

Research Article

Application of the Evidence Procedure to the Estimation of Wireless Channels

Dmitriy Shutin,¹ Gernot Kubin,¹ and Bernard H. Fleury^{2,3}

¹Signal Processing and Speech Communication Laboratory, Graz University of Technology, 8010 Graz, Austria

²Institute of Electronic Systems, Aalborg University, Fredrik Bajers Vej 7A, 9220 Aalborg, Denmark

³Forschungszentrum Telekommunikation Wien (ftw.), Donau City Strasse 1, 1220 Wien, Austria

Received 5 November 2006; Accepted 8 March 2007

Recommended by Sven Nordholm

We address the application of the Bayesian evidence procedure to the estimation of wireless channels. The proposed scheme is based on relevance vector machines (RVM) originally proposed by M. Tipping. RVMs allow to estimate channel parameters as well as to assess the number of multipath components constituting the channel within the Bayesian framework by locally maximizing the evidence integral. We show that, in the case of channel sounding using pulse-compression techniques, it is possible to cast the channel model as a general linear model, thus allowing RVM methods to be applied. We extend the original RVM algorithm to the multiple-observation/multiple-sensor scenario by proposing a new graphical model to represent multipath components. Through the analysis of the evidence procedure we develop a thresholding algorithm that is used in estimating the number of components. We also discuss the relationship of the evidence procedure to the standard minimum description length (MDL) criterion. We show that the maximum of the evidence corresponds to the minimum of the MDL criterion. The applicability of the proposed scheme is demonstrated with synthetic as well as real-world channel measurements, and a performance increase over the conventional MDL criterion applied to maximum-likelihood estimates of the channel parameters is observed.

Copyright © 2007 Dmitriy Shutin et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Deep understanding of wireless channels is an essential prerequisite to satisfy the ever-growing demand for fast information access over wireless systems. A wireless channel contains explicitly or implicitly all the information about the propagation environment. To ensure reliable communication, the transceiver should be constantly aware of the channel state. In order to make this task feasible, accurate channel models, which reproduce in a realistic manner the channel behavior, are required. However, efficient joint estimation of the channel parameters, for example, number of the multipath components (model order), their relative delays, Doppler frequencies, directions of the impinging wavefronts, and polarizations, is a particularly difficult task. It often leads to analytically intractable and computationally very expensive optimization procedures. The problem is often relaxed by assuming that the number of multipath components is fixed, which simplifies optimization in many cases [1, 2]. However, both underspecifying and overspecifying the model order leads to significant performance degradation: residual intersymbol interference impairs the performance

of the decoder in the former case, while additive noise is injected in the channel equalizer in the latter: the excessive components amount only to the random fluctuations of the background noise. To amend this situation, empirical methods like cross-validation can be employed (see, e.g., [3]). Cross-validation selects the optimal model by measuring its performance over a validation data set and selecting the one that performs the best. In case of practical multipath channels, such data sets are often unavailable due to the time-variability of the channel impulse responses. Alternatively, one can employ model selection schemes in the spirit of Ockham's razor principle: simple models (in terms of the number of parameters involved) are preferred over more complex ones. Examples are the Akaike information criterion (AIC) and minimum description length (MDL) [4, 5]. In this paper, we show how the Ockham principle can be effectively used to perform estimation of the channel parameters coupled with estimating the model order, that is, the number of wavefronts.

Consider a certain class of parametric models (hypotheses) \mathcal{H}_i defined as the collection of prior distributions $p(\mathbf{w}_i | \mathcal{H}_i)$ for the model parameters \mathbf{w}_i . Given the measurement

data \mathcal{Z} and a family of conditional distributions $p(\mathcal{Z} | \mathbf{w}_i, \mathcal{H}_i)$, our goal is to infer the hypothesis $\widehat{\mathcal{H}}$ and the corresponding parameters $\widehat{\mathbf{w}}$ that maximize the posterior

$$\{\widehat{\mathbf{w}}, \widehat{\mathcal{H}}\} = \arg \max_{\mathbf{w}_i, \mathcal{H}_i} \{p(\mathbf{w}_i, \mathcal{H}_i | \mathcal{Z})\}. \quad (1)$$

The key to solving (1) lies in inferring the corresponding parameters \mathbf{w}_i and \mathcal{H}_i from the data \mathcal{Z} , which is often a nontrivial task. As far as the Bayesian methodology is concerned, there are two ways this inference problem can be solved [6, Section 5]. In the *joint estimation method*, $p(\mathbf{w}_i, \mathcal{H}_i | \mathcal{Z})$ is maximized directly with respect to the quantities of interest \mathbf{w}_i and \mathcal{H}_i . This often leads to computationally-intractable optimization algorithms. Alternatively, one can rewrite the posterior $p(\mathbf{w}_i, \mathcal{H}_i | \mathcal{Z})$ as

$$p(\mathbf{w}_i, \mathcal{H}_i | \mathcal{Z}) = p(\mathbf{w}_i | \mathcal{Z}, \mathcal{H}_i) p(\mathcal{H}_i | \mathcal{Z}) \quad (2)$$

and maximize each term on the right-hand side sequentially from right to left. This approach is known as the *marginal estimation method*. Marginal estimation methods (MEM) are well exemplified by expectation-maximization (EM) algorithms and used in many different signal processing applications (see [2, 3, 7]). MEMs are usually easier to compute, however they are prone to land in a local rather than global optimum. We recognize the first factor on the right-hand side of (2) as a parameter posterior, while the other one is a posterior for different model hypotheses. It is the maximization of $p(\mathcal{H}_i | \mathcal{Z})$ that guides our model selection decision. Then, the data analysis consists of two steps [8, Chapter 28], [9]:

- (1) inferring the parameters under the hypothesis \mathcal{H}_i

$$\begin{aligned} p(\mathbf{w}_i | \mathcal{Z}, \mathcal{H}_i) &= \frac{p(\mathcal{Z} | \mathbf{w}_i, \mathcal{H}_i) p(\mathbf{w}_i | \mathcal{H}_i)}{p(\mathcal{Z} | \mathcal{H}_i)} \\ &\equiv \frac{\text{Likelihood} \times \text{Prior}}{\text{Evidence}}, \end{aligned} \quad (3)$$

- (2) comparing different model hypotheses using the model posterior

$$\begin{aligned} p(\mathcal{H}_i | \mathcal{Z}) &\propto p(\mathcal{Z} | \mathcal{H}_i) p(\mathcal{H}_i) \\ &\equiv \text{Evidence} \times \text{Hypothesis Prior}. \end{aligned} \quad (4)$$

In the second stage, $p(\mathcal{H}_i)$ measures our subjective prior over different hypotheses before the data is observed. In many cases it is reasonable to assign equal probabilities to different hypotheses, thus reducing the hypothesis selection to selecting the model with the highest evidence $p(\mathcal{Z} | \mathcal{H}_i)$.¹ The evidence can be expressed as the following integral:

$$p(\mathcal{Z} | \mathcal{H}_i) = \int p(\mathcal{Z} | \mathbf{w}_i, \mathcal{H}_i) p(\mathbf{w}_i | \mathcal{H}_i) d\mathbf{w}_i. \quad (5)$$

¹ In the Bayesian literature, the evidence is also known as the likelihood for the hypothesis \mathcal{H}_i .




FIGURE 1: An equivalent baseband model of the radio channel with receiver matched filter (MF) front-end.


The evidence integral (5) plays a crucial role in the development of Schwarz's approach to model order estimation [10] (Bayesian information criterion), as well as in a Bayesian interpretation of Rissanen's MDL principle and its variations [5, 11, 12]. Maximizing (5) with respect to the unknown model \mathcal{H}_i is known as the evidence maximization procedure, or evidence procedure (EP) [13, 14].

Equations (3), (4), and (5) form the theoretical framework for our joint model and parameter estimation. The estimation algorithm is based on relevance vector machines. Relevance vector machines (RVM), originally proposed by Tipping [15], are an example of the marginal estimation method that, for a set of hypotheses \mathcal{H}_i , iteratively approximates (1) by alternating between the model selection, that is, maximizing (5) with respect to \mathcal{H}_i , and inferring the corresponding model parameters from maximization of (3). RVMs have been initially proposed to find sparse solutions to general linear problems. However, they can be quite effectively adapted to the estimation of the impulse response of wireless channels, thus resulting in an effective channel parameter estimation and model selection scheme within the Bayesian framework.

The material presented in the paper is organized as follows: Section 2 introduces the signal model of the wireless channel and the used notation; Section 3 explains the framework of the EP in the context of wireless channels. In Section 4 we explain how model selection is implemented within the presented framework and discuss the relationship between the EP and the MDL criterion for model selection. Finally, Section 5 presents some application results illustrating the performance of the RVM-based estimator in synthetic as well as in actual wireless environments.

2. CHANNEL ESTIMATION USING PULSE-COMPRESSION TECHNIQUE

Channel estimation usually consists of two steps: (1) sending a specific sounding sequence $s(t)$ through the channel and observing the response $y(t)$ at the other end, and (2) estimating the channel parameters from the matched-filtered received signal $z(t)$ (Figure 1). It is common to represent the multipath channel response as the sum of delayed and weighted Dirac impulses, with each impulse representing one individual multipath component (see, e.g., [16, Section 5]). Such special structure of the channel impulse response implies that the filtered signal $z(t)$ should have a sparse structure. Unfortunately, this sparse structure is often obscured by additive noise and temporal dispersion due to the finite bandwidth of the transmitter and receiver hardware. This

FIGURE 2: Sounding sequence $s(t)$.

motivates the application of algorithms capable of recovering this sparse structure from the measurement data.

Let us consider an equivalent baseband channel sounding scheme shown in Figure 1. The sounding signal $s(t)$ (Figure 2) consists of periodically repeated burst waveforms $u(t)$, that is, $s(t) = \sum_{i=-\infty}^{\infty} u(t - iT_f)$, where $u(t)$ has duration $T_u \leq T_f$ and is formed as $u(t) = \sum_{m=0}^{M-1} b_m p(t - mT_p)$. The sequence $b_0 \cdots b_{M-1}$ is the known sounding sequence consisting of M chips, and $p(t)$ is the shaping pulse of duration T_p , $MT_p = T_u$. Furthermore, we assume that the receiver (Rx) is equipped with a planar antenna array consisting of P sensors located at positions $\mathbf{s}_1, \dots, \mathbf{s}_P \in \mathbb{R}^2$ with respect to an arbitrary reference point. Let us now assume that the maximum absolute Doppler frequency of the impinging waves is much smaller than the inverse of a single burst duration $1/T_u$. This low Doppler frequency assumption is equivalent to assuming that, within a single observation window equivalent to the period of the sounding sequence, we can safely neglect the influence of the Doppler shifts.

The received signal vector $\mathbf{y}(t) \in \mathbb{C}^{P \times 1}$ for a single burst waveform is given as [2]

$$\mathbf{y}(t) = \sum_{l=1}^L a_l \mathbf{c}(\phi_l) e^{j2\pi\nu_l t} u(t - \tau_l) + \boldsymbol{\eta}(t). \quad (6)$$

Here, a_l , τ_l , and ν_l are respectively the complex gain, the delay, and the Doppler shift of the l th multipath component. The P -dimensional complex vector $\mathbf{c}(\phi_l) = [c_1(\phi_l), \dots, c_P(\phi_l)]^T$ is the steering vector of the array. Provided the coupling between the elements can be neglected, its components are given as $c_p(\phi_l) = f_p(\phi_l) \exp(j2\pi\lambda^{-1}(\mathbf{e}(\phi_l), \mathbf{s}_p))$ with λ , $\mathbf{e}(\phi_l)$ and $f_p(\phi_l)$ denoting the wavelength, the unit vector in \mathbb{R}^2 pointing in the direction of the incoming wavefront determined by the azimuth ϕ_l , and the complex electric field pattern of the p th sensor, respectively. The additive term $\boldsymbol{\eta}(t) \in \mathbb{C}^{P \times 1}$ is a vector-valued complex white Gaussian noise process, that is, the components of $\boldsymbol{\eta}(t)$ are independent complex Gaussian processes with double-sided spectral density N_0 .

