

Using the Electromagnetic Spectrum

M. ACCETTULLI-HUBER¹, D. ALLWRIGHT², D. S. BERMAN¹, E.
CAMPILLO-FUNOLLET^{3†}, P. DATE⁴, and R. GONZO⁵,

¹ *Queen Mary University of London, London, UK*

² *University of Oxford, Oxford, UK*

³ *University of Sussex, Brighton, UK*

⁴ *Brunel University, London, UK*

⁵ *Trinity College Dublin, Dublin, Ireland*

(Communicated to MIIR on 28 June 2021)

Study Group: ESGI 165

Communicated by: Bernard Piette

Industrial Partner: DSTL.

Presenter: Emma Bowley, Chris Williams

Team Members: Manuel Accettulli-Huber, David Allwright, Stefano de Angelis, David Berman, Eduard Campillo-Funollet, Paresh Date, Riccardo Gonzo, David Torkington

Industrial Sector: Defence.

Tools: Fisher information matrix, Markov chain Monte Carlo, particle filter

Key Words: source detection, electromagnetic spectrum

MSC2020 Codes: 60J20, 78A46.

Summary

We consider the problem of estimating the power and position of an unknown transmitter using only observations on the received power for several sensors, without any directional information. We describe several existing propagation models, and we compute the Fisher information for the free-space loss model. For more complex model, we present some results using a Markov Chain Monte Carlo method, and we describe the application of a particle filter to estimate posterior distributions for the transmitter and channel characteristics.

† Corresponding author: e.campillo-funollet@sussex.ac.uk

1 Introduction

1.1 Problem description

Consider an area with an unknown number of transmitters. DSTL is interested in estimating the use of the electromagnetic spectrum in the area. The available information is limited to observations of the received power—but not directionality—as a function of the frequency for a given number of sensors, scattered on the area of interest, but not uniformly distributed. For example, the sensors can be concentrated to one side of the area of interest, with the goal to estimate information about the electromagnetic spectrum before moving closer to the transmitters, see Figure 1. For this problem, we are not interested in developing new propagation models for electromagnetic waves; we will use existing models such as the models in the WINNER II report [2].

DSTL summarised the challenge into five goals:

- (1) **Modelling Uncertainty:** The propagation models are very useful, but we need to understand the uncertainty when we do not know the true values of the parameters to put into them. For example, we may know the location of a transmitter, but have a probability distribution for its height. We may have a situation where several of the variables in the model are in terms of distributions and we need to understand how to combine them.
- (2) **Finding Transmitters:** If we have a set of measurements from our sensors, can we produce estimates of where the transmitters might be? The sensors may only record the power that they have received and not directionality making the challenge harder.
- (3) **Placing Sensors:** Can we work out the density of sensors needed to determine the spectrum usage within acceptable bounds? How will our uncertainty about the environment vary as we change the sensor density? If we know something about the topology, can we work out the best locations for the sensors?
- (4) **Find Spectrum Occupancy:** Given a sparse set of sensor readings, can we work out what parts of the spectrum are being used and the power density across an area?
- (5) **Safe to Transmit:** Building on the questions above, can we take the set of sensor readings and deduce how safe it is to transmit if we know the susceptibilities of the equipment? For example, if a piece of equipment is operating at a particular frequency we may not be able to transmit at frequencies close to that or risk causing interference. What is the best way to present this to people (or machines!) tasked with making decisions?

The work during the European Study Group with Industry (ESGI) 165 focused on modelling uncertainty and finding the position of the transmitters. We note that once the position of the transmitters is estimated, we could apply a propagation model to estimate the electromagnetic spectrum at a new location, and we could push forward the uncertainty using the same model.

1.2 Model set up

Suppose in some channel (i.e. a combination of centre frequency f_c and bandwidth) we have power measurements P_{r_i} at N receiver points \mathbf{x}_i for $i = 1, 2, \dots, N$. In order to understand the electromagnetic spectrum in the area, our goal is to estimate what would be the received power at some other point or points. Note that these points are generally outside the region covered by the \mathbf{x}_i , e.g. outside their convex hull and possibly at distances that may be larger than the distances between the \mathbf{x}_i . Among the possible unknowns are:

- (1) the number of transmitters operating in that channel;
- (2) their positions (including their height above the terrain);
- (3) their powers;
- (4) their directionality of their transmission;
- (5) whether they have direct line-of-sight (LOS) to our receiver positions (which affects what path loss model is appropriate);
- (6) the variability in time of any of these, e.g. whether they are moving transmitters, whether their transmit power varies according to what use they are making of the channel, how the path loss varies (fading) etc.

