algorithm. An observation that has been incorrectly assigned to a target may lead to a catastrophic failure of the filter, i.e. the filter might diverge from the true state and never recover. For RF sensors data association is most often trivial to solve as, for the majority of transceivers, each measurement comes with an identifier of the agent that has generated the observation. In such a case the only issue to be careful of, is to ensure that there are no two agents using the same identifier.

If the messages do not come with an identifier, more sophisticated data association algorithms have to be used. Gating techniques, for example using the Mahalanobis distance, are a common tool for data association, but these techniques are not further discussed as they are not relevant in the scope of the thesis.

5.3.3 Data fusion

> *Data fusion* is the process of combining information from a number of different sources to provide a robust and complete description of an environment or process of interest [28].

To obtain information about the state of interest, it is often necessary to fuse data coming from different sensors. This is especially true for range-only localisation, where multiple range-only sensors may observe a target at the same time from different positions. Different architectures are possible depending on the computational capabilities of the sensor nodes. Sensor nodes can transmit the measurement directly to a central computer, or they can calculate and transmit their own belief about the state of interest. A fully centralised architecture is used for the CPS, as the computational power in the sensor nodes themselves is not sufficient to calculate a belief about the state of interest.

Fully centralised architecture

This architecture is based on the independent likelihood pool. The assumption here is that given the true state $\mathbf{x}$, the information obtained from the ith sensor (information source) is independent of the information obtained from other sources. Therefore it is possible to obtain an update for the state by multiplying the prior distribution of the state with the individual likelihoods from each sensor and normalizing using the marginal distribution $p(\mathbf{Z}^n)$. This is described in equation (5.12) and is known as the *independent likelihood pool* [28]. Using this equation it is possible to combine the measurements from any number of sensors in a simple manner. A graphic representation can be seen in figure 5.18.

The conditional probabilities $p(\mathbf{z}_i|\mathbf{x})$ are known a priori as functions of both $\mathbf{z}_i$ and $\mathbf{x}$. These are simply the sensor models. Upon reception of a measurement, a likelihood-function $\lambda(\mathbf{x})$ can be constructed which now depends only on the state. This procedure is done for

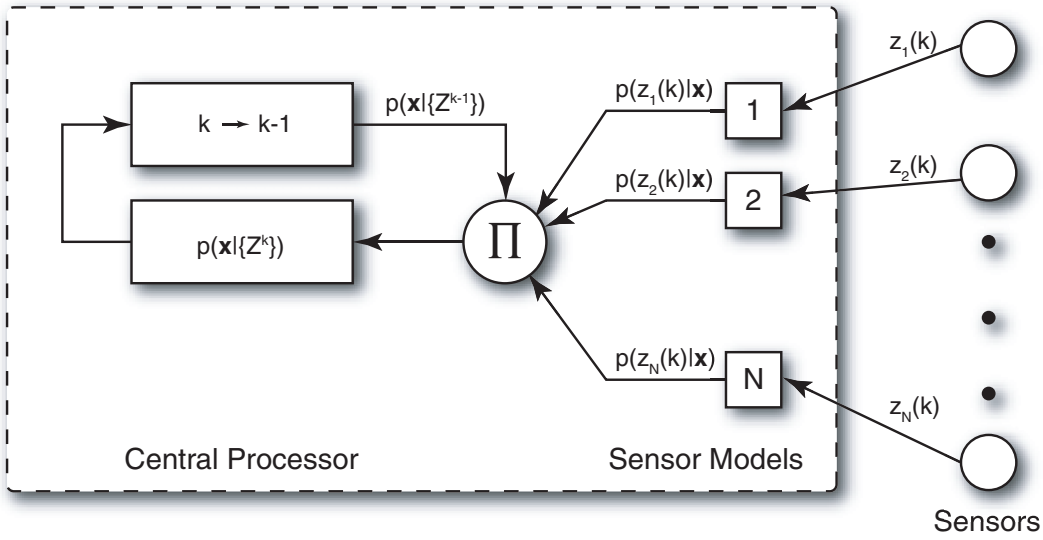

Figure 5.18: Fully centralised sensor fusion system — Adapted from [28] where a form using log-likelihoods is presented.

each individual sensor i. Taking the product of the likelihood functions for all sensors with the prior distribution for the state and normalising appropriately, the posterior distribution $p(\mathbf{x}|\mathbf{Z}^n)$ is obtained. This posterior describes the probability distribution of the state for the specific observation sequence used to obtain the sensor likelihoods.

$$p(\mathbf{x}|\mathbf{Z}^n) = [p(\mathbf{Z}^n)]^{-1} p(\mathbf{x}) \prod_{i=1}^{n} p(\mathbf{z}_i|\mathbf{x}) \tag{5.12}$$

Equation (5.12) is formulated to require that all past observations are stored, and that upon arrival of a new measurement the new total likelihood is recomputed using all measurements, including the latest measurement. Bayes theorem allows us to modify equation (5.12) so that new observations can be incorporated in an incremental manner, therefore avoiding the need to store all past observations. This is described by equation (5.13).

$$p(\mathbf{x}|\mathbf{Z}^k) = \frac{p(\mathbf{z}_k|\mathbf{x}) p(\mathbf{x}|\mathbf{Z}^{k-1})}{p(\mathbf{z}_k|\mathbf{Z}^{k-1})} \tag{5.13}$$

Now it is only necessary to compute and store the posterior distribution $p(\mathbf{x}_k|\mathbf{Z}^{k-1})$. This posterior serves as the new prior when new observations arrive. Using Bayes theorem the new observation $p(\mathbf{z}_k|\mathbf{x})$ is incorporated according to equation (5.13).

5.3.4 The effect of ranging errors and the baseline

In multi-sensor range-only localisation there are two important factors influencing the localisation process. Firstly, the accuracy of the sensor itself, i.e. the degree to which the sensor is capable of providing an accurate measurement. Ranging errors occur if the sensor is not able to measure accurately and therefore the estimated distance also becomes erroneous.

The second factor relates to the geometrical arrangement of the sensors. The geometrical sensor setup is often limited by the practical application. This situation might occur in a CPS installation where the sensors on the truck cannot be mounted as far apart as desired, but instead must be mounted close together for practical reasons. The baseline or the distance between different range-only sensors observing the same target, in relation to the distance observed, influences localisation accuracy. If the target is located perpendicular to the baseline, the ranging errors exhibit a different effect when compared to a target located co-linearly to the baseline.

This subsection examines the effect of these two factors on localisation accuracy when using range-only sensors. It does so in a qualitative manner and describes the effects occurring when choosing a particularly small baseline. For illustration purposes two range-only sensors are considered.
The following six cases are examined when using two range sensors for observations:

- Both sensors observe the target with correct measurements in a perpendicular arrangement. This case leads to a correctly evaluated position.
- One sensor observes the target with correct measurements in a perpendicular arrangement, while the second one observes a too large/too small range. This case leads to a shift in position.
- One sensor observes the target with correct measurements in a perpendicular arrangement, while the second one observes a too large/too small range. In this case the baseline is chosen to be smaller compared to the previous case. This case leads to a greater shift in position as in the case of a larger baseline.
- Both sensors observe the target with correct measurements in a co-linear arrangement. This case leads to a correctly evaluated position.
- One sensor observes the target with correct measurements in a co-linear arrangement, while the second one observes a too small range. This case leads to the split of the previously single hypothesis into two.