The receiver front-end consists of a matched filter (MF) matched to the transmitted sequence $u(t)$. Under the low Doppler frequency assumption the term $e^{j2\pi\nu_l t}$ stays time-invariant within a single burst duration, that is, equal to a complex constant that can be incorporated in the complex gain a_l . The signal $\mathbf{z}(t)$ at the output of the MF is then given

as

$$\mathbf{z}(t) = \sum_{l=1}^L a_l \mathbf{c}(\phi_l) R_{uu}(t - \tau_l) + \boldsymbol{\xi}(t), \quad (7)$$

where $R_{uu}(t) = \int u(t')u^*(t+t')dt'$ is the autocorrelation function of the burst waveform $u(t)$ and $\boldsymbol{\xi}(t) = \int \boldsymbol{\eta}(t')u^*(t+t')dt'$ is a spatially white P -dimensional vector with each element being a zero-mean wide-sense stationary (WSS) Gaussian noise with autocorrelation function

$$\begin{aligned} R_{\boldsymbol{\xi}\boldsymbol{\xi}}(t) &= E\{\boldsymbol{\xi}_p(t')\boldsymbol{\xi}_p^*(t+t')\} = N_0 R_{uu}(t), \\ E\{\boldsymbol{\xi}_p(t')\boldsymbol{\xi}_p(t+t')\} &= 0. \end{aligned} \quad (8)$$

Here $E\{\cdot\}$ denotes the expectation operator. Equation (7) states that the MF output is a linear combination of L scaled and delayed kernel functions $R_{uu}(t - \tau_l)$, weighted across sensors as given by the components of $\mathbf{c}(\phi_l)$ and observed in the presence of the colored noise $\boldsymbol{\xi}(t)$.

In practice, however, the output of the MF is sampled with the sampling period $T_s \leq T_p$, resulting in PN -tuples of the MF output, where N is the number of MF output samples. By collecting the output of each sensor into a vector, we can rewrite (7) in a vector form:

$$\mathbf{z}_p = \mathbf{K}\mathbf{w}_p + \boldsymbol{\xi}_p, \quad p = 1 \cdots P, \quad (9)$$

where we have defined

$$\begin{aligned} \mathbf{z}_p &= [z_p[0], z_p[1], \dots, z_p[N-1]]^T, \\ \mathbf{w}_p &= [a_1 c_p(\phi_1), \dots, a_L c_p(\phi_L)]^T, \\ \boldsymbol{\xi}_p &= [\xi_p[0], \xi_p[1], \dots, \xi_p[N-1]]^T. \end{aligned} \quad (10)$$

The additive noise vectors $\boldsymbol{\xi}_p$, $p = 1 \cdots P$, possess the following properties that will be exploited later:

$$\begin{aligned} E\{\boldsymbol{\xi}_p\} &= \mathbf{0}, \quad E\{\boldsymbol{\xi}_m \boldsymbol{\xi}_k^H\} = \mathbf{0}, \quad \text{for } m \neq k, \\ E\{\boldsymbol{\xi}_p \boldsymbol{\xi}_p^H\} &= \boldsymbol{\Sigma} = N_0 \boldsymbol{\Lambda}, \quad \text{where } \Lambda_{i,j} = R_{uu}((i-j)T_s). \end{aligned} \quad (11)$$

Note that (12) follows directly from (8). The matrix \mathbf{K} , also called the design matrix, accumulates the shifted and sampled versions of the kernel function $R_{uu}(t)$. It is constructed as $\mathbf{K} = [\mathbf{r}_1, \dots, \mathbf{r}_L]$, with $\mathbf{r}_l = [R_{uu}(-\tau_l), R_{uu}(T_s - \tau_l), \dots, R_{uu}((N-1)T_s - \tau_l)]^T$.

In general, the channel estimation problem is posed as follows: given the measured sampled signals \mathbf{z}_p , $p = 1 \cdots P$, determine the order L of the model and estimate optimally (with respect to some quality criterion) all multipath parameters a_l , τ_l , and ϕ_l , for $l = 1 \cdots L$. In this contribution, we restrict ourselves to the estimation of the model order L along with the vector \mathbf{w}_p , rather than of the constituting parameters τ_l , ϕ_l , and a_l . We will also quantize, although arbitrarily fine,² the search space for the multipath delays τ_l . Thus, we

² There is actually a limit beyond which it makes no sense to make the search grid finer, since it will not decrease the variance of the estimates, which is lower-bounded by the Crammer-Rao bound [2].

do not try to estimate the path delays with infinite resolution, but rather fix the delay values to be located on a grid with a given mesh determining the quantization error. The size of the delay search space L_0 and the resulting quantized delays $\mathcal{T} = \{T_1, \dots, T_{L_0}\}$ form the initial model hypothesis \mathcal{H}_0 , which would manifest itself in the L_0 columns of the design matrix \mathbf{K} . This allows to formulate the channel estimation problem as a standard linear problem to which the RVM algorithm can be applied.

As it can be seen, our idea lies in finding the closest approximation of the continuous-time model (7) with the discrete-time equivalent (9). By incorporating the model selection in the analysis, we also strive to find the most compact representation (in terms of the number of components), while preserving good approximation quality. Thus, our goal is to estimate the channel parameters \mathbf{w}_p as well as to determine how many multipath components $L \leq L_0$ are present in the measured impulse response. The application of the RVM framework to solve this problem follows in the next section.

3. EVIDENCE MAXIMIZATION, RELEVANCE VECTOR MACHINES, AND WIRELESS CHANNELS

We begin our analysis following the steps outlined in Section 1. In order to ease the algorithm description we first assume that $P = 1$, that is, only a single sensor is used. Extensions to the case $P > 1$ are carried out later in Section 3.2. To simplify the notations we also drop the subscript index p in our further notations.

From (9) it follows that the observation vector \mathbf{z} is a linear combination of the vectors from the column-space of \mathbf{K} , weighted according to the parameters \mathbf{w} and embedded in the correlated noise ξ . In order to correctly assess the order of the model, it is imperative to take the noise process into account. It follows from (12) that the covariance matrix of the noise is proportional to the unknown spectral height N_0 , which should therefore be estimated from the data. Thus, the model hypotheses \mathcal{H}_i should include the term N_0 . In the following analysis we assume that $\beta = N_0^{-1}$ is Gamma-distributed [15], with the corresponding probability density function (pdf) given as

$$p(\beta | \kappa, \nu) = \frac{\kappa^\nu}{\Gamma(\nu)} \beta^{\nu-1} \exp(-\kappa\beta), \quad (13)$$

with parameters κ and ν predefined so that (13) accurately reflects our a priori information about N_0 . In the absence of any a priori knowledge one can make use of a noninformative (i.e., flat in the logarithmic domain) prior by fixing the parameters to small values $\kappa = \nu = 10^{-4}$ [15]. Furthermore, to steer the model selection mechanism, we introduce an extra parameter (hyperparameter) α_l , $l = 1 \dots L_0$, for each column in \mathbf{K} . This parameter measures the contribution or relevance of the corresponding weight w_l in explaining the data \mathbf{z} from the likelihood $p(\mathbf{z} | \mathbf{w}_i, \mathcal{H}_i)$. This is achieved by specifying the prior $p(\mathbf{w} | \boldsymbol{\alpha})$ for the model weights:

$$p(\mathbf{w} | \boldsymbol{\alpha}) = \prod_{l=1}^{L_0} \frac{\alpha_l}{\pi} \exp\left(-|w_l|^2 \alpha_l\right). \quad (14)$$

High values of α_l will render the contribution of the corresponding column in the matrix \mathbf{K} “irrelevant,” since the weight w_l is likely to have a very small value (hence they are termed *relevance hyperparameters*). This will enable us to prune the model by setting the corresponding weight w_l to zero, thus effectively removing the corresponding column from the matrix and the corresponding delay T_l from the delay search space \mathcal{T} . We also see that α_l^{-1} is nothing else as the prior variance of the model weight w_l . Also note that the prior (14) implicitly assumes statistical independence of the multipath contributions.

To complete the Bayesian framework, we also specify the prior over the hyperparameters. Similarly to the noise contribution, we assume the hyperparameters α_l to be Gamma-distributed with the corresponding pdf

$$p(\boldsymbol{\alpha} | \zeta, \epsilon) = \prod_{l=1}^L \frac{\zeta^\epsilon}{\Gamma(\epsilon)} \alpha_l^{\epsilon-1} \exp(-\zeta \alpha_l), \quad (15)$$

where ζ and ϵ are fixed at some values that ensure an appropriate form of the prior. Again, we can make this prior noninformative by fixing ζ and ϵ to small values, for example, $\epsilon = \zeta = 10^{-4}$.

Now, let us define the hypothesis \mathcal{H}_i more formally. Let $\mathcal{P}(\mathcal{S})$ be a power set consisting of all possible subsets of basis vector indices $\mathcal{S} = \{1 \dots L_0\}$, and $i \mapsto \mathcal{P}(i)$ the indexing of $\mathcal{P}(\mathcal{S})$ such that $\mathcal{P}(0) = \mathcal{S}$. Then for each index value i the hypothesis \mathcal{H}_i is the set $\mathcal{H}_i = \{\beta; \alpha_j, j \in \mathcal{P}(i)\}$. Clearly, the initial hypothesis $\mathcal{H}_0 = \{\beta; \alpha_j, j \in \mathcal{S}\}$ includes all possible potential basis functions.

Now we are ready to outline the learning algorithm that estimates the model parameters \mathbf{w} , β , and hyperparameters $\boldsymbol{\alpha}$ from the measurement data \mathbf{z} .

3.1. Learning algorithm

Basically, learning consists of inferring the values of \mathbf{w}_i and the hypothesis \mathcal{H}_i that maximize the posterior (2): $p(\mathbf{w}_i, \mathcal{H}_i | \mathcal{Z}) \equiv p(\mathbf{w}_i, \boldsymbol{\alpha}_i, \beta | \mathbf{z})$. Here $\boldsymbol{\alpha}_i$ denotes the vector of all evidence hyperparameters associated with the i th hypothesis. The latter expression can also be rewritten as

$$p(\mathbf{w}, \boldsymbol{\alpha}, \beta | \mathbf{z}) = p(\mathbf{w} | \mathbf{z}, \boldsymbol{\alpha}, \beta) p(\boldsymbol{\alpha}, \beta | \mathbf{z}). \quad (16)$$

The explicit dependence on the hypothesis index i has been dropped to simplify the notation. We recognize that the first term $p(\mathbf{w} | \mathbf{z}, \boldsymbol{\alpha}, \beta)$ in (16) is the weight posterior and the other one $p(\boldsymbol{\alpha}, \beta | \mathbf{z})$ is the hypothesis posterior. From this point we can start with the Bayesian two-step analysis as has been indicated before.

Assuming the parameters $\boldsymbol{\alpha}$ and β are known, estimation of model parameters consists of finding values \mathbf{w} that maximize $p(\mathbf{w} | \mathbf{z}, \boldsymbol{\alpha}, \beta)$. Using Bayes' rule we can rewrite this posterior as

$$p(\mathbf{w} | \mathbf{z}, \boldsymbol{\alpha}, \beta) \propto p(\mathbf{z} | \mathbf{w}, \boldsymbol{\alpha}, \beta) p(\mathbf{w} | \boldsymbol{\alpha}, \beta). \quad (17)$$

Consider the Bayesian graphical model [17] in Figure 3. This graph captures the relationship between different variables involved in (16). It is a useful tool to represent the dependencies among the variables involved in the analysis in




FIGURE 3: Graph representing the discrete-time model of the wireless channel.

order to factor the joint density function into contributing marginals.

It immediately follows from the structure of the graph in Figure 3 that $p(\mathbf{z} | \mathbf{w}, \alpha, \beta) = p(\mathbf{z} | \mathbf{w}, \beta)$ and $p(\mathbf{w} | \alpha, \beta) = p(\mathbf{w} | \alpha)$, that is, \mathbf{z} and α are conditionally independent given \mathbf{w} and β , and \mathbf{w} and β are conditionally independent given α . Thus, (17) is equivalent to

$$p(\mathbf{w} | \mathbf{z}, \alpha, \beta) \propto p(\mathbf{z} | \mathbf{w}, \beta) p(\mathbf{w} | \alpha), \quad (18)$$

where the second factor on the right-hand side is given in (14). The first term is the likelihood of \mathbf{w} and β given the data. From (9) it follows that

$$p(\mathbf{z} | \mathbf{w}, \beta) = \frac{\exp\{- (\mathbf{z} - \mathbf{K}\mathbf{w})^H \beta \mathbf{\Lambda}^{-1} (\mathbf{z} - \mathbf{K}\mathbf{w})\}}{\pi^N |\beta^{-1} \mathbf{\Lambda}|}. \quad (19)$$

Since both right-hand factors in (18) are Gaussian densities, $p(\mathbf{w} | \mathbf{z}, \alpha, \beta)$ is also a Gaussian density with the covariance matrix Φ and mean μ given as

$$\Phi = (\mathbf{A} + \beta \mathbf{K}^H \mathbf{\Lambda}^{-1} \mathbf{K})^{-1}, \quad (20)$$

$$\mu = \beta \Phi \mathbf{K}^H \mathbf{\Lambda}^{-1} \mathbf{z}. \quad (21)$$

The matrix $\mathbf{A} = \text{diag}(\alpha)$ is a diagonal matrix that contains the evidence parameters α_l on its main diagonal. Clearly, μ is a maximum a-posteriori (MAP) estimate of the parameter vector \mathbf{w} under the hypothesis \mathcal{H}_i , with Φ being the covariance matrix of the resulting estimates. This completes the model fitting step.