Because of the number and nature of the unknowns and the noise levels (e.g. 8 dB) in the path loss models, it is plausible at the outset that any method of making useful progress with these questions will need to make use of any *prior* information that may be available. In Section 4 we investigate these difficulties using a free-space path loss model.

The prior information could take many forms but could include for instance:

- (1) Geographical information to know which parts of the area are undulating or flat; open country, forest, suburban, urban etc. This enables us to know what path loss models are appropriate.
- (2) Transmitter location information, e.g. that certain areas are more or less likely to contain transmitters.
- (3) Transmitter power information, for instance about the type of transmitter, its centre frequency and the likelihood to transmit within a power range.

To simplify the problem, DSTL have advised that we should assume *omnidirectional* transmitters, and initially assume only one transmitter in a channel, and take path-loss models like those presented in the WINNER II Channel models report [2]. We should also assume that the receiver positions are known, and that transmit power does not vary too rapidly—in particular that there is not fast power control. The measurements at the receiver positions may not be synchronized; we need a model that can cope with the situation where at each time step we only get signals from *some* of the sensors.

The path loss models presented in [2] take a standard form (measured in decibels)

$$\text{PL} = A \log_{10}(d) + B + C \log_{10} \left(\frac{f_c}{5.0} \right) + X, \quad (1.1)$$

where d is the horizontal range in metres, A includes the path loss exponent, f_c is the

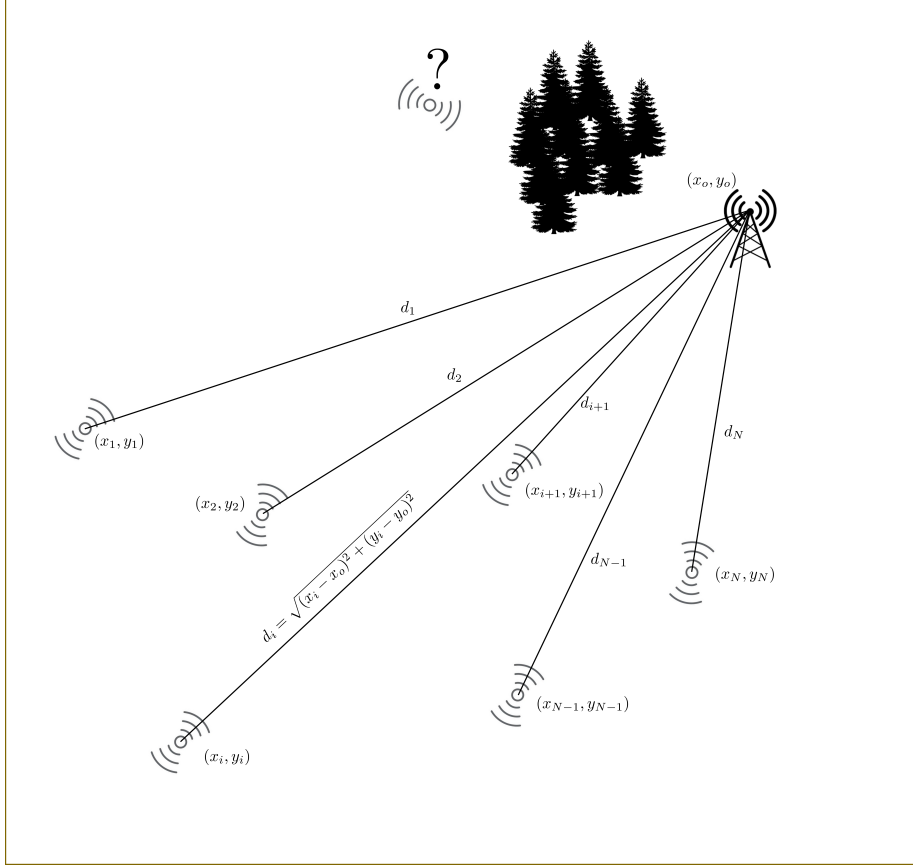


Figure 1. A number of sensors in known locations, labeled with (x_i, y_i) , receive the signal from a transmitter located at (x_o, y_o) . The goal of this challenge is to estimate the signal at an unreachable location, and to provide uncertainty quantification on the estimate accounting for unknown characteristics of the area, e.g. the presence of a forest.

centre frequency in GHz, and X is an environment-specific term. Each such formula comes with a certain frequency band and range that it is recommended for. DSTL is also interested in accounting for uncertainty on A, B, C and X .