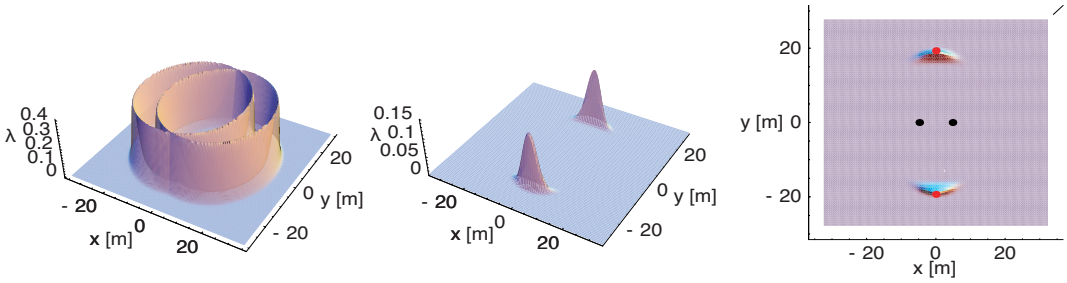

Figure 5.19: Range-only localisation with perpendicular sensor arrangement and correct measurements — (left) Individual sensor likelihood functions for the two sensors shown superimposed. (middle) Resulting combined likelihood function. (right) Resulting combined likelihood function with sensor positions (black dots) and true positions (red dots).

- One sensor observes the target with correct measurements in a co-linear arrangement, while the second one observes a too large range. This leads to the weakening of the true hypothesis.

The following examples assume that the sensors have uni-modal sensor likelihood functions. This is mostly not true in a real application as already seen previously for RF sensors, but this simplification allows the influences of ranging errors and the baseline to be seen. Other factors, such as previously described bias, are also not considered here.

Perpendicular sensor arrangement with correct measurements

In the case of both sensors observing the target and measuring the correct range, the true position will be recovered. As figure 5.19 shows, in a two sensor scenario it may be not possible to eliminate secondary hypotheses.

Perpendicular sensor arrangement with one sensor observing an incorrect range (large baseline)

If one of the sensors observes an incorrect range, the combined likelihood of both sensors will lead to a shifted position estimate. The following situation is depicted in figure 5.20. The left sensor observes a range that is too large. Alternatively, this figure may also describe the situation where the right sensor measures a too short range. It can be seen that the shift tends to be away from the sensor with a range estimate that is too large, or alternatively, towards the sensor with the range observation that is too small.

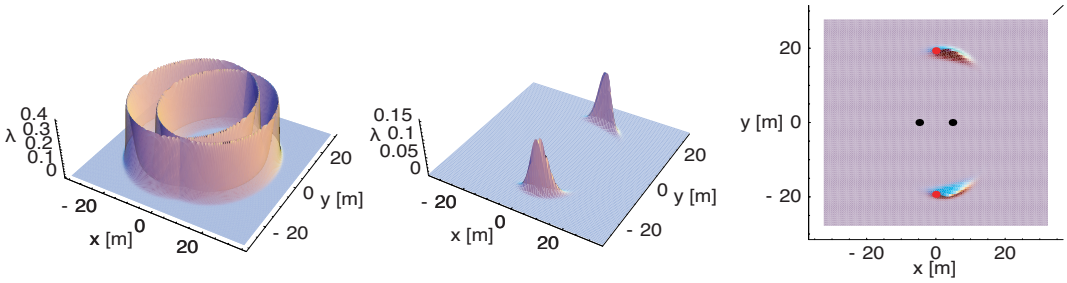

Figure 5.20: Range-only localisation with perpendicular sensor arrangement and one sensor observing an incorrect range (large baseline) — (left) Individual sensor likelihood functions for the two sensors shown superimposed. (middle) Resulting combined likelihood function. (right) Resulting combined likelihood function with sensor positions (black dots) and true positions (red dots).

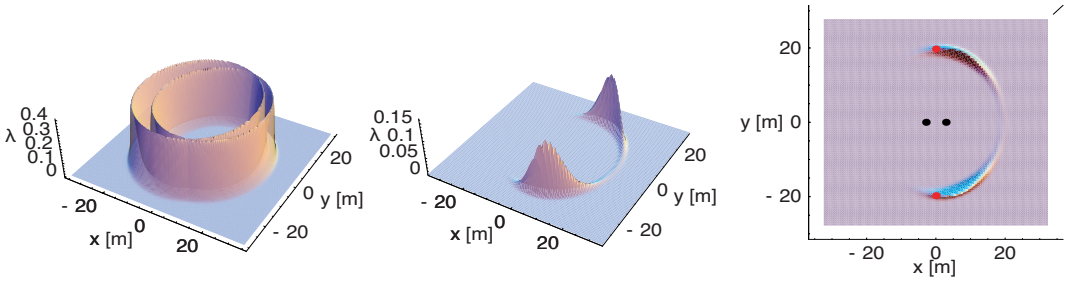

Figure 5.21: Range-only localisation with perpendicular sensor arrangement and one sensor observing an incorrect range (small baseline) — (left) Individual sensor likelihood functions for the two sensors shown superimposed. (middle) Resulting combined likelihood function. (right) Resulting combined likelihood function with sensor positions (black dots) and true positions (red dots).

Perpendicular sensor arrangement with one sensor observing an incorrect range (small baseline)

The effects in this case are the same as for the previous case. As figure 5.21 shows however, the effects are now more pronounced and the shift occurring is larger compared to the case with the larger baseline. The error in the range estimate was taken to be the same as in the previous example. In the case shown here it is now also possible to see that the hypotheses are becoming more spread out and are not as pronounced as in the previous case. This is not a direct effect of the baseline, but occurs because the combination of baseline and range measurements used in this example leads to solutions that are located towards the baseline axis. As the next example shows more clearly, such arrangements lead to spread likelihood functions.

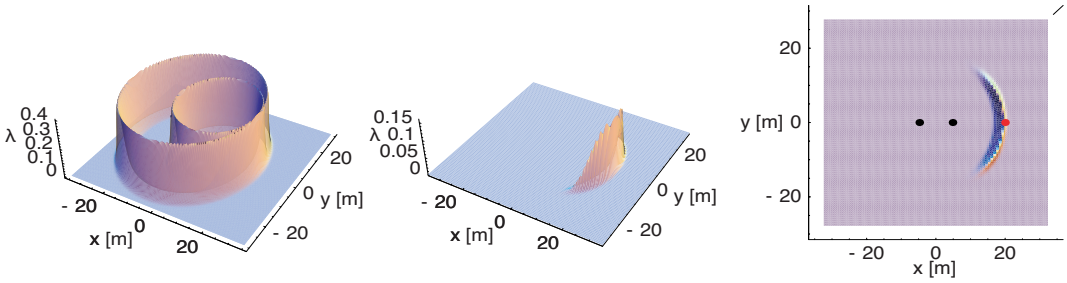

Figure 5.22: Range-only localisation with co-linear sensor arrangement and correct measurements — (left) Individual sensor likelihood functions for the two sensors shown superimposed. (middle) Resulting combined likelihood function. (right) Resulting combined likelihood function with sensor positions (black dots) and true position (red dot).

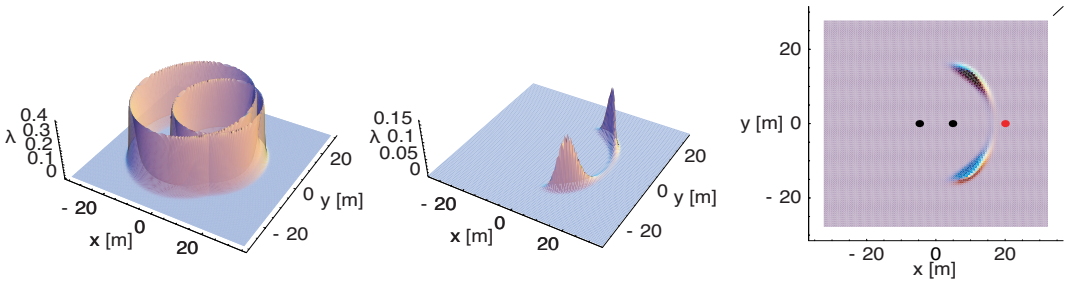

Figure 5.23: Range-only localisation with co-linear sensor arrangement and incorrect measurements (hypothesis split) — (left) Individual sensor likelihood functions for the two sensors shown superimposed. (middle) Resulting combined likelihood function (right) Resulting combined likelihood function with sensor positions (black dots) and true position (red dot).