Our next step is to find parameters α and β that maximize the hypothesis posterior $p(\alpha, \beta | \mathbf{z})$ in (16). This density function can be represented as $p(\alpha, \beta | \mathbf{z}) \propto p(\mathbf{z} | \alpha, \beta) p(\alpha, \beta)$, where $p(\mathbf{z} | \alpha, \beta)$ is the evidence term and $p(\alpha, \beta) = p(\alpha) p(\beta)$ is the hypothesis prior. As it was mentioned earlier, it is quite reasonable to choose noninformative priors since we would like to give all possible hypotheses \mathcal{H}_i an equal chance of being valid. This can be achieved by setting ζ , ϵ , κ , and ν to very small values. In fact, it can be easily concluded (see derivations in the appendix) that maximum of the evidence $p(\mathbf{z} | \alpha, \beta)$ coincides with the maximum of $p(\mathbf{z} | \alpha, \beta) p(\alpha, \beta)$ when $\zeta = \epsilon = \kappa = \nu = 0$, which effectively results in the noninformative hyperpriors for α and β .

This formulation of prior distributions is related to automatic relevance determination (ARD) [14, 18]. As a consequence of this assumption, the maximization of the model

posterior is equivalent to the maximization of the evidence, which is known as the evidence procedure [13].

The evidence term $p(\mathbf{z} | \alpha, \beta)$ can be expressed as

$$\begin{aligned} p(\mathbf{z} | \alpha, \beta) &= \int p(\mathbf{z} | \mathbf{w}, \beta) p(\mathbf{w} | \alpha) d\mathbf{w} \\ &= \frac{\exp(-\mathbf{z}^H (\beta^{-1} \mathbf{\Lambda} + \mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H)^{-1} \mathbf{z})}{\pi^N |\beta^{-1} \mathbf{\Lambda} + \mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H|}, \end{aligned} \quad (22)$$

which is equivalent to (5), where conditional independencies between variables have been used to simplify the integrands. In the Bayesian literature this quantity is known as *marginal likelihood* and its maximization with respect to the unknown hyperparameters α and β is a *type-II maximum likelihood* method [19]. To ease the optimization, several terms in (22) can be expressed as a function of the weight posterior parameters μ and Φ as given by (20) and (21). Then, by taking the derivatives of the logarithm of (22) with respect to α and β and by setting them to zero, we obtain its maximizing values as (see also the appendix)

$$\alpha_l = \frac{1}{\Phi_{ll} + |\mu_l|^2}, \quad (23)$$

$$\beta^{-1} = \frac{\text{tr}[\Phi \mathbf{K}^H \mathbf{\Lambda}^{-1} \mathbf{K}] + (\mathbf{z} - \mathbf{K}\mu)^H \mathbf{\Lambda}^{-1} (\mathbf{z} - \mathbf{K}\mu)}{N}. \quad (24)$$

In (23) μ_l and Φ_{ll} denote the l th element of, respectively, the vector μ , and the main diagonal of the matrix Φ . Unlike the maximizing values obtained in the original RVM paper [15, equation (18)], (24) is derived for the extended, more general case of colored additive noise ξ with the corresponding covariance matrix $\beta^{-1} \mathbf{\Lambda}$ arising due to the MF processing at the receiver. Clearly, if the noise is assumed to be white, expressions (23) and (24) coincide with those derived in [15]. Also note that α and β are dependent as it can be seen from (23) and (24).

Thus, for a particular hypothesis \mathcal{H}_i the learning algorithm proceeds by repeated application of (20) and (21), alternated with the update of the corresponding evidence parameters α_i and β from (23) and (24), as depicted in Figure 4, until some suitable convergence criterion has been satisfied. Provided a good initialization of $\alpha_i^{[0]}$ and $\beta^{[0]}$ is chosen,³ the scheme in Figure 4 converges after j iterations to the stationary point of the system of coupled equations (20), (21), (23), and (24). Then, the maximization (1) is performed by selecting the hypothesis that results in the highest posterior (2).

In practice, however, we will observe that during the reestimation some of the hyperparameters α_l diverge, or, in fact, become numerically indistinguishable from infinity given the computer accuracy.⁴ The divergence of some of the hyperparameters enables us to approximate (1) by performing an

³ Later in Section 5 we consider several rules for initializing the hyperparameters.

⁴ In the finite sample size case, however, this will only happen in the high SNR regime. Otherwise, α_l will take large but still finite values. In Section 4.1 we elaborate more on the conditions that lead to convergence/divergence of this learning scheme.




FIGURE 4: Iterative learning of the parameters; the superscript $[j]$ denotes the iteration index.

on-line model selection: starting from the initial hypothesis \mathcal{H}_0 , we prune the hyperparameters that become larger than a certain threshold as the iterations proceed by setting them to infinity. In turn, this sets the corresponding coefficient w_l to zero, thus “switching off” the l th column in the kernel matrix \mathbf{K} and removing the delay T_l from the search space \mathcal{T} . This effectively implements the model selection by creating smaller hypotheses $\mathcal{H}_i < \mathcal{H}_0$ (with fewer basis functions) without performing an exhaustive search over all the possibilities. The choice of the threshold will be discussed in Section 4.

3.2. Extensions to multiple channel observations

In this subsection we extend the above analysis to multiple channel observations or multiple antenna systems. When detecting multipath components any additional channel measurement (either in time, by observing several periods of the sounding sequence $u(t)$, or in space, by using multiple sensor antenna) can be used to increase detection quality. Of course, it is important to make sure that the multipath components are time-invariant within the observation interval. The basic idea how to incorporate several channel observations is quite simple: in the original formulation each hyperparameter α_l was used to control a single weight w_l and thus the single component. Having several channel observations, a single hyperparameter α_l now controls weights representing contribution of *the same* physical multipath component, but present in the different channel observations.

Usage of a single parameter in this case expresses the channel coherence property in the Bayesian framework. The corresponding graphical model that illustrates this idea for a single hyperparameter α_l is depicted in Figure 5. It is interesting to note that similar ideas, though in a totally different context, were adapted to train neural networks by allowing a single hyperparameter to control a group of weights [18]. Note that it is also possible to introduce an individual hyperparameter $\alpha_{p,l}$ for each weight $w_{p,l}$, but this eventually decouples the problem into P separate one-dimensional problems and, as the result, any dependency between the consecutive channels is ignored.

Now, let us return to (9). It can be seen that the weights \mathbf{w}_p capture the structure induced by multiple antennas. However, for the moment we ignore this structure and treat the components of \mathbf{w}_p as a wide-sense stationary (WSS)




FIGURE 5: Usage of α_l in a multiple-observation discrete-time wireless channel model to represent P coherent channel measurements.

process over the individual channels, $p = 1 \cdots P$. We will also allow each sensor to have a different MF. This might not necessarily be the case for wireless channel sounding, but thus a more general situation can be considered. Different matched filters result in different design matrices \mathbf{K}_p , and thus different noise covariance matrices Σ_p , $p = 1 \cdots P$. We will however require that the variance of the input noise remains the same and equals $N_0 = \beta^{-1}$ for all channels, so that $\Sigma_p = N_0 \Lambda_p$, and the noise components are statistically independent among the channels. Then, by defining

$$\begin{aligned} \tilde{\Sigma} &= \beta^{-1} \begin{bmatrix} \Lambda_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \Lambda_P \end{bmatrix}, & \tilde{\mathbf{A}} &= \underbrace{\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ & \ddots \\ \mathbf{0} & \mathbf{A} \end{bmatrix}}_{P \times P \text{ block matrix}}, \\ \tilde{\mathbf{K}} &= \begin{bmatrix} \mathbf{K}_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \mathbf{K}_P \end{bmatrix}, & \tilde{\mathbf{z}} &= \begin{bmatrix} \mathbf{z}_1 \\ \vdots \\ \mathbf{z}_P \end{bmatrix}, & \tilde{\mathbf{w}} &= \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_P \end{bmatrix}, \end{aligned} \quad (25)$$

we rewrite (9) as

$$\tilde{\mathbf{z}} = \tilde{\mathbf{K}} \tilde{\mathbf{w}} + \tilde{\xi}. \quad (26)$$

A crucial point of this system representation is that the hyperparameters α_l are shared by P channels as it can be seen in the structure of the matrix $\tilde{\mathbf{A}}$. This will have a corresponding effect on the hyperparameter reestimation algorithm.

From the structural equivalence of (9) and (26) we can easily infer that (20) and (21) are modified as follows:

$$\Phi_p = (\mathbf{A} + \beta \mathbf{K}_p^H \Lambda_p^{-1} \mathbf{K}_p)^{-1}, \quad (27)$$

$$\mu_p = \beta \Phi_p \mathbf{K}_p^H \Lambda_p^{-1} \mathbf{z}_p, \quad p = 1 \cdots P. \quad (28)$$

The expressions for the hyperparameter updates become a bit more complicated but are still straight-forward to compute. It is shown in the appendix that

$$\alpha_l = \frac{P}{\sum_{p=1}^P (\Phi_{p,ll} + |\mu_{p,l}|^2)}, \quad (29)$$

$$N_0 = \beta^{-1}$$

$$= \frac{1}{NP} \left(\sum_{p=1}^P \text{tr}[\Phi_p \mathbf{K}_p^H \Lambda_p^{-1} \mathbf{K}_p] + \sum_{p=1}^P (\mathbf{z}_p - \mathbf{K}_p \boldsymbol{\mu}_p)^H \Lambda_p^{-1} (\mathbf{z}_p - \mathbf{K}_p \boldsymbol{\mu}_p) \right), \quad (30)$$

where $\mu_{p,l}$ is the l th element of the MAP estimate of the parameter vector \mathbf{w}_p given by (28), and $\Phi_{p,ll}$ is the l th element on the main diagonal of Φ_p from (27). Comparing the latter expressions with those developed for the single channel case, we observe that (29) and (30) use multiple channels to improve the estimates of the noise spectral height and channel weight hyperparameters. They also offer more insight into the physical meaning of the hyperparameters $\boldsymbol{\alpha}$. On the one hand, the hyperparameters are used to regularize the matrix inversion (27), needed to obtain the MAP estimates of the parameters $w_{p,l}$ and their corresponding variances. On the other hand, they act as the inverse of the second noncentral moments of the coefficients $w_{p,l}$, as can be seen from (29).

4. MODEL SELECTION AND BASIS PRUNING

The ability to select the best model to represent the measured data is an important feature of the proposed scheme, and thus it is paramount to consider in more detail how the model selection is effectively achieved. In Section 3.1 we have briefly mentioned that during the learning phase many of the hyperparameters α_l 's tend to large values, meaning that the corresponding weights w_l 's will cluster around zero according to the prior (14). This will allow us to set these coefficients to zero, thus effectively pruning the corresponding basis function from the design matrix. However the question how large a hyperparameter has to grow in order to prune its corresponding basis function has not yet been discussed. In the original RVM paper [15], the author suggests using a threshold α_{th} to prune the model. The empirical evidence collected by the author suggests setting the threshold to "a sufficiently large number" (e.g., $\alpha_{\text{th}} = 10^{12}$). However, our theoretical analysis presented in the following section will show that such high thresholds are only meaningful in very high SNR regimes, or if the number of channel observations P is sufficiently large. In more general, and often more realistic, scenarios such high thresholds are absolutely impractical. Thus, there is a need to study the model selection problem in the context of the presented approach more rigorously.