For particular kinds of terrain and transmitter, the WINNER II report [2] provides values for the parameters in (1.1). For example,

$$PL = 25.1 \log_{10}(d) + 55.4 - 0.13(h_{bs} - 25) \log_{10}(d/100) - 0.9(h_{ms} - 1.5) + 21.3 \log_{10}(f_c/5.0). \quad (1.2)$$

is the WINNER II model for scenario D1 NLOS, where D1 means a “Rural Macro-cell”, defined for ranges up to 10 km with low building density, and NLOS means Non Line-of-Sight. This model is intended for use over ranges d between 50 m and 5 km, over the frequency range 2–6 GHz, and with typical transmitter (base station) height h_{bs} of 32 m and typical receiver height h_{ms} of 1.5 m. The WINNER II document contains about 20 such formulae for different environments, all for the 2–6 GHz range. The ITU (International Telecommunication Union) produces similar recommended propagation

models for all relevant parts of the spectrum. For a given path loss PL the received power is modelled (again in dB) as

$$P_r = P_t - \text{PL} - \text{SF} \quad (1.3)$$

where SF is a random shadow fading term. In the WINNER II models SF is taken to be Gaussian with mean 0 and a specified standard deviation in each scenario, and in the D1 NLOS scenario this is 8 dB. It represents a combination of many factors, including for instance multipath propagation. It is because multipath propagation is effectively *multiplicative* noise that SF appears additively in the received power in dB. The WINNER II report does not seem to include recommendations as to what should be assumed about how SF is correlated in time, or in space. In practice the correlation in time and space of the different contributions to shadow fading will depend on the position and motion of the scatterers concerned. Multipath involving a scatterer near to the receiver is likely to vary more significantly with receiver position than that from a scatterer farther away. The contributions to multipath scattering from fixed objects (buildings, power lines etc.) will clearly be steady, but the contributions from moving objects (traffic, wind turbines etc.) will vary on the timescale of their motion. So in the context of DSTL's problem where we are trying to sense an *unknown* EM environment, it is clear that correlations of the shading in space and time are going to be difficult to model. In what follows we shall treat the shading as uncorrelated, but it should be borne in mind that this could be a considerable oversimplification. There are also elements of the system that are not represented in the WINNER II models at all, but that will still affect the fall-off in power with distance away from the transmitter. For instance, some of the models include the height of the transmitter above the terrain but none of them include the vertical aperture of the transmitter. However, that vertical aperture will affect the ability of the transmitter to direct energy into the horizontal plane, and so will affect the range-power curve. These “model error” terms also have to be represented by the shadow fading noise term SF; ; see [3] for models including other factors.

2 Bayesian Update Model

In this section, we will describe the application of a Bayesian update model (see for instance [4]) to obtain estimates for the transmitter position—and if necessary, for the channel parameters—and update them when new information is recorded. Note that here we are using the word *model* with two different meanings: the propagation model, describing how the signal of the transmitter propagates across the area, and the Bayesian update model, describing how to assimilate new information to update our estimates.

Assume that we have one transmitter and N sensors. $P_{r_i}(k)$ is the measurement at sensor i at time k , $P_t(k)$, $d_i(k)$, $f_c(k)$, $h_{bs}(k)$ are defined similarly. At each time k , the following equation describes the measurements:

$$P_{r_i}(k) = P_t(k) - \left(25.1 \log_{10} d_i(k) + 55.4 - 0.13(h_{bs}(k) - 25) \log_{10} \frac{d_i(k)}{100} - 0.9(h_{ms} - 1.5) + 21.3 \log_{10} \frac{f_c(k)}{5} + W_i(k) \right), \quad i = 1, 2, \dots, N. \quad (2.1)$$