Co-linear sensor arrangement with correct measurements

Figure 5.22 shows how a co-linearly aligned solution can occur. As opposed to the perpendicular case, it is possible to eliminate the secondary hypothesis when using only two sensors, as this case shows. This solution is highly unstable as the next two examples illustrate.

Co-linear sensor arrangement with incorrect measurements (hypothesis split)

In figure 5.23, the situation of a previously single hypothesis being split into two distinct hypotheses is shown. This error occurs when the left sensor observes a shorter distance, or the right sensor observes a larger distance. In the case shown in the figure, the left sensor observes a shorter distance. This leads to an overall combined likelihood function which has a distance smaller than the true distance, and can be seen by the fact that some part of the likelihood function lies between the right sensor and the true position.

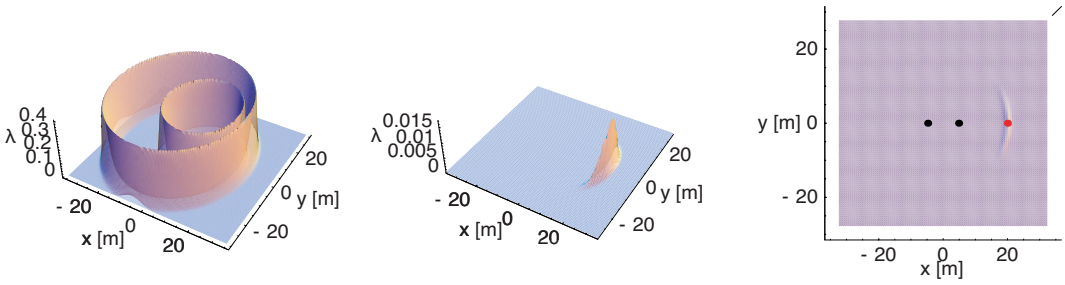

Figure 5.24: Range-only localisation with co-linear sensor arrangement and incorrect measurements (hypothesis weakening) — (left) Individual sensor likelihood functions for the two sensors shown superimposed. (middle) Resulting combined likelihood function (right) Resulting combined likelihood function with sensor positions (black dots) and true position (red dot).

Co-linear sensor arrangement with incorrect measurements (hypothesis weakening)

The opposite of the previous situation is shown in figure 5.24. The resulting combined likelihood function is not as pronounced when compared to the case with the correct measurements (see figure 5.22). Furthermore, this case will lead to a radial shift in the estimated position. Although not so pronounced and hardly visible, the position has shifted away from the sensors, as the left sensor observed a larger range in this example. As a general rule for this case, the resulting position will lie on the sensor axis between the two points of the two individual measurements where the distance is the smallest.

To summarise the effects shown in the examples, it can be seen that range-only localisation is susceptible to errors in the range. Such errors will lead to a shift of the estimated position. In some cases, depending on the geometrical arrangement, a previous single hypothesis could be split into two hypotheses. The length of the baseline plays an important role, and larger baselines are preferred as they counteract the errors that occur due to incorrectly measured ranges. The examples presented show the extreme cases where the target was either located perpendicular to the baseline or co-linear. It is understood that there will be an infinite number of possible locations, but the effects can be extrapolated.

So far only errors in the estimated range have been considered, and it was assumed that by using the likelihood function of a measurement, the true position could be correctly retrieved. It is clear that the uncertainty will increase when sensor bias is also considered.

5.3.5 Using external information

The application of a CPS as presented in chapter 2 is well suited to expand the presented range-only localisation approach with additional information obtained from external observers. So far, the localisation approach has been based on having a centralised system with multiple fixed RF readers, and it was assumed that they would all be attached to the same truck in the CPS application. The localisation was performed relative to the truck's position. The reader is reminded that the trucks and light vehicles can have absolute position information from Global Positioning System (GPS) sensors.

Now given that under certain circumstances it will be possible for trucks to exchange information about the measurements they obtain, and given that two or more trucks may observe the same agent, the localisation approach can be extended to include the measurement observation obtained from the external observer.

Figure 5.25 illustrates a configuration where two independent tracking systems make measurements about a common state of interest, i.e. the two trucks independently measure the position of the same agent.
The two independent observers can use the information in three different ways:

- Raw measurements can be exchanged between the two systems: To be able to use the raw measurements, the unconditional (joint) sensor likelihood functions corresponding to the sensor with which the measurement were taken must be available. This means that either this additional information has to be communicated as well, or it must be present a priori, for example stored in a database in each individual system.

- Conditional sensor likelihood functions describing the likelihood given the measurement can be exchanged: To overcome the possible need to communicate the unconditional joint sensor likelihood, the system that generated the measurement may instantiate it and thus create the conditional sensor likelihood given the measurement. This conditional measurement is then communicated, using possibly less data for communication than in the previous case.

- Exchange of combined conditional sensor likelihood functions: Upon simultaneous reception of multiple measurements from multiple sensors of one tracking system, the measurements are instantiated and the resulting combined conditional sensor likelihood function is then communicated.

In either case the systems need to communicate their own position, so that the transmitted information can be incorporated in a consistent manner. The task consists of transforming the measurement taken by the external observer ($\mathbf{z}_{meas,extobs}$), into a *virtual* measurement