Below, we present two methods for implementing model selection within the proposed algorithm. The first method relies on the statistical properties of the hyperparameters α_l , when the update equations (27), (28), (29), and (30) converge to a stationary point. The second method exploits the

relationship that we will establish between the proposed scheme and the minimum description length principle [4, 8, 20, 21], thus linking the EP to this classical model selection approach.

4.1. Statistical analysis of the hyperparameters in the stationary point

The decision to keep or to prune a basis function from the design matrix is based purely on the value of the corresponding hyperparameter α_l . In the following we analyze the convergence properties of the iterative learning scheme depicted in Figure 4 using expressions (27), (28), (29), and (30), and the resulting distribution of the hyperparameters once convergence is achieved.

We start our analysis of the evidence parameters α_l by making some simplifications to make the derivations tractable.

- (i) P channels are assumed.
- (ii) The same MF is used to process each of the P sensor output signals, that is, $\mathbf{K}_p = \mathbf{K}$ and $\boldsymbol{\Sigma}_p = \boldsymbol{\Sigma} = \beta^{-1} \mathbf{I}$, $p = 1 \dots P$.
- (iii) The noise covariance matrix $\boldsymbol{\Sigma}$ is known, and $\mathbf{B} = \boldsymbol{\Sigma}^{-1}$.
- (iv) We assume the presence of a single multipath component, that is, $L = 1$, with known delay τ . Thus, the design matrix is given as $\mathbf{K} = [\mathbf{r}(\tau)]$, where $\mathbf{r}(\tau) = [R_{uu}(-\tau), R_{uu}(T_s - \tau), \dots, R_{uu}((N-1)T_s - \tau)]^T$ is the associated basis function.
- (v) The hyperparameter associated with this component is denoted as α .

Our goal is to consider the steady-state solution α_∞ for hyperparameter α in this simplified scenario. In this case (27) and (28) simplify to

$$\begin{aligned} \phi &= (\alpha + \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau))^{-1}, \\ \boldsymbol{\mu}_p &= \phi \mathbf{K}^H \mathbf{B} \mathbf{z}_p = \frac{\mathbf{r}(\tau)^H \mathbf{B} \mathbf{z}_p}{\alpha + \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau)}, \quad p = 1 \dots P. \end{aligned} \quad (31)$$

Inserting these two expressions into (29) yields

$$\alpha^{-1} = \frac{1}{\alpha + \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau)} + \frac{\sum_p |\mathbf{r}(\tau)^H \mathbf{B} \mathbf{z}_p / (\alpha + \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau))|^2}{P}. \quad (32)$$

From (32) the solution α_∞ is easily found to be

$$\alpha_\infty = \frac{(\mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau))^2}{(1/P) \sum_p |\mathbf{r}(\tau)^H \mathbf{B} \mathbf{z}_p|^2 - \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau)}. \quad (33)$$

A closer look at (33) reveals that the right-hand side expression might not always be positive since the denominator can be negative for some values of \mathbf{z}_p . This contradicts the assumption that the hyperparameter α is positive.⁵ A further

⁵ Recall that α^{-1} is the prior variance of the corresponding parameter w . This constrains α to be nonnegative.

analysis of (32) reveals that (29) converges to (33) if and only if the denominator of (33) is positive:

$$\frac{1}{P} \sum_p |\mathbf{r}(\tau)^H \mathbf{B} \mathbf{z}_p|^2 > \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau). \quad (34)$$

Otherwise, the iterative learning scheme depicted in Figure 4 diverges, that is, $\alpha_\infty = \infty$. This can be inferred by interpreting (29) as a nonlinear dynamic system that, at iteration j , maps $\alpha^{[j-1]}$ into the updated value $\alpha^{[j]}$. The nonlinear mapping is given by the right-hand side of (29), where the quantities Φ_p and μ_p depend on the values of the hyperparameters at iteration $j - 1$. In Figure 6 we show several iterations of this mapping that illustrate how the solution trajectories evolve. If condition (34) is satisfied, the sequence of solutions $\{\alpha^{[j]}\}$ converges to a stationary point (Figure 6(a)) given by (33). Otherwise, $\{\alpha^{[j]}\}$ diverges (Figure 6(b)). Thus, (32) is a stationary point only provided the condition (34) is satisfied:

$$\alpha_\infty = \begin{cases} \frac{(\mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau))^2}{\sum_p |\mathbf{r}(\tau)^H \mathbf{B} \mathbf{z}_p|^2 / P - \mathbf{r}(\tau)^H \mathbf{B} \mathbf{r}(\tau)}; & \text{cond. (34) is satisfied,} \\ \infty; & \text{otherwise.} \end{cases} \quad (35)$$

Practically, this means that for a given measurement \mathbf{z}_p , and known noise matrix \mathbf{B} , we can immediately decide whether a given basis function $\mathbf{r}(\tau)$ should be included in the basis by simply checking if (34) is satisfied or not.

A similar analysis is performed in [22], where the behavior of the likelihood function with respect to a single parameter is studied. The obtained convergence results coincide with ours when $P = 1$. Expression (34) is, however, more general and accounts for multiple channel observations and colored noise. In [22] the authors also suggest that testing (34) for a given basis function $\mathbf{r}(\tau)$ is sufficient to find a sparse representation and no further pruning is necessary. In other words, each basis function in the design matrix \mathbf{K} is subject to the test (34) and, if the test fails, that is, (34) does not hold for the basis function under test, the basis function is pruned.

In case of wireless channels, however, we have experimentally observed that even in simulated high-SNR scenarios such pruning results in a significantly overestimated number of multipath components. Moreover, it can be inferred from (34) that, as the SNR increases, the number of functions pruned with this approach decreases, resulting in less and less sparse representations. This motivates us to perform a more detailed analysis of (35).

Let us slightly modify the assumptions we made earlier. We now assume that the multipath delay τ is unknown. The design matrix is constructed similarly but this time $\mathbf{K} = [\mathbf{r}_l]$, where

$$\mathbf{r}_l = [R_{uu}(-T_l), \dots, R_{uu}((N-1)T_s - T_l)]^T \quad (36)$$

is the basis function associated with the delay $T_l \in \mathcal{T}$ used in our discrete-time model. Under these assumptions the input signal \mathbf{z}_p is nothing else but the basis function $\mathbf{r}(\tau)$ scaled





FIGURE 6: Evolution of the two representative solution trajectories for two cases: (a) $\{\alpha^{[j]}\}$ converges, (b) $\{\alpha^{[j]}\}$ diverges.

and embedded in the additive complex zero-mean Gaussian noise with covariance matrix Σ , that is,

$$\mathbf{z}_p = w_p \mathbf{r}(\tau) + \xi_p. \quad (37)$$

Let us further assume that $w_p \in \mathbb{C}$, $p = 1 \dots P$, are unknown but fixed complex scaling factors. In further derivations we assume, unless explicitly stated otherwise, that the condition (34) is satisfied for the basis \mathbf{r}_l . By plugging (37)

into (33) and rearranging the result with respect to α_∞^{-1} we arrive at

$$\begin{aligned} \alpha_\infty^{-1} &= \frac{|\mathbf{r}_l^H \mathbf{B} \mathbf{r}(\tau)|^2 \sum_p |w_p|^2}{P |\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l|^2} \\ &+ \frac{2 \sum_p \text{Re} \{w_p \mathbf{r}_l^H \mathbf{B} \mathbf{r}(\tau) \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P |\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l|^2} \\ &+ \frac{\mathbf{r}_l^H \mathbf{B} (\sum_p \xi_p \xi_p^H) \mathbf{B} \mathbf{r}_l}{P |\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l|^2} - \frac{1}{\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l}. \end{aligned} \quad (38)$$

Now, we consider two scenarios. In the first scenario $\tau = T_l \in \mathcal{T}$, that is, the discrete-time model matches the observed signal. Although unrealistic, this allows to study the properties of α_∞^{-1} more closely. In the second scenario, we study what happens if the discrete-time model does not match perfectly the measured signal. This case helps us to define how the model selection rules have to be adjusted to consider possible misalignment of the path component delays in the model.

4.1.1. Model match: $\tau = T_l$

In this situation, $\mathbf{r}_l = \mathbf{r}(\tau)$, and thus (38) can be further simplified according to

$$\begin{aligned} \alpha_\infty^{-1} &= \frac{\sum_p |w_p|^2}{P} + \frac{2 \sum_p \text{Re} \{w_p \xi_p \mathbf{B} \mathbf{r}_l\}}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)} \\ &+ \frac{\mathbf{r}_l^H \mathbf{B} (\sum_p \xi_p \xi_p^H) \mathbf{B} \mathbf{r}_l}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)^2} - \frac{1}{\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l}, \end{aligned} \quad (39)$$

where the only random quantity is the additive noise term ξ_p . This allows us to study the statistical properties of the finite stationary point in (35).

Equation (39) shows how the noise and multipath component contribute to α_∞^{-1} . If all w_p are set to be zero, that is, there is no multipath component, then $\alpha_\infty^{-1} = \alpha_n^{-1}$ reflects only the noise contribution:

$$\alpha_n^{-1} = \frac{\mathbf{r}_l^H \mathbf{B} (\sum_p \xi_p \xi_p^H) \mathbf{B} \mathbf{r}_l}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)^2} - \frac{1}{\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l}. \quad (40)$$

On the other hand, in the absence of noise, that is, in the infinite SNR case, the corresponding hyperparameter α_∞^{-1} includes the contribution of the multipath component⁶ α_s^{-1} :

$$\alpha_s^{-1} = \frac{\sum_p |w_p|^2}{P} + \frac{2 \sum_p \text{Re} \{w_p \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)}. \quad (41)$$

In a realistic case, both noise and multipath component are present, and α_∞^{-1} consists of the sum of two contributions

$\alpha_\infty^{-1} = \alpha_s^{-1} + \alpha_n^{-1}$. Both quantities α_s^{-1} and α_n^{-1} are random variables with pdf's depending on the number of channel observations P , the basis function \mathbf{r}_l , and the noise covariance matrix Σ . In the sequel we analyze their statistical properties.

We first consider α_s^{-1} . The first term on the right-hand side of (41) is a deterministic quantity that equals the average power of the multipath component. The second one, on the other hand, is random. The product $\text{Re} \{w_p \xi_p^H \mathbf{B} \mathbf{r}_l\}$ in (41) is recognized as the cross-correlation between the additive noise term and the basis function \mathbf{r}_l . It is Gaussian distributed with expectation and variance given as

$$E \left\{ \frac{2 \sum_p \text{Re} \{w_p \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)} \right\} = 0, \quad (42)$$

$$E \left\{ \left(\frac{2 \sum_p \text{Re} \{w_p \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)} \right)^2 \right\} = \frac{2 \sum_p |w_p|^2}{P^2 (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)},$$

respectively, where $E\{\cdot\}$ denotes the expectation operator. Thus, α_s^{-1} is distributed as

$$\alpha_s^{-1} \sim \mathcal{N} \left(\frac{\sum_p |w_p|^2}{P}, \frac{2 \sum_p |w_p|^2}{P^2 (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)} \right), \quad (43)$$

which is a normal distribution with the mean given by the average power of the multipath component and variance proportional to this power.