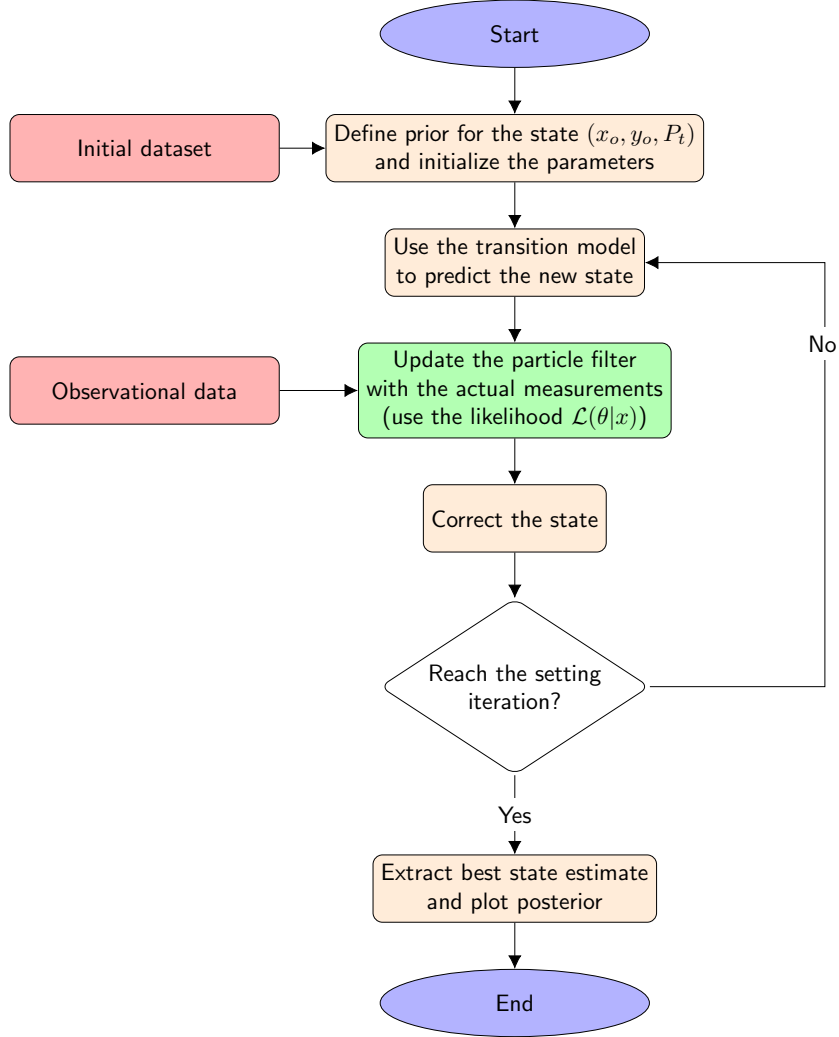


Figure 2. Workflow for the recursive Bayesian particle filter

$$d_i(k) = \sqrt{(x_i(k) - x_o(k))^2 + (y_i(k) - y_o(k))^2}.$$

$W_i(k)$ is a zero mean Gaussian noise process, with variance σ_i^2 . This represents inaccuracies in our model as well as observation noise. For the D1 NLOS path loss model, σ is 8 dB. If necessary, we can assume W_i, W_j to be correlated and build that into the Bayesian update model.

Given the estimates at time k , $(P_t(k), x_o(k), y_o(k))$, the transition equation,

$$\begin{bmatrix} P_t(k+1) \\ x_o(k+1) \\ y_o(k+1) \end{bmatrix} = \begin{bmatrix} P_t(k) \\ x_o(k) \\ y_o(k) \end{bmatrix} + V(k), \quad (2.2)$$

describe the parameter update to time $k+1$. Here, $V(k)$ is a zero mean noise process with

compatible dimensions, representing modelling inaccuracies. Covariance matrix of $V(k)$ can be ‘learnt’ through Bayesian update or set initially through method of moments; the former method requires more computational power, whilst the latter is easier to compute but incorporates less information from the observations.

Apart from the measurement and transition equations above, we set $P_t(0), x_o(0), y_o(0)$, the initial—*a priori*—estimates. These, along with (constant) covariance matrices of $W(k), V(k)$ can be found for example using method of moments or quasi-maximum likelihood.