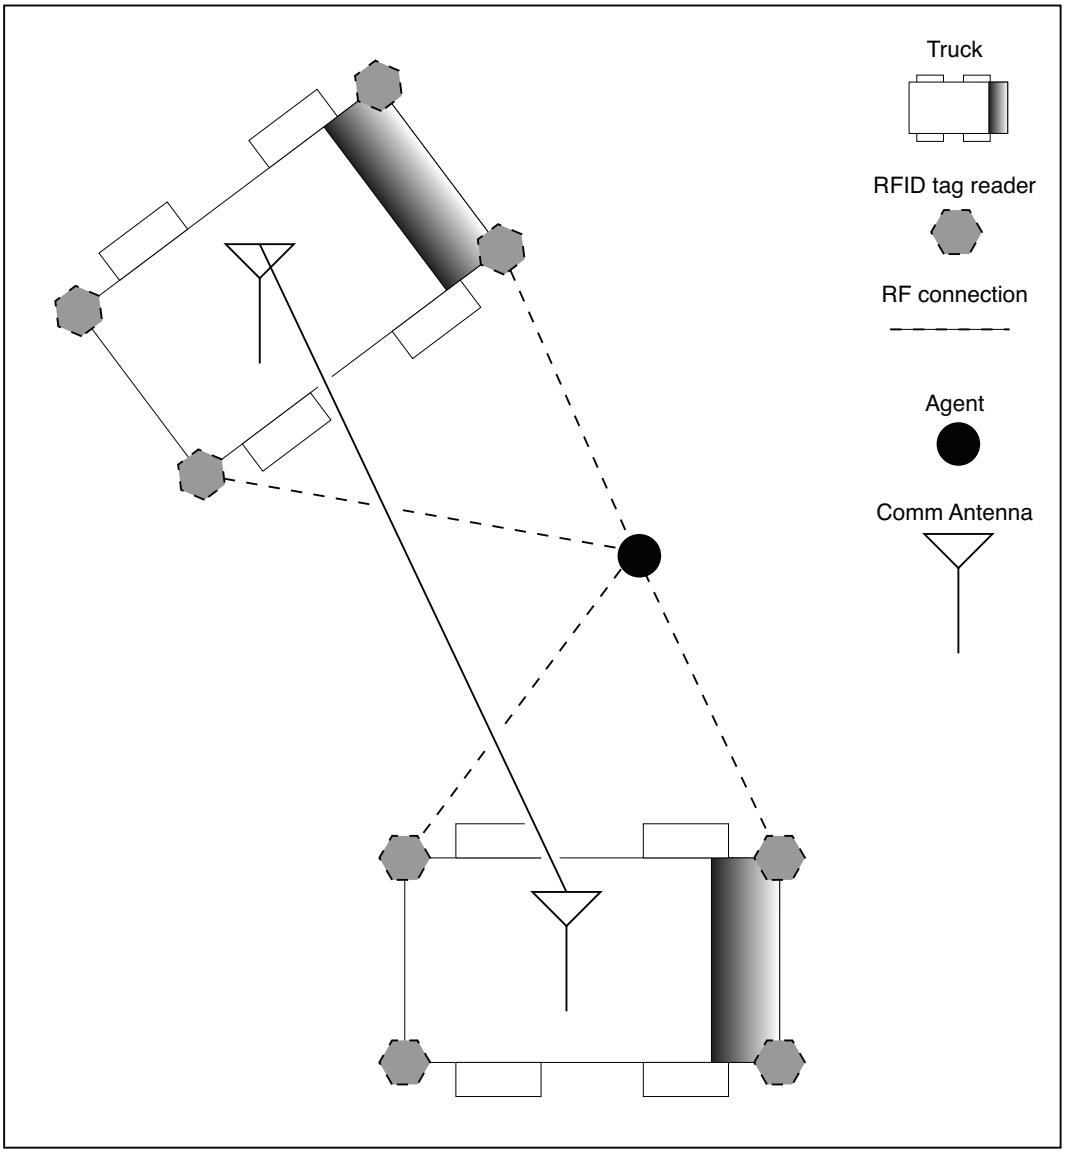

Figure 5.25: External observer — Two trucks observe the same agent. They can exchange information through a secondary communication link, as shown here, or through the RF tag sensor network as possibly formed by suitable RFID technlogy.

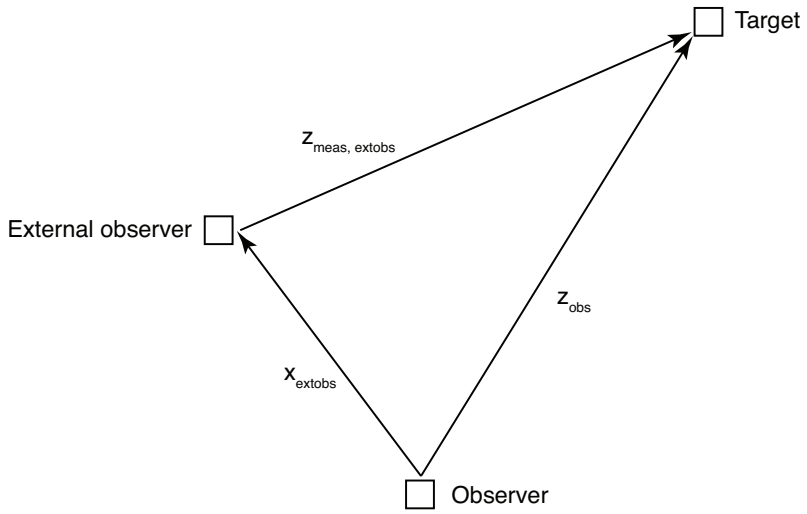

Figure 5.26: Vector addition — Transforming an external measurement. A *virtual* measurement can be created if knowledge about the position of the external is present and the measurement taken by the external observer is communicated and available.

($\mathbf{z}_{obs}$), appearing to be taken by the observer. This is calculated by adding the vectors that describe the state of the external observer ($\mathbf{x}_{extobs}$), and the measurement taken by this observer. This is shown in figure 5.26.

$$\mathbf{z}_{obs} = \mathbf{z}_{meas,extobs} + \mathbf{x}_{extobs} \tag{5.14}$$

The corresponding PDF for the virtual measurement is then obtained through a convolution of the PDFs.

$$p(\mathbf{z}_{obs}) = p(\mathbf{z}_{meas,extobs}) * p(\mathbf{x}_{extobs}) \tag{5.15}$$

Subsection 4.3.2 showed that a convolution has a smoothing effect. This means that the quality of a virtual measurement is heavily influenced by the quality of the position information of the external observer. The following example shows the basic principle to generate such a virtual measurement. It also highlights the effect of the quality of the position information of the external observer.

Example — Creating a virtual measurement

Figure 5.27 shows the process of obtaining a *virtual measurement*. An observer (black dot) tracks the position of a target (red dot). The same target is also seen by another external observer (blue dot). This external observer can communicate its observations

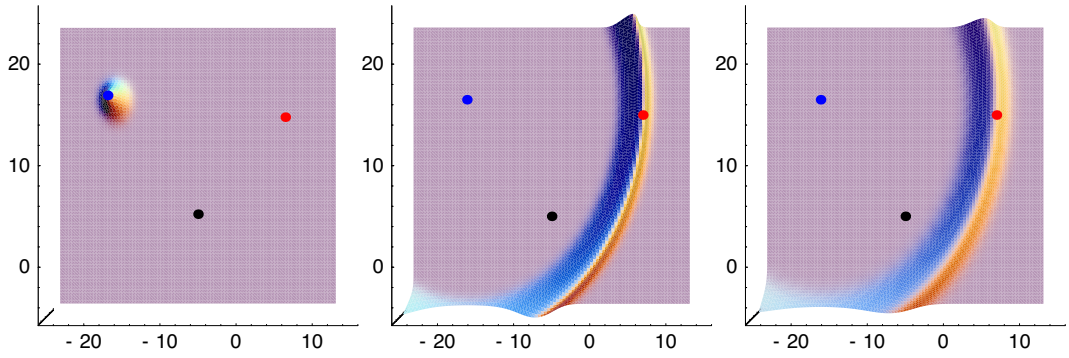

Figure 5.27: External observation — The position of the external observer (blue dot) is known with corresponding uncertainty (left). The external observer observes the target (red dot) and generates a conditional sensor likelihood function (middle). Another observer (black dot) can generate a *virtual* observation (right) by combining the knowledge about the position of the external observer and the measurement likelihood function obtained from the external observer. This virtual observation can then be incorporated into its belief about the state of the target, in the same way as for any other measurement.

to the first observer. To obtain a virtual measurement it needs to know the position of the external observer. If the external observer now communicates its observation, it is possible to construct a virtual measurement by convolving the state PDF corresponding to the external observer's position with the sensor likelihood function obtained from the measurement of the external observer. The resulting measurement can then be incorporated using the same method as for the other measurements. The virtual measurement in figure 5.27 is very similar to the measurement taken by the external observer. This is due to the fact that the position information of the external observer was of good quality.

The quality of the position information regarding the external observer strongly determines the quality of the resulting virtual observation, as figure 5.28 shows. In this example the position uncertainty of the external observer (blue dot) is assumed to be higher than in figure 5.27, which can be seen in the spread of the state PDF of the external observer. The observation likelihood function obtained by the external observer, is exactly the same as for the previous situation, but the resulting virtual measurement is now not as pronounced as for the previous situation.

This example illustrates how important it is to have an external observer with good position information. Using only visual comparison of the resulting two virtual measurement likelihood functions, it can be seen how the quality of the position information influences the final observation likelihood function. This observation also translates into the tracking behaviour when such virtual observations are incorporated. Better information about the position of the external observer leads to improved virtual measurements.