Now, let us consider the term α_n^{-1} . In (40) the only random element is $\sum_{p=1}^P \xi_p \xi_p^H$. This random matrix is known to have a complex Wishart distribution [23, 24] with the scale matrix Σ and P degrees of freedom. Let us denote

$$\mathbf{c} = \frac{\mathbf{B} \mathbf{r}_l}{\sqrt{P \mathbf{r}_l^H \mathbf{B} \mathbf{r}_l}}, \quad x = \mathbf{c}^H \sum_{p=1}^P \xi_p \xi_p^H \mathbf{c}. \quad (44)$$

It can be shown that x is Gamma-distributed, that is, $x \sim \mathcal{G}(P, \sigma_c^2)$, with the shape parameter P and the scale parameter σ_c^2 given as

$$\sigma_c^2 = \mathbf{c}^H \Sigma \mathbf{c} = \frac{1}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)}. \quad (45)$$

The pdf of x reads

$$p(x | P, \sigma_c^2) = \frac{x^{P-1}}{\Gamma(P) (\sigma_c^2)^P} e^{-x/\sigma_c^2}. \quad (46)$$

The mean and the variance of x are easily computed to be

$$\begin{aligned} E\{x\} &= P \sigma_c^2 = \frac{1}{\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l}, \\ \text{Var}\{x\} &= P (\sigma_c^2)^2 = \frac{1}{P (\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)^2}. \end{aligned} \quad (47)$$

Taking the term $-1/(\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)$ in (40) into account, we introduce a variable $\tilde{\alpha}_n^{-1}$: a zero mean random variable with the pdf

$$p_{\tilde{\alpha}_n^{-1}}(x | P, \sigma_c^2) = \frac{(x - E\{x\})^{P-1}}{\Gamma(P) (\sigma_c^2)^P} e^{-(x - E\{x\})/\sigma_c^2}, \quad (48)$$

⁶ Actually, the second term in the resulting expression vanishes in a perfectly noise-free case, and then $\alpha_s^{-1} = \sum_p |w_p|^2/P$.

which is equivalent to (46), but shifted so as to correspond to a zero-mean distribution. However, it is known that only positive values of α_n^{-1} occur in practice. The probability mass of the negative part of (48) equals the probability that the condition (34) is not satisfied and the resulting α_∞ eventually diverges to infinity and is pruned. Taking this into account the pdf of α_n^{-1} reads

$$p_{\alpha_n^{-1}}(x) = P_n \delta(x) + (1 - P_n) \mathcal{I}^+(x) \tilde{p}_{\alpha_n^{-1}}(x | P, \sigma_c^2), \quad (49)$$

where $\delta(\cdot)$ denotes a Dirac delta function, P_n is defined as

$$P_n = \int_{-1/(\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)}^0 \tilde{p}_{\alpha_n^{-1}}(x | P, \sigma_c^2) dx, \quad (50)$$

and $\mathcal{I}^+(\cdot)$ is the indicator function of the set of positive real numbers:

$$\mathcal{I}^+(x) = \begin{cases} 0 & x \leq 0, \\ 1 & x > 0. \end{cases} \quad (51)$$

A closer look at (49) shows that as P increases the variance of the Gamma distribution decreases, with α_n^{-1} concentrating at zero. In the limiting case as $P \rightarrow \infty$, (49) converges to a Dirac delta function localized at zero, that is, $\alpha_n = \infty$. This allows natural pruning of the corresponding basis function. This situation is equivalent to averaging out the noise, as the number of channel observations grows. Practically, however, P stays always finite, which means that (43) and (49) have a certain finite variance.

The pruning problem can now be approached from the perspective of classical detection theory. To prune a basis function, we have to decide if the corresponding value of α^{-1} has been generated by the noise distribution (49), that is, the *null hypothesis*, or by the pdf of $\alpha_s^{-1} + \alpha_n^{-1}$, that is, the *alternative hypothesis*. Computing the latter is difficult. The problem might be somewhat relaxed by taking the assumption that α_s^{-1} and α_n^{-1} are statistically independent. However proving the plausibility of this assumption is difficult. Even if we were successful in finding the analytical expression for the pdf of the alternative hypothesis, such model selection approach is hampered by our inability to evaluate (43) since the gains w_p 's are not known a priori. However, we can still use (49) to select a threshold.

Recall that the presented algorithm allows to learn (estimate) the noise spectral height $N_0 = \beta^{-1}$ from the measurements. Assuming that we know β , and, as a consequence, the whole matrix \mathbf{B} then, for any basis function \mathbf{r}_l in the design matrix \mathbf{K} and the corresponding hyperparameter α_l , we can decide with a priori specified probability ρ that α_l is generated by the distribution (49). Indeed, let α_{th}^{-1} be a ρ -quantile of (49) such that the probability $P(\alpha^{-1} \leq \alpha_{\text{th}}^{-1}) = \rho$. Since (49) is known exactly, we can easily compute α_{th}^{-1} and prune all the basis functions for which $\alpha_l^{-1} \leq \alpha_{\text{th}}^{-1}$.

4.1.2. Model mismatch: $\tau \neq T_l$

The analysis performed above relies on the knowledge that the true multipath delay τ belongs to \mathcal{T} . Unfortunately,

this is often unrealistic and the model mismatch $\tau \notin \mathcal{T}$ must be considered. To be able to study how the model mismatch influences the value of the hyperparameters we have to make a few more assumptions. Let us for simplicity select the model delay T_l to be a multiple of the chip period T_p . We will also need to assume a certain shape of the correlation function $R_{uu}(t)$ to make the whole analysis tractable. It may be convenient to assume that the main lobe of $R_{uu}(t)$ can be approximated by a raised cosine function with period $2T_p$. This approximation makes sense if the sounding pulse $p(t)$ defined in Section 2 is a square root raised cosine pulse. Clearly, this approximation can also be applied for other shapes of the main lobe, but the analysis of quality of such approximation remains outside the scope of this paper.

Just as in the previous case, we can split the expression (38) into the multipath component contribution α_s^{-1}

$$\alpha_s^{-1} = \frac{|\gamma(\tau)|^2 \sum_p |w_p|^2}{P} + \frac{2 \sum_p \text{Re} \{w_p \gamma(\tau) \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P |\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l|}, \quad (52)$$

where

$$\gamma(\tau) = \frac{\mathbf{r}_l^H \mathbf{B} \mathbf{r}(\tau)}{\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l}, \quad (53)$$

and the same noise contribution α_n^{-1} defined in (40). It can be seen that the $\gamma(\tau)$ makes (52) differ from (41), and as such it is the key to the analysis of the model mismatch. Note that this function is bounded as $|\gamma(\tau)| \leq 1$, with equality following only if $\tau = T_l$. Note also that in our case for $|\tau - T_l| < T_p$ the correlation $\gamma(\tau)$ is strictly positive.

Due to the properties of the sounding sequence $u(t)$, the magnitude of $R_{uu}(t)$ for $|t| > T_p$ is sufficiently small and in our analysis of model mismatch can be safely assumed to be zero. Furthermore, if \mathbf{r}_l is chosen to coincide with the multiple of the sampling period $T_l = lT_s$, then it follows from (12) that the product $\mathbf{r}_l^H \mathbf{B} = \mathbf{r}_l^H \Sigma^{-1} = \beta \mathbf{e}_l^H$ is a vector with all elements being zero except the l th element, which is equal to β . Thus, the product $\mathbf{r}_l^H \mathbf{B} \mathbf{r}(\tau)$ for $|\tau - T_l| < T_p$ must have a form identical to that of the correlation function $R_{uu}(t)$ for $|t| < T_p$. It follows that when $|\tau - T_l| \geq T_p$ the correlation $\gamma(\tau)$ can be assumed to be zero, and it makes sense to analyze (52) only when $|\tau - T_l| < T_p$. In Figure 7 we plot the correlation functions $R_{uu}(t)$ and $\gamma(\tau)$ for this case.

Since the true value of τ is unknown, we assume this parameter to be random, uniformly distributed in the interval $[T_l - T_p, T_l + T_p]$. This in turn induces corresponding distributions for the random variables $\gamma(\tau)$ and $\gamma(\tau)^2$, which enter, respectively, the second and first terms on the right-hand side of (52).

It can be shown that in this case $\gamma(\tau) \sim \mathcal{B}(0.5, 0.5)$, where $\mathcal{B}(0.5, 0.5)$ is a Beta distribution [25] with both distribution parameters equal to 1/2. The corresponding pdf $p_\gamma(x)$ is given in this case as

$$p_\gamma(x) = \frac{1}{B(0.5, 0.5)} x^{-1/2} (1-x)^{-1/2}, \quad (54)$$

where $B(\cdot, \cdot)$ is a Beta-function [26] with $B(0.5, 0.5) = \pi$.




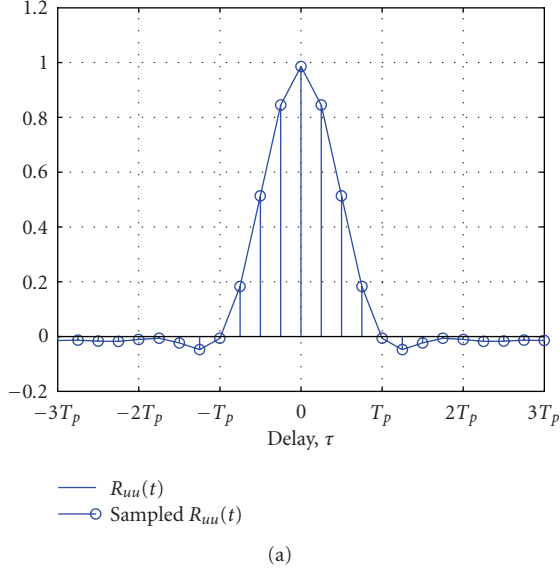


FIGURE 7: Evaluated correlation functions (a) $R_{uu}(t)$ and (b) $\gamma(\tau)$.

It is also straight-forward to compute the pdf of the term $\gamma(\tau)^2$:

$$p_{\gamma^2}(x) = \frac{1}{\pi} x^{-3/4} (1 - \sqrt{x})^{-1/2}. \quad (55)$$

The corresponding empirical and theoretical pdf's of $\gamma(\tau)$ and $\gamma(\tau)^2$ are shown in Figure 8.

Now we have to find out how this information can be utilized to design an appropriate threshold. In the case of a perfectly matched model the threshold is selected based on the noise distribution (49). In the case of a model mismatch, the term (52) measures the amount of the interference resulting from the model imperfection.

Indeed, if $|\tau - T_l| \geq T_p$, then the resulting $\gamma(\tau) = 0$, and thus $\alpha_s^{-1} = 0$. The corresponding evidence parameter α_∞^{-1} is then equal to the noise contribution α_n^{-1} only and will be pruned using the method we described for the matched

FIGURE 8: Comparison between the empirical and theoretical pdf's of (a) $\gamma(\tau)$ and (b) $\gamma(\tau)^2$ for the cosine approximation case. To compute the histogram $N = 5000$ samples were used.

model case. If however, $|\tau - T_l| < T_p$, then a certain fraction of α_s^{-1} will be added to the noise contribution α_n^{-1} , thus causing the interference. In order to be able to take this interference into account and adjust the threshold accordingly, we propose the following approach.

The amount of interference added is measured by the magnitude of α_s^{-1} in (52). It consists of two terms: the first one is the multipath power, scaled by the factor $\gamma(\tau)^2$:

$$\gamma(\tau)^2 \frac{\sum_p |w_p|^2}{P}. \quad (56)$$

The second term is a cross product between the multipath component and the additive noise, scaled by $\gamma(\tau)$:

$$\gamma(\tau) \frac{2 \sum_p \text{Re} \{w_p \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P(\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)}. \quad (57)$$

Both terms have the same physical interpretation as in (41), but with scaling factors $\gamma(\tau)$ depending on the true value of τ .

We see that in (52) there are quite a few unknowns: we do not know the true multipath delay τ , the multipath gains w_p , as well as the instantaneous noise value ξ . To be able to circumvent these uncertainties, we consider the large sample size case, that is, $P \rightarrow \infty$ and invoke the law of large numbers to approximate (56) and (57) by their expectations.

First of all, using (42) it is easy to see that

$$E \left\{ \gamma(\tau) \frac{2 \sum_p \text{Re} \{w_p \xi_p^H \mathbf{B} \mathbf{r}_l\}}{P(\mathbf{r}_l^H \mathbf{B} \mathbf{r}_l)} \right\} = 0. \quad (58)$$

The other term (56) converges to $\gamma(\tau)^2 E\{|w_p|^2\}$ as P grows. So, even in the high SNR regime and infinite number of channel observations P the term (56) does not go to zero. In order to assess how large it is, we approximate the gains of the multipath component w_p by the corresponding MAP estimate μ_p obtained with (28).

The correlation function $\gamma(\tau)$ can also be taken into account. Since we know the distributions of both $\gamma(\tau)$ and $\gamma(\tau)^2$, we can summarize these by the corresponding mean values. In fact, we will need the mean only for $\gamma(\tau)^2$ since it enters the irreducible part of α_s^{-1} .

In our case it is computed as

$$E\{\gamma(\tau)^2\} = \int_0^1 \frac{x}{\pi} x^{-3/4} (1 - \sqrt{x})^{-1/2} dx = \frac{3}{8}. \quad (59)$$

Having obtained the mean, we can approximate the interference $\hat{\alpha}_s^{-1}$ due to the model mismatch as

$$\hat{\alpha}_s^{-1} = \frac{3}{8} \times \frac{\sum_{p=0}^{P-1} |\mu_p|^2}{P}. \quad (60)$$

The final threshold that accounts for the model mismatch is then obtained as

$$\hat{\alpha}_{\text{th}}^{-1} = \hat{\alpha}_s^{-1} + \alpha_{\text{th}}^{-1}, \quad (61)$$

where α_{th}^{-1} is the threshold developed earlier for the matched model case.