The Bayesian update algorithm can be implemented as a particle filter (see for instance [5]). Figure 2 describes the workflow of a generic particle filter. At the start, we use an initial dataset, based on *a priori* knowledge, to define the prior state of (x_o, y_o, P_t) , describing the position and power of the transmitter. The transition model (2.2) provides a new candidate state, that we refine using the likelihood of the new observations. Once we stop receiving information, we extract the estimates as well as information about the uncertainty from the posterior distribution. Note that within the particle filter the states are described in terms of probability distributions, and in consequence they contain at the same time information about the best estimates and the uncertainty.

Once we have updated distributions of all the l.h.s. parameters of (2.2), we can find distribution of power at a coordinate (x^*, y^*) from the transmitter by applying the propagation model,

$$P^*(k) = P_t(k) - \left(25.1 \log_{10} d^*(k) + 55.4 - 0.13(h_{bs}(k) - 25) \log_{10} \frac{d^*}{100} - 0.9(h_{ms} - 1.5) + 21.3 \log_{10} \frac{f_c}{5} \right), \quad (2.3)$$

where $d^* = \sqrt{(x^* - x_o(k))^2 + (y^* - y_o(k))^2}$.

3 Free-space path loss

In order to gain insight into the problem, we study the free-space path loss model. Following the notation introduced in the previous section, we have

$$P_{r_i} = P_t - 10n \log_{10} d_i. \quad (3.1)$$

In this we have incorporated the B, C and X terms from (1.1) into our P_t , whilst n is the propagation exponent and depends on the environment, typically $n \in [2, 4]$.

3.1 A simple model with Gaussian errors

Given the positions and measurements of N sensors, and assuming Gaussian distributed errors, the (negative) loglikelihood for the problem is

$$L(P_t, n, (x_o, y_o); P_{r_1}, \dots, P_{r_N}, (x_1, y_1), \dots, (x_N, y_N)) \propto \sum_{i=1}^N (P_{r_i} - P_t + 10n \log_{10} d_i)^2, \quad (3.2)$$




Figure 3. Posterior with Friis model, fixed n and power range, and one sensor. Green dot is the true position of the emitter, blue dot is the sensor.

where for simplicity we absorbed the proportionality constant in (3.1) into P_t and n . Assuming a uniform prior, (3.2) can be interpreted as the non-normalised posterior distribution for the parameters $P_t, n, (x_o, y_o)$ given the observations P_{r_i} .

To illustrate the results using (3.2) we first fix n and set a range of possible emitting powers $P_t \in [-10, 10]$. Figure 3 depicts the log-likelihood under this conditions. We can see that it is not possible to infer any information about the location of the emitter, except for a cut-off near the sensor. The cut-off here is due to the limited range of transmitting powers that we are considering.

Similarly, we can add more sensors to see how the posterior distribution for the location of the emitter changes. Figure 4 shows the posteriors for three and six sensors.

Using the posterior distributions we can apply the forward model (3.1) to compute the expected power at any given location, see Figure 5. Note that the expected power here is high due to applying the Friis model even if the distance to the emitter is small. Although the overall profile of the expected value would not change, the actual value will be smaller if a cut-off is applied. The free-space path loss model, as well as any of the models in the WINNER II report [2], cannot be applied below a given threshold for the distance to the transmitter.

3.2 Bayesian inference

If we interpret (3.2) as the Bayesian posterior with uniform priors, we can apply a Markov Chain Monte Carlo (MCMC) method to explore the posterior distribution for the position of the transmitter. In this context, (3.2) characterises our target probability distribution, and the MCMC method provides a numerical approximation.




Figure 4. Posteriors with Friis model, fixed n and power range. Green dot is the true position of the emitter, blue dots are the sensors.

We implemented (3.2) using the Python module *pyMC3* [6], suitable for MCMC methods with propagation models that are computationally cheap to evaluate.




Figure 5. Expected power using the Friis model, fixed n and power range. Green dot is the true position of the emitter, blue dots are the sensors.