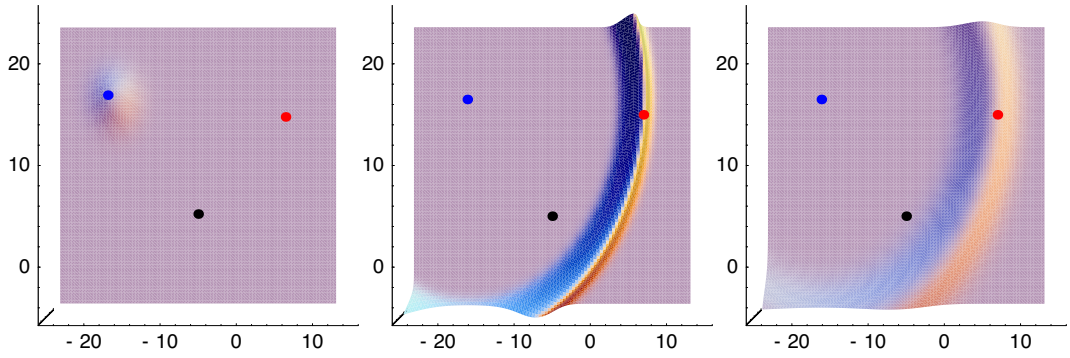

Figure 5.28: External observation with decreased position accuracy — The position of the external observer (blue dot) is now known with less certainty (left) as compared to the example in figure 5.27. The external observer makes the same observation about the target (red dot) and generates exactly the same conditional sensor likelihood function (middle) as in figure 5.27. Combining the knowledge about the position of the external observer and the measurement likelihood function the other observer (black dot) can generate a *virtual* observation (right) which is of lesser quality than in figure 5.27.

5.4 Bayesian Decision Theory for model selection

Using the alternative RF sensor model description, it is possible to generate sensor models based on real data that are very specific to the environment from which the data were taken. Having such a specific model will provide higher localisation accuracy, as the variance of the data used in the model generation process should be lower. The truck in the CPS application, or in general the mobile agent/robot, will not be stationary. Instead, it will be moving around the environment. It is expected that an RF model obtained in one area will not necessarily perform as well in a different environment. To overcome this problem, a possible approach would be to have a multitude of sensor models, each for a specific region, environment or environmental condition, and to choose a suitable model depending on the circumstances. In the process of characterising a sensor, a set of different sensor likelihood functions might be obtained for the sensor, each appropriate for certain operating conditions.

Figure 5.29 shows an example of two different sensor likelihood functions, which in this case, are obtained from two different sensors mounted at different heights. For any localisation and tracking application it is crucial that the correct sensor likelihood-function is applied in the filtering process so that the best position estimate can be obtained (see subsection 4.4.1).

This section shows how under certain circumstances the correct RF sensor model to be used, can be identified on-line through the application of BDT. Given external observations and a set of sensor measurements, the task consists of identifying the correct or best model.

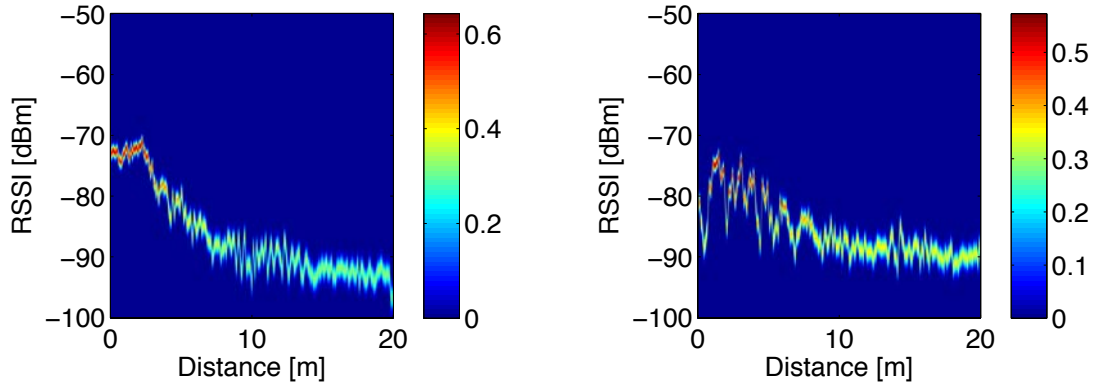

Figure 5.29: Two different sensor likelihood functions — These two likelihood functions correspond to (a) condition 1 and (b) condition 2, which for example could be obtained from different mounting positions or on dry/wet soil.

This procedure should be performed while the localisation filter is in operation. If the external observation by itself completely determines the choice of the correct model, i.e. based on the external observation only one sensor-model is applicable, this task is trivial. Often this will not be the case, and it will be necessary to use the actual sensor measurements to aid in the identification process.

5.4.1 Bayesian decision theory for two and multiple sensor models

Assume that a sensor can operate under two conditions, and each of these conditions requires the use of a different sensor likelihood function. The set of possible sensor models consists in this case of two models, which can be written as $\mathcal{M} = \{M_1, M_2\}$. A measurement z is now taken, and based on this measurement a decision about which sensor model is more suitable may be obtained by using the following decision rule:

Choose model M_1 if

$$\mathcal{P}(M_1|z) > \mathcal{P}(M_2|z) \tag{5.16}$$

and model M_2 otherwise.

The conditional probability $\mathcal{P}(M_i|z)$ for the i-th sensor may be obtained using Bayes theorem. In this case, the conditional probability is equivalent to the posterior probability obtained using Bayes theorem, when a prior assumption about the probability of the i-th model $\mathcal{P}(M_i)$, and a conditional probability distribution $p(z|M_i)$ describing the probability of measurements given the i-th model are available. If subsequently, more measurements become available, it is possible to refine the estimates about the probabilities by treating the conditional probabilities $\mathcal{P}(M_i|z)$ as new priors and recursively estimating new posteriors.

The most difficult problem is to obtain the conditional probability distribution $p(z|M_i)$ as the models M_i are a function of other variables ϕ, or eventually of the state vector $\mathbf{x}$. To be precise, the conditional probability distribution should be written as $p(z|\phi, \mathbf{x})$. It now becomes apparent that knowledge of the variables ϕ and the state $\mathbf{x}$ must be present to be able to evaluate the conditional probability distribution $p(z|M_i)$. This can be interpreted as: given the model M_i, what is the probability of making the observation z?

As the choice of the model is based on maximising the posterior probability, the above stated decision rule for two sensors can be easily extended to the multiple sensors case. In the multiple models case, the i-th model is chosen if

$$\mathcal{P}(M_i|z) > \mathcal{P}(M_j|z) \qquad \text{for all } j \neq i \tag{5.17}$$

holds true.

5.4.2 Decision regions and decision boundaries for two and multiple sensor models

The boundary where the conditional sensor probability distributions are equal, as given by equation (5.18), is called the decision boundary. This is the boundary where the decision in favour of one model or against is taken.

$$\mathcal{P}(M_i|z) = \mathcal{P}(M_j|z) \qquad \text{for all } j \neq i \tag{5.18}$$

The areas on either side of a boundary are called decision regions and describe the areas or regions where a decision in favour of a specific model will be taken. An example of such a situation as it might occur for three possible RF sensor model candidates is shown in figure 5.30.

As mentioned previously, one main difficulty is to obtain the conditional probability distribution $p(z|M_i)$. In the case of the RF sensor under consideration here, the conditional probability distribution must be written as $\mathcal{P}(z|M_i, d, \alpha)$ to indicate that the measurement z depends on both the model M_i, the distance d and the azimuth angle α between sender and receiver. The conditional probability distributions $\mathcal{P}(z|M_i, d, \alpha)$ can be obtained after an external observation provides the distance d and the azimuth angle α.

Figure 5.31 shows the decision boundaries when applied to the two sensor models obtained from real data, which are shown in figure 5.29. The vertical boundaries occur when the two sensor model *swap* position as can be seen in the zoom in figure 5.32.