4.2. Improving the learning algorithm to cope with the model selection

In the light of the model selection strategy considered here, we anticipate two major problems arising with the learning algorithm discussed in Section 3. The first one is the estimation of the channel parameters that requires computation of the posterior (27). Even for the modest sizes of the hypothesis \mathcal{H}_i (from 100 to 200 basis functions), the matrix inversion is computationally very intensive. This issue becomes even more critical if we consider a hardware implementation of the estimation algorithm. The second problem arises

due to the nonvanishing correlation between the basis vectors \mathbf{r}_l constituting the design matrix \mathbf{K} . A very undesirable consequence of this correlation is that the evidence parameters α_l associated with these vectors become also correlated, and thus no longer represent the contribution of a single basis function. As a consequence the developed model selection rules are no longer applicable.

It is, however, possible to circumvent these two difficulties by modifying the learning algorithm as discussed below. The basic idea consists of estimating the channel parameters for each basis independently. In other words, instead of solving (27), (28), (29), and (30) jointly for all L basis functions, we find a solution for each basis vector separately. First, the new data vector $\mathbf{x}_{p,l}$ for the l th basis is computed as

$$\mathbf{x}_{p,l} = \mathbf{z}_p - \sum_{k=1, k \neq l}^L \mathbf{r}_k \mu_{p,k}. \quad (62)$$

This new data vector $\mathbf{x}_{p,l}$ now contains the information relevant to the basis \mathbf{r}_l only. It is then used to update the corresponding posterior statistics as well as evidence parameters exclusively for the l th basis as follows:

$$\begin{aligned} \Phi_l &= (\alpha_l + \beta \mathbf{r}_l^H \mathbf{\Lambda}^{-1} \mathbf{r}_l)^{-1}, \\ \mu_{p,l} &= \beta \Phi_l \mathbf{r}_l^H \mathbf{\Lambda}^{-1} \mathbf{x}_{p,l}, \quad p = 1 \cdots P. \end{aligned} \quad (63)$$

Note that expressions (63) are now scalars, unlike their matrix counterparts (27) and (28). Similarly, we update the evidence parameters as

$$\alpha_l = \frac{P}{\sum_{p=1}^P (\Phi_l + |\mu_{p,l}|^2)}. \quad (64)$$

Updates (63) and (64) are performed for all L components sequentially. Once all components are updated, we update the noise hyperparameter N_0 :

$$\begin{aligned} N_0 &= (\beta^{-1}) \\ &= \frac{1}{NP} \left(\sum_{p=1}^P \text{tr}[\Phi(\mathbf{K})^H \mathbf{\Lambda}^{-1} \mathbf{K}] \right. \\ &\quad \left. + \sum_{p=1}^P (\mathbf{z}_p - \mathbf{K} \boldsymbol{\mu}_p)^H \mathbf{\Lambda}^{-1} (\mathbf{z}_p - \mathbf{K} \boldsymbol{\mu}_p) \right). \end{aligned} \quad (65)$$

The above updating procedures constitute a single iteration of the modified learning algorithm. This iteration is repeated until some suitable convergence criterion is satisfied. Note that the procedure described here is an instance of the SAGE algorithm. This opens a potential to unite both SAGE and evidence procedure, allowing to implement simultaneous parameter and model order estimation within the SAGE framework.

This iterative method, also known as successive interference cancellation, allows solving both anticipated problems. First of all, there is no need to compute matrix inversion at

each iteration. Second, the obtained values of α now reflect the contribution of a single basis function only, since they were estimated while the contribution of other bases was canceled in (62).

Now, at the end of each iteration, once the new value of the noise is obtained using (65), we can decide to prune some of the components, as described in Section 4.1.

4.3. MDL principle and evidence procedure

The goal of this section is to establish a relationship between the classical information-theoretic criteria for model selection, such as minimum description length (MDL) [4, 5, 8, 20], and the evidence procedure discussed here. For simplicity we will only consider a single channel observation case, that is, $P = 1$. Extension to the case $P > 1$ is straightforward.

The MDL criterion was originally formulated from the perspective of coding theory as a solution to the problem of balancing the code length and the resulting length of the data encode with this code. This concept however can naturally be transferred to general model selection problems.

In terms of parameter estimation theory, we can interpret the length of the encoded data as the parameter likelihood evaluated at its maximum. The length of the code is equivalent to what is known in the literature as the *stochastic complexity* [11, 20, 21]. The Bayesian interpretation of the stochastic complexity term obtained for likelihood functions from an exponential family (see [20] for more details) is of particular interest for our problem at hand. The description length in this case is given as

$$\begin{aligned} \text{DL}(\mathcal{H}_i) &= \underbrace{-\log(p(\mathbf{z} | \mathbf{w}_{\text{MAP}}, \mathcal{H}_i))}_{\text{model performance}} \\ &+ \underbrace{\frac{L}{2} \log \frac{N}{2\pi} - \log(p(\mathbf{w}_{\text{MAP}} | \mathcal{H}_i)) + \log(\sqrt{|\mathbf{I}_1(\mathbf{w}_{\text{MAP}})|})}_{\text{stochastic complexity}}. \end{aligned} \quad (66)$$

Here $\mathbf{I}_1(\mathbf{w}_{\text{MAP}})$ is the Fisher information matrix of a single sample evaluated at the MAP estimate of the model parameter vector, and $p(\mathbf{w}_{\text{MAP}} | \mathcal{H}_i)$ is the corresponding prior for this vector.

Thus, joint model and parameter estimation schemes should aim at minimizing the DL so as to find the compromise between the model fit (likelihood) and the number of the parameters involved. The latter is directly proportional to the stochastic complexity term.

We will now show that the EP employed in our model selection scheme results in a very similar expression.

Let us once again come back to the evidence term (22). To exemplify the main message that we want to convey here, we will compute the integral in (22) differently. For each model hypothesis defined as in Section 3, let us define $\Delta(\mathbf{w}_i) = -\log(p(\mathbf{z} | \mathbf{w}_i, \beta_i)) - \log(p(\mathbf{w}_i | \alpha_i))$. Then (22) can be expressed as

$$p(\mathbf{z} | \alpha_i, \beta_i) = \int \exp(-\Delta(\mathbf{w}_i)) d\mathbf{w}_i. \quad (67)$$

Now we proceed by computing the integral (67) using a Laplace method [8, Chapter 27], also known as a saddle-point approximation. The essence of the method consists of computing the second-order Taylor series around the argument that maximizes the integrand in (67), which is the MAP estimate of the model parameters μ_i given in (21). In our case, $\Delta(\mathbf{w}_i)$ is known to be quadratic, since both $p(\mathbf{z} | \mathbf{w}_i, \beta_i)$ and $p(\mathbf{w}_i | \alpha_i)$ are Gaussian, so the approximation is exact.

It is then easily verified that for the hypothesis \mathcal{H}_i with $|\mathcal{P}(i)| = L$ basis functions

$$\begin{aligned} p(\mathbf{z} | \alpha_i, \beta_i) &= \int \exp(-(\mathbf{w}_i - \mu_i)^H \Phi_i^{-1} (\mathbf{w}_i - \mu_i)) d\mathbf{w}_i \\ &\times \exp(-\Delta(\mu_i)) = \exp(-\Delta(\mu_i)) \pi^L |\Phi_i|. \end{aligned} \quad (68)$$

By taking the logarithm of (68) and changing the sign of the resulting expression we arrive at the final expression for the negative log-evidence:

$$\begin{aligned} -\log(p(\mathbf{z} | \alpha_i, \beta_i)) &= -\log(p(\mathbf{z} | \mu_i, \beta_i)) \\ &- \log(p(\mu_i | \alpha_i)) - L \log(\pi) \\ &- \log(|\Phi_i|). \end{aligned} \quad (69)$$

Noting that Φ_i has been computed using N data samples, and that in this case $\log(|\Phi_i/N|) = \log(|\mathbf{I}_1^{-1}(\mu_i)|)$, we rewrite (69) as

$$\begin{aligned} \text{DL}(\mathcal{H}_i) &= \underbrace{-\log(p(\mathbf{z} | \mu_i, \beta_i))}_{\text{model performance}} \\ &+ \underbrace{L \log\left(\frac{N}{\pi}\right) - \log(p(\mu_i | \alpha_i)) + \log(|\mathbf{I}_1(\mu_i)|)}_{\text{model complexity}}. \end{aligned} \quad (70)$$

We note that (66) and (70) are essentially similar, with the distinction that the latter accounts for complex data. Thus we conclude that maximizing evidence (or minimizing the negative log-evidence) is equivalent to minimizing the DL.

Let us now consider how this can be exploited in our case. In general, the MDL concept assumes presence of multiple *estimated* models. The model that minimizes the DL functional is then picked as the optimal one. In our case, evaluation of the DL functional for all possible hypotheses \mathcal{H}_i is way too complex. In order to make this procedure more efficient, we can exploit the estimated evidence information.

Consider the graph shown in Figure 9. Each node on the graph corresponds to a certain hypothesis \mathcal{H}_i consisting of $|\mathcal{P}_i|$ basis functions. An edge emanating from a node is associated with a certain basis function from the hypothesis \mathcal{H}_i . Should the path through the graph include this edge, the corresponding basis function would be pruned, leading to a new smaller hypothesis. Clearly, the optimal path through the graph should be the one that minimizes the DL criterion. Now, let us propose a strategy to find the optimal model without evaluating all possible paths through the graph.




FIGURE 9: Model selection by evidence evaluation.

At the initial stage, we start in the leftmost node, which corresponds to the full hypothesis \mathcal{H}_0 . We then proceed with the learning using the iterative scheme depicted in Figure 4 to obtain the estimates of the evidence parameters $\alpha_l, l \in \mathcal{P}(0)$, for each basis function in \mathcal{H}_0 . Once convergence is achieved, we evaluate the corresponding description length DL_0 for this hypothesis using (70). Since the optimal path should decrease the DL, the hypothesis at the next stage \mathcal{H}_i is selected by moving along the edge that corresponds to the basis function with the largest value of α (i.e., the basis function with the smallest evidence). For the newly selected hypothesis \mathcal{H}_i we again estimate the evidence parameters α_i and the corresponding description length DL_i . If $DL_0 < DL_i$, then the hypothesis \mathcal{H}_0 achieves the minimum of the description length and it is then selected as a solution. Otherwise, that is, if $DL_0 > DL_i$, we continue along the graph, each time pruning a basis function with the smallest evidence and comparing the description length at each stage. We proceed so until the DL does not decrease any more, or until we stop at the last node that has no basis functions at all. Such an empty hypothesis corresponds to the case when there is no structure in the observed data. In other words, it corresponds to the case when the algorithm failed to find any multipath components. This technique requires searching between L_0 to a maximum of $L_0(L_0 + 1)/2$ possible hypotheses, while a total search requires testing a total of 2^{L_0} different models.

5. APPLICATION OF THE RVM TO WIRELESS CHANNELS

The application of the proposed channel estimation scheme coupled with the considered model selection approach requires two major components: (1) it needs a proper construction of the kernel design matrix that is dense enough to ensure good delay resolution, and (2) the iterative nature of the algorithm requires a good initialization.

The construction of the design matrix \mathbf{K} can be done with various approaches, depending on how much a priori information we have about the possible positions of the multipath components. The columns of the matrix \mathbf{K} contain the shifted versions of the kernel $R_{uu}(nT_s - T_l), l = 1 \cdots L_0$, where T_l are the possible positions of the multipath

components that form the search space \mathcal{T} . The delays T_l can be selected uniformly to cover the whole delay span or might be chosen so as to sample some areas of the impulse response more densely, where multipath components are likely to appear. Note that the delays T_l are not constrained to fall on a regular grid. The power-delay profile (PDP) may be a good indicator of how to place the multipath components. Initialization of the model hyperparameters can also be done quite effectively. In the sequel we propose two different initialization techniques.

The simplest one consists of evaluating the condition (34) for all the basis functions in the already created design matrix \mathbf{K} . For those basis functions that satisfy condition (34), the corresponding evidence parameter is initialized using (33). Other basis functions are removed from the design matrix \mathbf{K} . Such initialization assumes that there is no interference between the neighboring basis functions. It makes sense to employ it when the minimal spacing between the elements in \mathcal{T} is at most half the duration of the sounding pulse T_p .