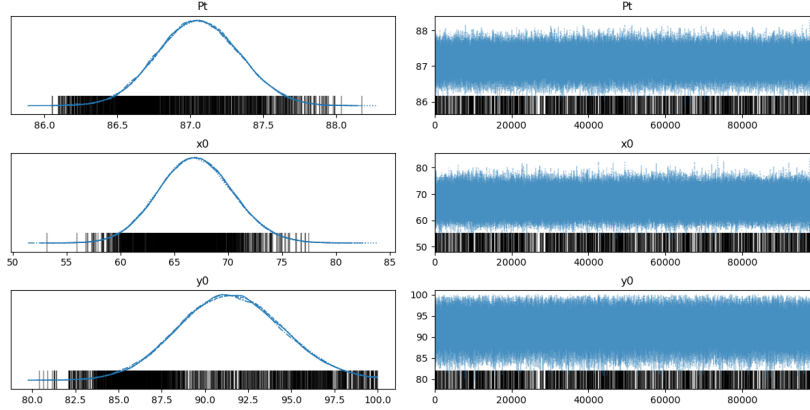


Figure 6. Posterior probability distribution for the position and power of the transmitter, approximated using a Markov Chain Monte Carlo algorithm implemented with *pyMC3*. The right-hand panel shows the traces of the Markov chain, ensuring a good mixing, i.e. a complete exploration of the probability space.

4 Fisher Information

One basic step in this problem is to compute the Fisher information matrix, expressing the information that the received powers give us about the locations and powers of transmitters. We begin with the case of just one transmitter, at position \mathbf{x}^* which has presumably been done often in the literature, but assuming that the shadow fading is independent at the N sensors the resulting Fisher information matrix with respect to

(\mathbf{x}^*, P_t) is

$$F = \frac{1}{\sigma^2} \begin{pmatrix} c^2 \sum_i \mathbf{v}_i \mathbf{v}_i^T & -c \sum_i \mathbf{v}_i \\ -c \sum_i \mathbf{v}_i^T & N \end{pmatrix}, \quad (4.1)$$

$$\mathbf{v}_i = \frac{\mathbf{x}^* - \mathbf{x}_i}{d_i^2}, \quad (4.2)$$

$$d_i = |\mathbf{x}^* - \mathbf{x}_i|. \quad (4.3)$$

In this $c = 10n/\ln(10) \approx 4.343n$ gives the dependence on the path-loss exponent. The sensors are at \mathbf{x}_i for $i = 1, 2, \dots, N$, and the distances $d_i = |\mathbf{x}^* - \mathbf{x}_i|$ are the ranges from transmitter to sensors. We are thinking of the \mathbf{x} as column vectors, the superscript T denotes transpose, and so the matrix F is 3×3 and is the information for the unknowns in the order \mathbf{x}^*, P_t . We have assumed that the shadow fading losses W_i are *independent*: this may not be a realistic assumption in practice, if, for instance, some of the fading is due to a terrain feature near the transmitter that affects *all* N transmission paths. If that were to apply then the information in the data would be less, and the resulting spread of the position distribution would be wider.

Since P_t is unknown, the Fisher information relating to \mathbf{x}^* is the Schur complement

$$F_{\mathbf{x}^*} = \frac{c^2}{\sigma^2} \sum_i (\mathbf{v}_i - \bar{\mathbf{v}})(\mathbf{v}_i^T - \bar{\mathbf{v}}^T), \quad (4.4)$$

$$\bar{\mathbf{v}} = \frac{1}{N} \sum_i \mathbf{v}_i. \quad (4.5)$$

This is proportional to the variance of a random vector that takes each value \mathbf{v}_i with probability $1/N$. In general terms therefore, this $F_{\mathbf{x}^*}$ is proportional to the square of the typical separation of the sensors, and inversely proportional to the fourth power of their range to the transmitter. So under these assumptions we see that the basic uncertainty about transmitter position is proportional to

$$\sigma \frac{1}{\sqrt{N}} \left(\frac{\ln(10)}{10n} \right) \frac{(\text{range})^2}{(\text{aperture})} (\text{geometric factors}), \quad (4.6)$$

and these dependencies are as to be expected. The inverse dependence on n is perhaps surprising: it results from the fact that if n is larger, the received power falls off more rapidly with distance, and therefore gives us more information about the range.