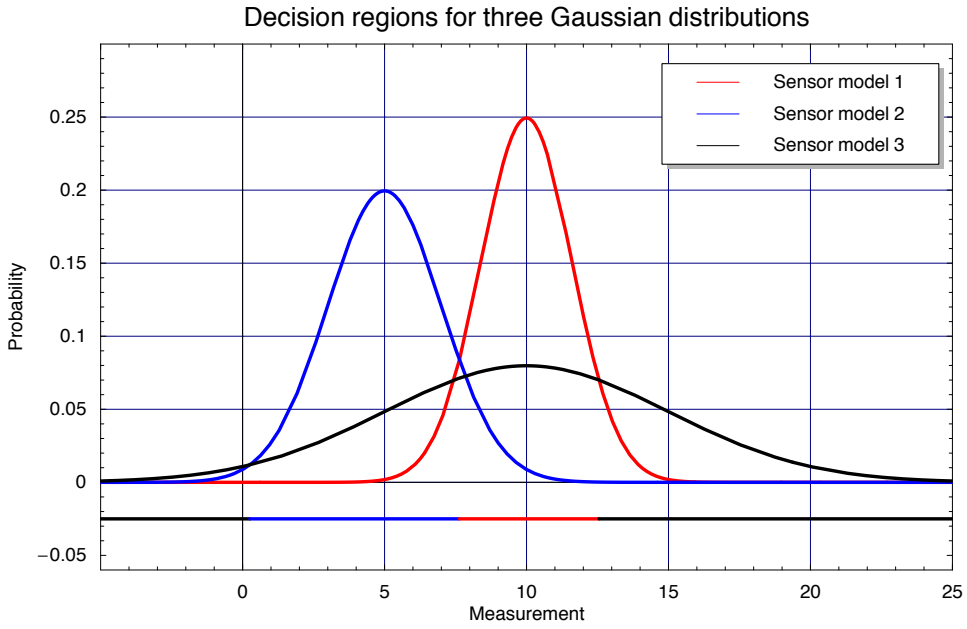

Figure 5.30: Decision regions and boundaries for three sensor model candidates — The three conditional distributions for an RF sensor depict the probability of making a certain measurement for a given distance d which has to be provided through external information. The decision regions are shown below the distributions in the corresponding colour. It is worth noting that although model 1 and model 3 have the same mean, model one will be preferred for measurements around the mean as this model is more informative, or higher peaked.

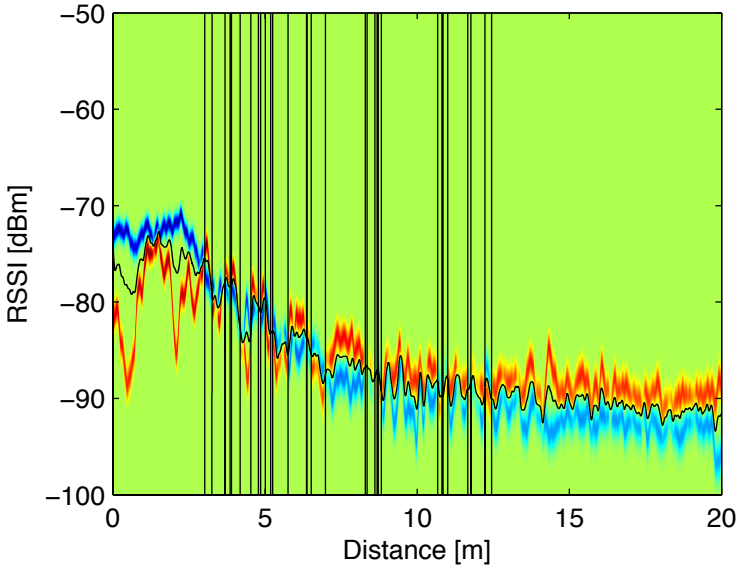

Figure 5.31: Decision boundaries for two sensor models obtained with real data — The decision boundaries corresponding to the two sensor likelihood-functions from figure 5.29 are shown in black. One model can be identified with the blue color tones, while the second one is coloured in the red tones.

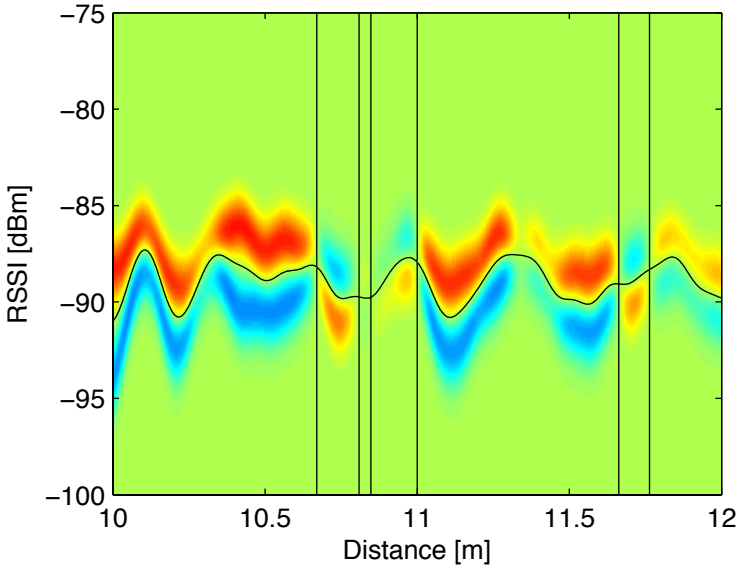

Figure 5.32: Decision boundaries for two sensor models obtained with real data (zoom) — This zoom shows how the vertical decision boundaries occur due to a change of position of the two likelihood-functions. Whereas one function was in the top and the other one was in the bottom of the picture, they change their position which in this case leads to a vertical decision boundary.

5.4.3 On-line sensor-model selection — examples

This subsection demonstrates on-line identification examples. The distance from a receiver to a sender is tracked using three different RF sensors that are assumed to be mounted at the same position, each sensor with an individual sensor model. An equivalent situation can be found where a single sensor tracks the distance, but three different sensor models are available for this sensor, and the task is to identify which one is the correct model. For the filtering process the sensor models are used as shown in figure 5.33. The receiver moves at a constant speed away from the transmitter and in this example geo-referenced measurements are received at regular intervals. Referenced, noisy RF measurements come with information about the true distance between transmitter and receiver. The distance covered is assumed to be from 2 m to 19 m and referenced measurements are obtained at 0.25 m intervals. To illustrate the task of identifying the model, in this simulation setup measurements are generated between 2 m and 5 m using sensor model 1, between 5 m and 10 m using sensor model 2 and between 10 m and 19 m using sensor model 3. This means that the identification should detect the change in the model to be used for each of the three distance intervals. Monte Carlo Markov Chain (MCMC) simulations with 1000 runs were conducted to investigate the behaviour of the identification algorithms.

Example 1 — uninformative prior

An uninformative prior is chosen at the reception of each referenced measurement, therefore basing the choice of the model at each interval entirely on the actual measurement. This corresponds to a maximum likelihood decision discarding any prior knowledge about the model. The result for 1000 runs can be seen in figure 5.34. Frequent switching of the model occurs even when the sensor model has not changed. It is apparent that the effect of each decision about which model is correct, is based solely on the current measurement. Even though each measurement is referenced, it is noisy, which often leads to an incorrect choice.

It can also be seen that certain models are similar for given distance intervals. For example, in the interval between 5 m and 7 m distance, model 2 and model 3 have similarities. This can be seen by the fact that although model 2 is the correct model, frequent switching to model 3 occurs.

Example 2 — maintaining posterior estimates

In this example an uninformative prior is chosen initially at the start of the run, but then the posterior estimates are maintained and used as priors for the next iteration. This has a stabilising effect in comparison to the previous example, as can be seen in figure 5.35.