In the case when the spacing is denser, it is better to use independent evidence initialization. This type of initialization is in fact coupled with the construction of the design matrix \mathbf{K} and relies on the successive interference cancellation scheme discussed in Section 4.2. To make the procedure work, we need to set the initial channel coefficients to zero, that is, $\mu_p \equiv 0$. The basis vectors \mathbf{r}_l are computed as usual according to the delay search space \mathcal{T} . The initialization iterations start by computing (62). The basis \mathbf{r}_l that is best aligned with the residual $\mathbf{x}_{p,l}$ is then selected. If the selected \mathbf{r}_l satisfies condition (34), it is included in the design matrix \mathbf{K} , and the corresponding parameters $\Phi_l, \mu_{p,l}$, and α_l are computed according to (63) and (64), respectively. These steps are continued until all bases with delays from the search space \mathcal{T} are initialized, or until the basis vector that does not satisfy the condition (34) is encountered.

Of course, in order to be able to use this initialization scheme, it is crucial to get a good initial noise estimate. The initial noise parameter $N_0^{[0]}$ can in most cases be estimated from the tails of the channel impulse response, where multipath components are unlikely to be present or too weak to be detected. Generally, we have observed that the algorithm is less sensitive to the initial values of the hyperparameters α , but proper initialization of the noise spectral height is crucial.

Now we can describe the simulation setup used to assess the performance of the proposed algorithm.

5.1. Simulation setup

The generation of the synthetic channel is done following the block-diagram shown in Figure 1: a single period $u(t)$ of the sounding sequence $s(t)$ is filtered by the channel with the impulse response $\mathbf{h}(t)$, and complex white Gaussian noise is added to the channel outputs to produce the received signal $\mathbf{y}(t)$. The received signal is then run through the MF. The continuous-time signals at the output of the MF are represented with cubic splines. The resulting spline representation

is then used to obtain the sampled output $z_p[n]$, $p = 1 \cdots P$, with $n = 0 \cdots N - 1$. Output signals $z_p[n]$ are then used as the input to the estimation algorithm.

For all P channel observations we use the same MF, and thus $\Phi = \Phi_p$, $\mathbf{K} = \mathbf{K}_p$, and $\Sigma = \Sigma_p$, $p = 1 \cdots P$. Without loss of generality, we assume a shaping pulse of the duration $T_p = 10$ nanoseconds. The sampling period is assumed to be $T_s = T_p/N_s$, where N_s is the number of samples per chip used in the simulations. The sounding waveform $u(t)$ consists of $M = 255$ chips. We also assume the maximum delay spread in all simulations to be $\tau_{\text{spread}} = 1.27$ microseconds. With these parameters, a one-sample/chip resolution results in $N = 128$ samples. The autocorrelation function $R_{uu}(t)$ is also represented with cubic splines, allowing a proper construction of the design matrix \mathbf{K} according to the predefined delays in \mathcal{T} . Realizations of the channel parameters $w_{l,p}$ are randomly generated according to (14).

The performance of the algorithm is also evaluated under different SNR's at the output of the MF, defined as

$$\text{SNR} = 10 \log_{10} \left(\frac{1/\alpha}{N_0} \right). \quad (71)$$

For simplicity, we assumed that in the case $L > 1$ all simulated multipath components have the same expected power α^{-1} . Although this is not always a realistic assumption, it ensures that all simulated multipath components present in the measurement will be "treated" equally.

5.2. Numerical simulations

Let us now demonstrate the performance of the model selection schemes discussed in Section 4 on synthetic, as well as on measured channels.

5.2.1. Multipath detection with the perfect model match

First we consider the distribution of the hyperparameters once the stationary point has been reached. In order to do that, we apply the learning algorithm to the full hypothesis \mathcal{H}_0 . The delays in \mathcal{H}_0 are evenly positioned over the length of the impulse response: $\mathcal{T} = \{lT_s; l = 0 \cdots N - 1\}$, that is, $L_0 = N$. Here, we simulate the channel with a single multipath component, that is, $L = 1$, having the delay τ' equal to a multiple of the sampling period T_s . Thus, in the design matrix \mathbf{K} corresponding to the full hypothesis \mathcal{H}_0 there will be a basis function that coincides with the contribution of the true multipath component. Once the parameters have been learned, we partition all the hyperparameters α into those attributed to the noise, that is, α_n , and one parameter that corresponds to the multipath component α_s , that is, the one associated with the delay $T_l = \tau'$.

In a next step, we compare the obtained histogram of α_n^{-1} with the theoretical pdf $p_{\alpha_n^{-1}}(x)$ given in (49). The corresponding results are shown in Figure 10(a). A very good match between the empirical and theoretical pdf's can be observed.

Similarly, we investigate the behavior of the negative log-evidence versus the size of the hypothesis. We consider a

similar simulation setup as above, however with more than just one multipath component to make the results more realistic. Figure 10(b) depicts the evaluated negative log-evidence (69) as a function of the model order, evaluated for a single realization, when the true number of components is $L = 20$, and the number of channel observations is $P = 5$.

Note that, as the SNR increases, there are fewer components subject to the initial pruning, that is, those that do not satisfy condition (34). We also observe that the minimum of the negative log-evidence (i.e., maximum of the evidence) becomes more pronounced as the SNR increases, which has an effect of decreasing the variance of the model order estimates.

In order to find the best possible performance of the algorithm, we first perform some simulations assuming that the discrete-time model (9) perfectly matches the continuous-time model (7), that is, $\tau_l \in \mathcal{T}$, $l = 1, \dots, L$. This is realized by drawing uniformly L out of N possible delay values in the interval $[0, T_s(N-1)]$. Again, $\mathcal{T} = \{lT_s; l = 0 \cdots N-1\}$. The number of multipath components in the simulated channels is set to $L = 5$ and the channel is sampled with $N_s = 2$ samples per chip.

In this simulation we evaluate the detection performance by counting the errors made by the algorithms. Two types of errors can occur: (a) an *insertion error*—an erroneous detection of a nonexisting component; (b) a *deletion error*—a loss of an existing component. The case when an estimated delay \hat{T}_l matches one of the true simulated delays is called a *hit*. We further define the *multipath detection rate* as the ratio between the number of hits to the true number of components L plus the number of insertion errors. It follows that the detection rate is equal to 1 only if the number of hits equals the true number of components. If, however, the algorithm makes any deletion or insertion errors, the detection rate is then strongly smaller than 1. We study the detection rates for both model selection schemes versus different SNR's. The presented results are averaged over 300 independent channel realizations.

We start with the model selection approach based on the threshold selection using the ρ -quantile of the noise distribution-quantile-based model selection. The results shown in Figure 11(a) are obtained for $\rho = 1 - 10^{-6}$ and different numbers of channel observations P . It can be seen that, as P increases, the detection rate significantly improves. To obtain the results shown in Figure 11(b) we fix the number of channel observations at $P = 5$ and vary the value of the quantile ρ . It can be seen that as ρ approaches unity, the threshold is placed higher, meaning that fewer noise components can be mistakenly detected as multipath components, thus slightly improving the detection rate. However higher thresholds require a higher SNR to achieve the same detection rate, as compared for the thresholds obtained with lower ρ .

The next plot in Figure 11(c) shows the multipath detection rate when the model is selected based on the evaluation of the negative log-evidence under different model hypotheses (negative log-evidence model selection). It is interesting to note that in this case the reported curves behave

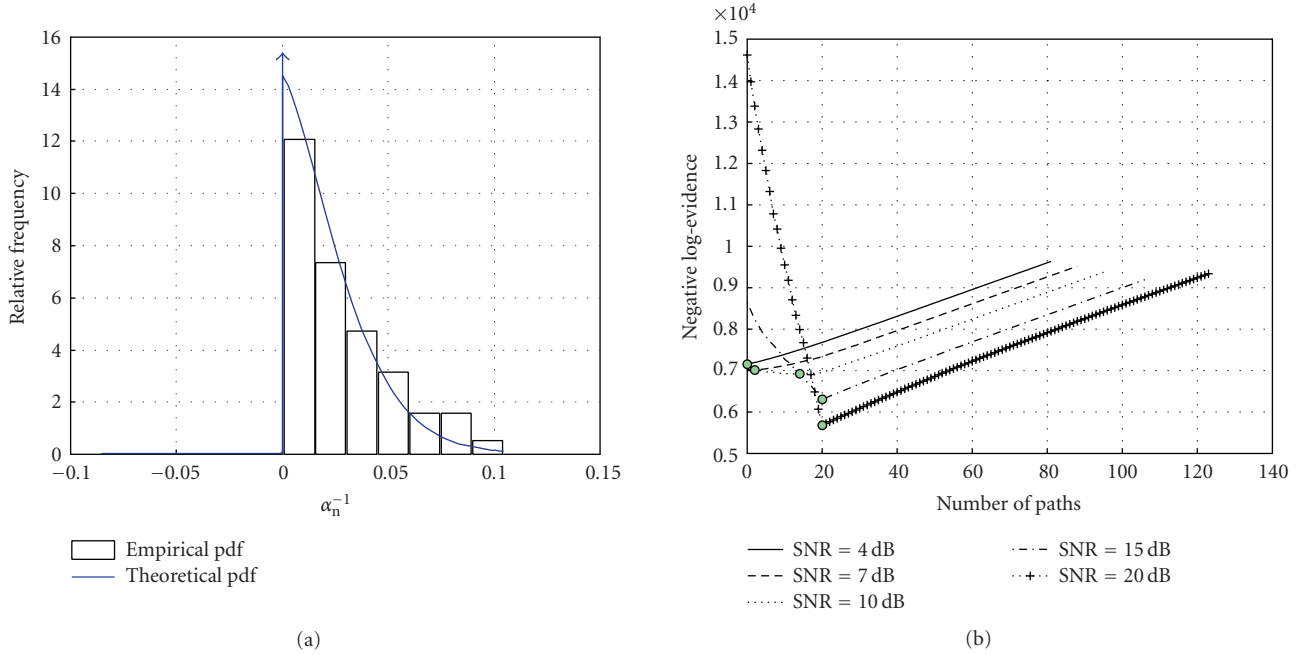


FIGURE 10: Evidence-based model selection criteria. (a) Empirical (bar plot) and theoretical (solid line) pdf's of hyperparameters α_n^{-1} (SNR = 10 dB, and $P = 10$). To compute the histogram $N = 500$ samples were used; (b) Negative log-evidence as a function of the model order (number of paths) for different SNR values ($P = 5$, and $L = 20$).




FIGURE 11: Multipath detection rates based on the EP. (a) Quantile-based model selection versus P : $\rho = 1 - 10^{-6}$, $L = 5$; (b) Quantile-based model selection versus ρ : $P = 5$, $L = 5$; (c) Negative log-evidence-based detection versus P .

quite differently from those shown in Figure 11(a). First, we see that for the case $P = 1$ the behavior of this method is slightly better, compared to the threshold-based method in Figure 11(a). But as P grows, the performance of the multipath detection does not increase proportionally, but rather exhibits a threshold-like behavior. In other words, multipath detection based on the negative log-evidence and alike MDL-based model selection requires the SNR above a certain threshold in order to operate reliably. Furthermore,

this threshold is independent of the number of channel observations P .

Thus from Figures 11(a) and 11(c) we can conclude that the quantile-based method performs better in a sense that it can always be improved by increasing the number of channel observations. Further, model selection using the thresholding approach can be performed on-line, concurrent with parameters estimation, while in the other case multiple models have to be learned.

Now, let us consider how the EP performs when the multipath component delays are on the real line, rather than on a discrete grid. Clearly, this case corresponds more to the real-life situation.

5.2.2. Multipath detection with the model mismatch

In the real world the delays of the multipath components do not necessarily coincide with the elements in \mathcal{T} used to approximate the continuous-time model (7). By using the discrete-time models to approximate the continuous-time counterparts, we would necessarily expect some performance degradation in terms of an increased number of components. This problem is similar to the problem that occurs in fractional delay filters (FDF) [27]. An FDF aims at approximating a delay that is not a multiple of the sampling period. As shown in [27], such filters have infinite impulse response. Though FIR approximations exist, they require several samples to represent a single delay.

Since there is an inevitable mismatch between the continuous-time and discrete-time models, it is worth asking how densely we should quantize the delay line to form the design matrix in order to achieve the best performance. It is convenient to select the delays in \mathcal{T} of the discrete-time model as a multiple of the sampling period T_s . As the sampling rate increases, the true delay values get closer to some elements in \mathcal{T} , thus approaching the continuous-time model (7).