We can illustrate this with some examples. In these we have generated $N = 20$ sensors at random in the rectangle $|x| < 40$, $|y| < 10$, represented by the black asterisks, and placed a transmitter on the y -axis at $y = 30$, the black circle. Then we have generated 100 samples of power data using the values of n and σ stated in the plot title, and made 100 maximum likelihood estimates (MLE) of the transmitter position, which are the red circles. In Figure 7 we see the case with path loss exponent $n = 4$ and shadow fading $\sigma = 4$ dB. The green circles indicate the 1-standard-deviation and 2-standard-deviation ellipses around the true transmitter position (computed from F), and the MLE are reasonably well clustered around the true position, and with an error distribution roughly as predicted by the Fisher information matrix. If we instead make the path-loss exponent $n = 2$ rather than 4, we get results like those illustrated in Figure 8. Here we




Figure 7. MLE transmitter positions (red circles) with path loss exponent $n = 4$ and shadow fading $\sigma = 4$ dB.




Figure 8. MLE transmitter positions (red circles) with path loss exponent $n = 2$ and shadow fading $\sigma = 4$ dB.

see not only the larger uncertainty region but also that some of the MLE lie to the south of the sensor region: this is the kind of thing that might be eliminated by a suitable prior. If we go back to $n = 4$ but put in larger shadow fading with $\sigma = 8$ dB then we get results typically like that shown in Figure 9, again with a large uncertainty and some MLE points south of the sensors.

4.1 Multiple transmitters

If more than one transmitter is using the same part of the radio spectrum, then the problem escalates considerably. This could happen if two transmitters are using exactly the same centre frequency and bandwidth, or if they are close in centre frequency and there is spectral leakage outside their nominal bandwidth, or if they are close enough in frequency that they lie within the receiver bandwidth of our sensors. A general type of




Figure 9. MLE transmitter positions (red circles) with path loss exponent $n = 4$ and shadow fading $\sigma = 8$ dB.

model to represent this situation has to allow for the fact that the *number* of transmitters is itself one of the unknowns. One possible approach is to use the theory of Random Finite Sets, which has been analysed, including for radar and sonar applications, in the thesis of Trevor Wood [1]. In our context, we would only be applying the model if we had detected *some* activity in the channel, so the random finite set (the transmitters) would be modelled as having $k \geq 1$ elements, perhaps with $k - 1$ being modelled as Poisson with a small mean. Here though, we just discuss the problem in very basic terms. Recall that we are using powers P in dB throughout, so if there are two received signals P_1 , P_2 dB in the same channel then the combined received signal in dB is

$$P_r = 10 \log_{10}(10^{P_1/10} + 10^{P_2/10}), \quad (4.7)$$

since the transmitters can be assumed to be incoherent. Unless P_1 and P_2 are within about 6 dB of each other, this P_r is within 1 dB of being simply $\max(P_1, P_2)$: the stronger signal dominates. So suppose there are just two transmitters in the channel, at positions \mathbf{x}_1^* and \mathbf{x}_2^* , with transmit powers P_{t1} and P_{t2} and a path loss exponent of n . Then we expect to get useful information about the position of the first transmitter from sensor points with $P_{t1} - 10n \log_{10} |\mathbf{x}_1^* - \mathbf{x}_i| > P_{t2} - 10n \log_{10} |\mathbf{x}_2^* - \mathbf{x}_i|$, and vice versa for the second. If transmitter 1 is stronger, $P_{t1} > P_{t2}$, then the region where sensors give useful information about the position of the weaker transmitter, transmitter 2, is the region

$$10 \log_{10} \left(\frac{|\mathbf{x}_2^* - \mathbf{x}_i|}{|\mathbf{x}_1^* - \mathbf{x}_i|} \right) < \frac{P_{t2} - P_{t1}}{n}. \quad (4.8)$$

This region will be a circle containing \mathbf{x}_2^* (but not centred on it). It is likely therefore that if we have no sensors inside that circle we will fail to detect the second transmitter at all: our sensors in that channel of the spectrum will detect the dominant transmitter but the signals of the second will be masked and will appear to be accounted for as simply shadow fading or other noise. In terms of the question whether it is safe to transmit in some area, this therefore illustrates a wider class of difficulty: there may be transmitters

that are masked from our sensors, but which make it unsafe to transmit in some area of interest.