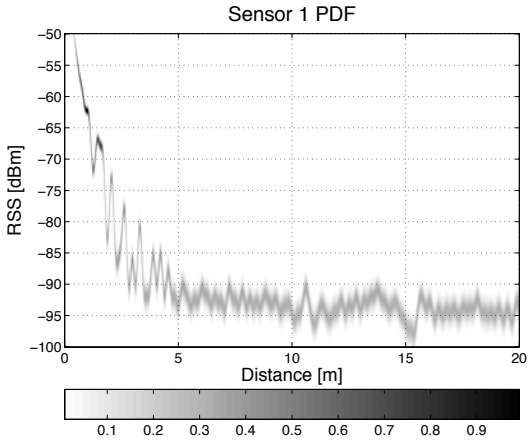

(a)

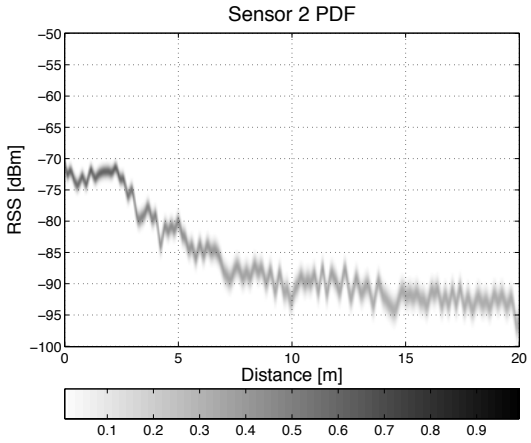

(b)

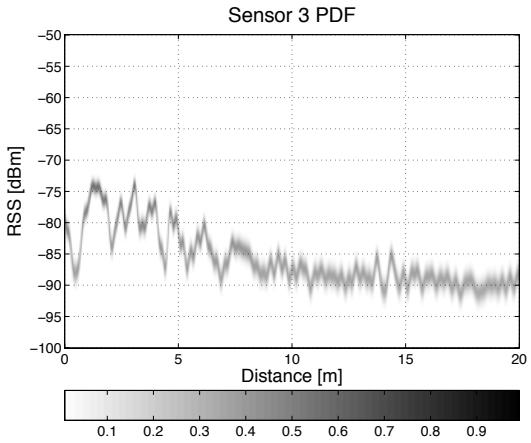

(c)

Figure 5.33: Sensor models for model identification — Three different sensor models are available for one sensor. It can be assumed, that depending on the environment, a switch to a more suitable sensor model may improve the localisation process.

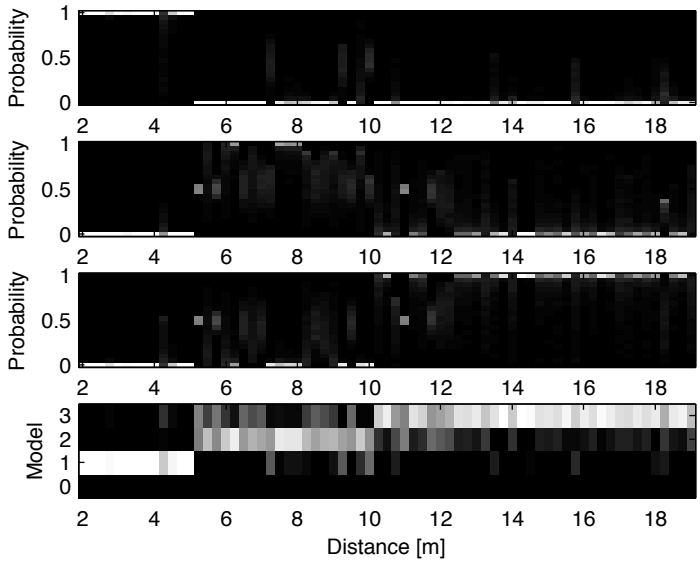

Figure 5.34: Choice of models (uninformative prior) — The top three plots show for each sensor model (1 is at the top, 3 is at the bottom) the probability as calculated at the reception of a referenced measurement based only on the measurement, i.e. using an uninformative prior. The bottom plot shows which model has been selected on the basis of the calculated probabilities. As can be seen, switching based on the actual measurement can lead to frequent switching even when the model has not changed. All plots are histograms of the simulation runs.

Nevertheless, a new failure mode can occur when the algorithm is implemented as described. It can occur that the actual switching transition happens too late. This means that the algorithm detects the model change some time after the change has occurred. If the model needs to be switched very often in the real application, then this can lead to the situation that some transitions are completely missed, as shown in this example.

The reason for this behaviour can be explained by looking at the fact that the probabilities used in the decision rule contain the complete history of the decisions taken. In this example as with the previous example, in the first interval, model 1 is often selected. Since the calculated posterior serves as the new prior, it is reinforced at each step in this interval that model 1 is the correct one, while number 2 and 3 should not be selected. This leads to numerically small values for model 2 and 3, while the value for model 1 will be very close to 1. If now a model transition occurs, it will take time (or a significant number of referenced observations) until this situation is reversed and a transition to the new correct model can occur. In the case presented here model 2 is correct for the second interval, but there have not been enough observations contributing to the belief that model 2 is correct to ensure the switch actually occurs. Switching is also hindered by the fact that model 2 is not the *clear* candidate for the second interval, which again can be seen from the previous example. Instead of switching, the old model is kept. The transition occurs only in the third interval, where the correct model is chosen, albeit with a long delay. It can be noted that the real transition occurs at 10 m, while for the MCMC simulations the first transitions occur no earlier than approximately 12 m. Furthermore, it can be seen how smooth the behaviour of the transition is for the 1000 runs.

5.4.4 Decision theory with variable switching threshold and minimum posterior threshold

When choosing the decision rule as given by equation (5.16) or equation (5.17), and recursively updating the posteriors two situations might occur:

- The model switches frequently between similar models, as shown in the first example and figure 5.34.
- The switching transition occurs with a delay as shown in the second example and figure 5.35.

If two models are similar it might happen that the decision boundary is crossed frequently. To avoid this behaviour, if it is undesired, a modified decision rule based on a threshold

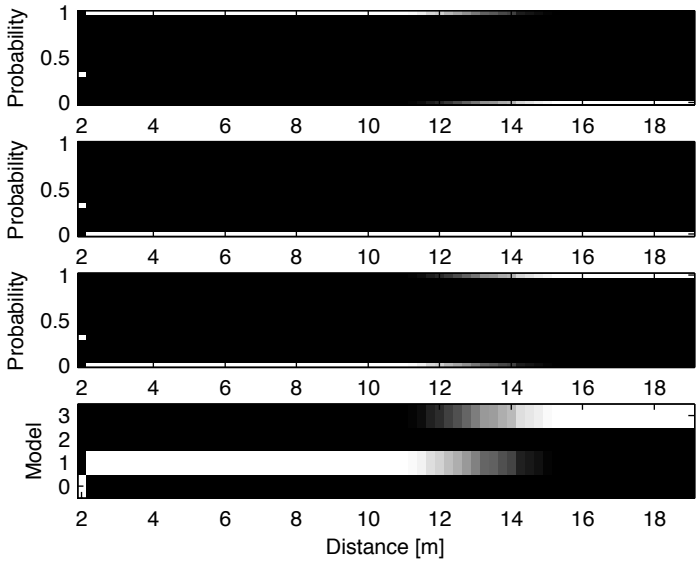

Figure 5.35: Choice of models (keeping posterior estimates) — The top three plots show for each sensor model (1 is at the top, 3 is at the bottom) the probability as calculated at the reception of a referenced measurement based on the measurement and the prior probability that this particular model is the correct model. The bottom plot shows again which model has been selected on the basis of the calculated probabilities. The switching occurs with a delay, and in the worst case a transition to a model (model 2) is skipped.

ε_{switch} can be used as follows, if the posteriors are kept and recursively used as the new priors, i.e. when Bayes Rule is applied.

Choose model M_i if

$$\mathcal{P}(M_i|z) > \varepsilon_{switch} \qquad \text{where } \varepsilon_{switch} > 0.5 \tag{5.19}$$

The choice of a threshold based on equation (5.19) ensures that, in the case of more than two models, no two models can have a posterior probability greater than the threshold.

The second situation is based on the fact that the belief that a certain model is correct can become over-confident. This happens because the posterior for a certain model keeps increasing, and in the limit approaches 1. Such a situation occurs when a model has been used correctly for a longer period of time, causing the belief that this particular model is correct to increase. If the model is switched now, it will take a certain amount of time for belief in the new correct model to be built up until it reaches the switching point. This delay means that the switching is not responsive. To counteract such a behaviour, the decision making process should always keep some *skepticism* about which model is the correct model. This can happen by introducing a lower limit ε_{min}. Individual beliefs can not become smaller than this limit, and if they do become smaller, they are set to this limit value. The justification for the above mentioned skepticism lies in the fact that prior knowledge about the system is available that assumes the model will switch at some point in time. Therefore, it is reasonable to keep some minor, but not too small, degree of belief in all the other models.

The following pseudo-algorithm incorporates the two presented modifications. It takes as its inputs the set of models (M), the indicator which is the presently used model (m_{old}), the set of priors (*prior*), which corresponds to the models, as well as the thresholds ε_{switch} and ε_{min}. It receives external referenced measurements and iterates as long as no model switching is detected. At each iteration, upon reception of a referenced measurement, it calculates a posterior for all models. Subsequently it checks whether any of the posteriors is smaller than the threshold ε_{min} and sets them to the threshold value if this is the case. The final step is based on the threshold ε_{switch}, where it is evaluated whether the model has changed. If this is the case the function is exited, and the new model along with the posteriors is returned. Otherwise the next iteration occurs.

Both thresholds are tuning parameters, which must be adapted to the models used and the frequency with which the referenced measurements are available.

Algorithm 5.4.1: ModelIdentify$(M, m_{old}, prior, \varepsilon_{switch}, \varepsilon_{min})$

$n \leftarrow$ NumberOf(M)
$m_{new} \leftarrow m_{old}$
$posterior \leftarrow prior$
while $(m_{new} = m_{old})$
do
- $z \leftarrow$ ReferencedMeasurement()
- **for** $i \leftarrow 1$ **to** n
  - **do** $posterior_i \leftarrow$ CalcPosterior$(M_i, posterior_i, z)$
- **for** $i \leftarrow 1$ **to** n
  - **do** **if** $posterior_i < \varepsilon_{min}$ **then** $posterior_i \leftarrow \varepsilon_{min}$
- **for** $i \leftarrow 1$ **to** n
  - **do** $posterior_i \leftarrow$ Normalize$(posterior)$
- **for** $i \leftarrow 1$ **to** n
  - **do** **if** $posterior_i > \varepsilon_{switch}$ **then** $m_{new} \leftarrow i$

return $(m_{new}, posterior)$

Example 3 — minimum posterior threshold

The introduction of a threshold for the minimum posterior makes the model identification process more responsive. Figure 5.36 shows that the model transition is detected shortly after the model has changed. Also, the erratic switching behaviour between the models as in example 1 is no longer present.

Example 4 — modified switching threshold

Introducing a switching threshold greater than 0.5 delays the switching transition marginally, but makes the process more stable. In figure 5.37, the transitions are detected correctly as in the previous example.

The differences between example 3 and example 4 are subtle and can be best seen when comparing the transition from model 1 to model 2. While, in example 3, there is a slight probability that the transition incorrectly occurs from model 1 to model 3 and then to

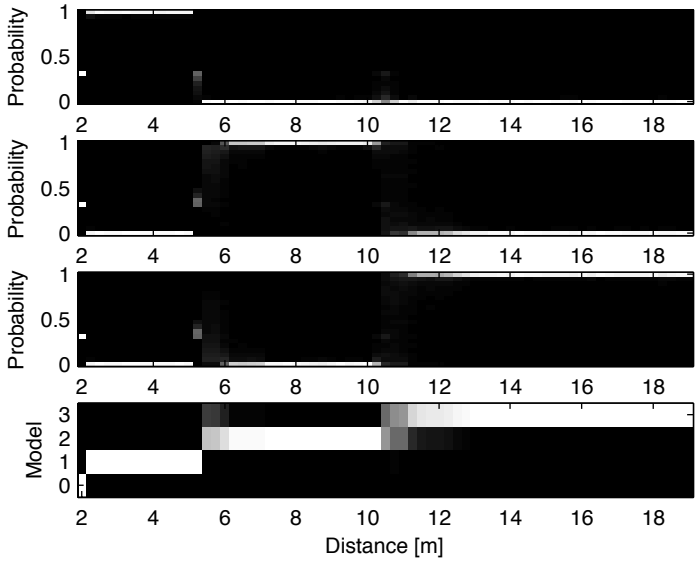

Figure 5.36: Choice of models (minimum posterior threshold) — The top three plots show for each sensor model (1 is at the top, 3 is at the bottom) the probability as calculated at the reception of a referenced measurement based on the measurement and the prior probability, that this particular model is the correct model, using a minimum posterior threshold. An overall responsive and stable behaviour can be seen.

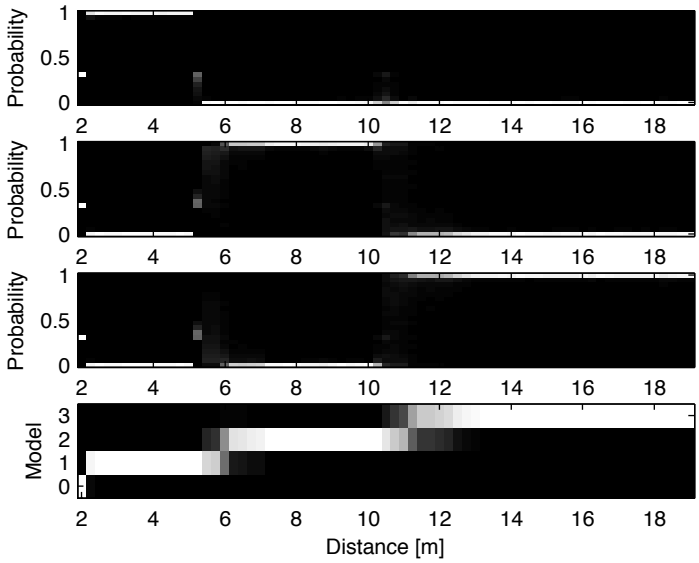

Figure 5.37: Choice of models (modified switching threshold) — The top three plots show for each sensor model (1 is at the top, 3 is at the bottom) the probability as calculated at the reception of a referenced measurement based on the measurement and the prior probability, that this particular model is the correct model, using a minimum posterior threshold and an additional variable switching threshold. Again an overall responsive and stable behaviour can be seen.

model 2, this case is not apparent in example 4. This however comes at the cost that the transition in the example occurs with a slight delay.

This section showed the possibility of identifying on-line the correct sensor model to be used in a tracking filter. Using the two thresholds ε_{switch} and ε_{min}, the behaviour of the on-line detection algorithm can be tuned to suit the application. The beauty of the solution presented lies in its simplicity. Bayesian Decision Theory, with minor adjustments for practical applications, provides the basis for the algorithm shown in this subsection.

5.5 Summary

In this chapter an alternative sensor representation which is based on geo-referenced measurements, and which can approximate multi-modal distributions as occurring with RF sensors, was described. This sensor model describes received signal intensity as a function of distance and azimuth angle. It also incorporates uncertainty and can be used in Particle Filter (PF) or Histogram Filter (HF) tracking architectures. The most important property of this model is that it reflects the signal propagation correctly based on training data, as it is actually an interpolant rather than a regression curve. The comparison of the new sensor description with the commonly used nth power model highlighted the reason why high-accuracy is not achievable in localisation applications with the latter model. The fact that the signal mean function is not correctly represented by the nth power model leads to low accuracy in the localisation filter.

Multi-sensor localisation to obtain the position of an agent using range-only sensors brought an insight into how this can be achieved and what effects can occur when the range obtained from the sensors is not correct. With these types of sensors it is absolutely necessary to fuse the information from multiple spatially separated sensors to resolve the bearing uncertainty. As an extension to a centralised approach, it was shown how measurements from external observers can be incorporated as virtual measurements to improve an agent's position estimate. This approach strongly depends on the quality of the position information available about the external observer.

Bayesian Decision Theory was used as a tool to identify the correct sensor model to be used in given environmental circumstances. This is of practical relevance where it is necessary to select the correct candidate from multiple available sensor models.

The following chapter will illustrate the main theoretical considerations and contributions from this chapter using simulations and experimental results.