5 Conclusion and discussion

In this report, we presented several approaches to the problem of estimating the use of the electromagnetic spectrum from the received power for several sensors. We focused on the estimation of the location of unknown transmitters, since it is a critical step for the overall goal. Once the location of the transmitters is known, a propagation model would provide the electromagnetic spectrum usage for the area. Simple propagation models provide analytical insights, for instance the Fisher information described in Section 4; for more complex models, we proposed numerical inference methods such as particle filters (Section 2) or Markov Chain Monte Carlo methods (Section 3).

One aspect not explored in the report is how one combines information from different frequencies. One possibility that should be explored is how to infer the characteristics of a channel by incorporating information from different frequencies. Some channel characteristics will be frequency dependent and other characteristics will not. Thus by combining the information from different frequencies we can learn more about the channel and incorporate this into the Bayesian update described in Section 2.

The methods we have described provide a probability distribution of received power at a given point for a given frequency. One can then ask the question, what is the expectation value of the power at a given frequency for the given point? Then the obvious thing to do is to take the expectation value for each frequency using the calculated distribution. We will assume there are n -frequency bands thus providing us with n almost independent expectations. The independence may not be true for all points, since a shadow might be a shadow for all frequencies but the idea is that different transmitters using different frequencies would be independent. In any case independence is not needed for what follows. Given n estimators it is well known that we can improve on the naive maximum likelihood estimation if $n > 2$ by using what is known as shrinkage estimators as introduced by James and Stein. One minimises the **overall** risk of the estimators across all frequencies by taking the naive expectation value and then shrinking using the James-Stein formula [8, 9]. It is important to realise that although this will improve the estimation overall it will not do so for a particular fixed frequency. Thus the use of shrinkage depends very much on the question you wish to ask. If you want to ask what is the expected power at this given frequency then don't use shrinkage estimators. If you want to ask, out of all the possible frequencies what is the estimated power at the frequency with the largest or smallest value, then use shrinkage.

Although we focused on estimating the presence of one transmitter, the problem of estimating the total number of transmitters was also part of the challenge. Due to the logarithmic scale of the measurements, in the presence of two transmitters, the most powerful one dominates the signal in each sensor. We propose two approaches to tackle this problem: first, study the model without noise to understand how many solutions it has depending on the number of sensors, and second, use a model selection approach to select the number of sensors; see [7] for an introduction to model selection methods.

There are several optimisation problems connected with this challenge. In particular,

one can consider the optimal position of new sensors—possibly constrained to accessible areas—in order to maximise the information gain about the position of the transmitter, or the information gain about the channel characteristics.

A more general propagation model would consider a loss depending on the position, $l : \mathbf{R}^2 \rightarrow \mathbf{R}$, and compute the path loss as a line integral along the path from the transmitter to the receiver. Although it poses a challenging inference problem, the loss function could be described with a prior accounting for what it is known about the terrain, and then estimated using the Bayesian update approach.

6 Acknowledgements

The authors would like to thank the organisers of the European Study Group with Industry 165, and the DSTL problem proposers for their help during the Study Group.

References

- [1] Wood, T. 2011. Random finite sets for multitarget tracking with applications. Oxford University D. Phil. Thesis
- [2] Meinilä J, Kyösti P, Jämsä T and Hentilä L. 2009. WINNER II channel models. *Radio Technologies and Concepts for IMT-Advanced*, **2**, 39-92.
- [3] Munoz D, Lara FB, Vargas C and Enriquez-Caldera R. 2009. Position location techniques and applications. Academic Press.
- [4] Beck JL. 2011. Bayesian updating, model class selection and robust stochastic predictions of structural response. In *Proceedings of the 8th European Conference on Structural Dynamics*. K VIV Royal Flemish Society of Engineers.
- [5] Reich S and Cotter C. 2014. Probabilistic forecasting and Bayesian data assimilation. Cambridge University Press.
- [6] Salvatier J, Wiecki TV and Fonnesbeck C. 2016. Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*. **2**, e55.
- [7] Konishi S and Kitagawa G. 2008. Information criteria and statistical modeling. Springer Science & Business Media; 2008.
- [8] Stein, C. 1956. Inadmissibility of the usual estimator for the mean of a multivariate distribution, *Proc. Third Berkeley Symp. Math. Statist. Prob.*, 1, pp. 197–206, MR 0084922.
- [9] Wasserman L, Chapter 7.6, All of Nonparametric Statistics. Springer Texts in Statistic; 2006.