We simulate a channel with a single multipath component that has a random delay, uniformly distributed in the interval $[0, \tau_{\text{spread}}]$.

The criterion used here to assess the performance of the algorithm is the probability of correct path extraction. This probability is defined to be the conditional probability that, given any path is detected, the algorithm finds exactly one component with the absolute difference between the estimated and the true delay less than the chip pulse duration T_p . Notice that the probability of correct path extraction is conditioned on the path detection, that is, it is evaluated for the cases when the estimation algorithm is able to find at least one component.

It is also interesting to compare the performance of the EP with other parameter estimation techniques. Here we consider the SAGE algorithm [2] that has become a popular multipath parameter estimation technique. The SAGE algorithm, however, does not provide any information about the number of multipath components. To make the comparison fair, we augment it with the standard MDL criterion [4, 5] to perform model selection.

Thus, we are going to compare three different model selection algorithms: the quantile-based (or threshold-based) scheme with a preselected quantile $\rho = 1 - 10^{-6}$, the SAGE + MDL method, and negative log-evidence method. We are also going to use the threshold-based method to demonstrate the difference between two EP initialization schemes: the joint initialization, and the independent initialization, discussed in Section 5. In all simulations the

negative log-evidence method was initialized using independent initialization.

We start with channels sampled with $N_s = 1$ sample/chip resolution and $P = 5$ channel observations. We see that the shown methods have different probabilities of path detection (Figure 12(a)), that is, they require different SNR to achieve the same path detection probability. The threshold-based methods can be, however, adjusted by selecting the quantile ρ appropriately. As we see, with $\rho = 1 - 10^{-6}$, the threshold-based and SAGE+MDL methods achieve the same probabilities of path detection. The resulting probabilities of correct path extraction are shown in Figure 12(b). Note that for low SNR comparison of the methods is meaningless, since too few paths are detected. However, above SNR ≈ 15 dB, with all methods we can achieve similar high path detection probability, which allows direct comparison of the correct path extraction probabilities. We can hence infer that, in this regime, model selection with negative log-evidence is superior to other methods, since it has higher probabilities of path extraction. In other words, this means that at higher SNR this method will introduce fewer artifacts.

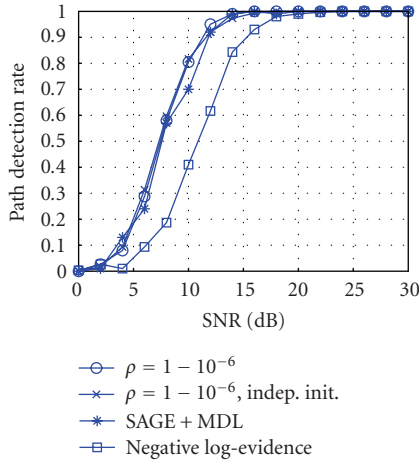
What is also important is that as the SNR increases, the correct path extraction rate drops. This happens simply because our model has a fixed resolution in the delay. As the result, at the higher SNR several components from the our model are used to approximate a single one with a delay between the sampling instances. This leads to the degradation of the correct path extraction rate since the number of components is overestimated.

Now, let us increase the sampling rate and study the case $N_s = 2$ (Figures 12(c) and 12(d)). We see that the probabilities of path extraction are now higher for all methods. A slight difference between the two EP initialization schemes can also be observed. Note however that the performance increase is higher for the SAGE + MDL and negative log-evidence algorithms, which both rely on the same model selection concept.


Finally, the last case with $N_s = 4$ is shown in Figures 12(e) and 12(f). Again SAGE + MDL and negative log-evidence schemes achieve higher correct path extraction probabilities as compared to the threshold-based method. The performance of the latter also increases with the sampling rate, but unfortunately not as fast as that of the description length-based model selection. These plots also demonstrate the difference between the two proposed initializations of the EP. In Figure 12(e) we see that in this case the independent initialization outperforms the joint one. As already mentioned, this distinction becomes noticeable, once the basis functions in \mathbf{K} exhibit significant correlation, which is the case for $N_s \gtrsim 2$.

5.3. Results for measured channels


We also apply the proposed algorithm to the measured data collected in in-door environments. Channel measurements were done with the MIMO channel sounder PropSound manufactured by Elektrobit Oy. The basic setup for channel sounding is equivalent to the block-diagram shown in Figure 1. In the conducted experiment the sounder operated




(a)




(b)



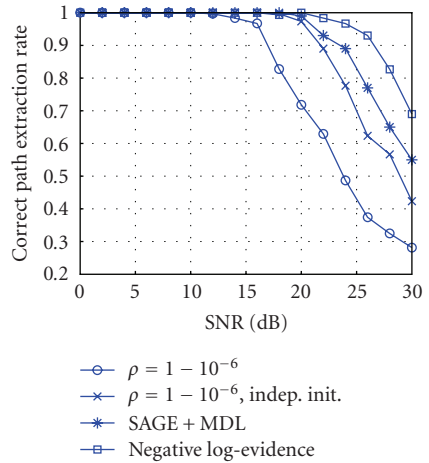
(c)



(d)



(e)



(f)

FIGURE 12: Comparison of the model selection schemes in a single path scenario. (a), (c), (e) path detection probability, and (b), (d), (f) probability of correct path extraction for $P = 5$, and (a), (b) $N_s = 1$; (c), (d) $N_s = 2$; and (e), (f) $N_s = 4$.

at the carrier frequency 5.2 GHz with a chip period of $T_p = 10$ nanoseconds. The output of the matched filter was sampled with the period $T_s = T_p/2$, thus resulting in a resolution of 2 samples per chip. The sounding sequence consisted of $M = 255$ chips, resulting in the burst waveform duration of $T_u = MT_p = 0.255$ microseconds.

Based on visual inspection of the PDP of the measured channels, the delays T_l in the search space \mathcal{T} are positioned uniformly in the interval between 250 nanoseconds and 1000 nanoseconds, with spacing between adjacent delays equal to T_s . This corresponds to the delay search space \mathcal{T} consisting of 151 elements. The initial estimate of the noise floor is obtained from the tail of the measured PDP. The algorithm stops once the relative change of the evidence parameters between two successive iterations is smaller than 0.0001%. The corresponding detection results for different number of channel observations are shown in Figure 13.

When $P = 1$ (see Figure 13(a)), the independent initialization results in only 9 basis functions constituting the initial hypothesis \mathcal{H}_0 . The final estimated number of components is found to be $L = 8$. As expected, increasing the number of channel observations P makes it possible to detect and estimate components with smaller SNR. For the case of $P = 5$ we detect already $L = 12$ components (Figure 13(b)), and for $P = 32$, $L = 15$ components (Figure 13(c)). This shows that increasing the number of observations not necessarily brings a proportional increase of the detected components, thus suggesting that there might be a limit given by the true number of multipath components.

6. CONCLUSION

This paper demonstrates the application of the evidence procedure to the analysis of wireless channels. The original formulation of this method, known as relevance vector machines, was reformulated to cope with the estimation of wireless channels. We extended the method to the complex domain and colored additive noise. We further extended the RVM to multiple channels by proposing a new graphical Bayesian model, where a single evidence parameter controls each multipath component observed with multiple channels. To our knowledge this is a new concept that can be useful not only for estimation, but also for simulating wireless channels.

Evidence parameters were originally introduced to control the sparsity of the model. Assuming a single path scenario we were able to find the statistical laws that govern the values of the evidence parameters once the estimation algorithm has converged to the stationary point. It was shown that in low SNR scenarios the evidence parameters do not attain infinite values, as has been assumed in the Tipping's original RVM formulation, but stay finite with values depending on the particular SNR level. This knowledge enabled us to develop model selection rules based on the discovered statistical laws behind the evidence parameters.

In order to be able to apply these rules in practice, we also proposed a modified learning algorithm that exploits the principle of successive interference cancellation. This




FIGURE 13: Multipath detection results for quantile-based method with $\rho = 1 - 10^{-6}$.

modification not only allows to avoid computationally intensive matrix inversions, but also removes the interference between the neighboring basis functions in the design matrix.

Model mismatch case was also considered in our analysis. We were able to assess the possible influence of the finite algorithm resolution and, to some extent, take it into account by adjusting the corresponding model selection rules.

We also showed the relationship between the EP and the classical model selection based on the MDL criterion. It was found that the maximum of the evidence corresponds to the minimum of the corresponding description length criterion. Thus, EP can be used as the classical MDL-like model selection scheme, but also allows faster and more efficient threshold-based implementation.

The EP framework was also compared with the multipath estimation using the SAGE algorithm augmented with the MDL criterion.

According to the simulation results, the description length-based methods, that is, negative log-evidence and SAGE + MDL method, give better results in terms of the achieved probabilities of correct path extraction. They also improve faster as the sampling rate grows. However, these model selection strategies require learning multiple models in parallel, which, of course, imposes additional computational load. The threshold-based method, on the other hand, allows to perform model selection on-line, thus being more efficient, but its performance increase with the growing sampling rate is more modest. The performance of the threshold-based method also depends on the value of the quantile ρ . In our simulations we set $\rho = 1 - 10^{-6}$, which results in the same probability of the path detection as in the SAGE + MDL algorithm. However, other values of ρ can be used, thus giving a way to further optimize the performance of the threshold-based method.

The comparison between the SAGE and EP schemes clearly shows that estimating evidence parameters really pays off. Introducing them in the computation of the model complexity, as it is done in the negative log-evidence approach, results in the best performance, compared to the other two methods. Although the negative log-evidence method needs a slightly higher SNR to reliably detect channels, it however results in the highest probability of the path extraction.

To summarize, we think that the EP is a very promising method that can be superior to the standard model selection algorithms like MDL, both in accuracy and in computational efficiency. It also offers a number of possibilities: the evidence parameters can also be estimated within the SAGE framework, thus extending the list of multipath parameters and enabling on-line model selection within the SAGE algorithm. As a consequence, this would allow to adapt the design matrix by estimating the delays τ_l from the data. The threshold-based method also opens perspectives for on-line remodeling, that is, adding or removing components during the estimation of the model parameters, which might result in much better and sparser models. Since the evidence parameters reflect the contribution of the multipath components, they might also be useful in applications, where it is necessary

to define some measure of confidence for a multipath component.

APPENDIX

EVIDENCE UPDATE EXPRESSIONS

To derive the update expressions for the evidence parameters in the multiple channels case, we first rewrite (22) using the definitions (25). Since both terms under the integral are Gaussian densities, the result can be easily evaluated as

$$\begin{aligned} p(\tilde{\mathbf{z}} | \boldsymbol{\alpha}, \beta) &= \int p(\tilde{\mathbf{z}} | \tilde{\boldsymbol{\omega}}, \beta) p(\tilde{\boldsymbol{\omega}} | \boldsymbol{\alpha}) d\tilde{\boldsymbol{\omega}} \\ &= \frac{\exp(-\tilde{\mathbf{z}}^H (\beta^{-1} \tilde{\mathbf{\Lambda}} + \tilde{\mathbf{K}} \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{K}}^H)^{-1} \tilde{\mathbf{z}})}{\pi^{PN} |\beta^{-1} \tilde{\mathbf{\Lambda}} + \tilde{\mathbf{K}} \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{K}}^H|}. \end{aligned} \quad (\text{A.1})$$

For the sake of completeness we also consider hypermodel priors $p(\boldsymbol{\alpha}, \beta)$ in the derivation of the hyperparameter update expressions. Thus, our goal is to find the values of $\boldsymbol{\alpha}$ and β that maximize $\mathcal{L}(\boldsymbol{\alpha}, \beta | \tilde{\mathbf{z}}) = \log(p(\tilde{\mathbf{z}} | \boldsymbol{\alpha}, \beta) p(\boldsymbol{\alpha}, \beta))$. This is achieved by taking the partial derivatives of $\mathcal{L}(\boldsymbol{\alpha}, \beta | \tilde{\mathbf{z}})$ with respect to $\boldsymbol{\alpha}$ and β , and equating them to zero [19]. It is convenient to maximize $\mathcal{L}(\boldsymbol{\alpha}, \beta | \tilde{\mathbf{z}})$ with respect to $\log(\alpha_l)$ and $\log(\beta)$ since the derivatives of the prior terms in the logarithmic domain are simpler.

First we prove the following matrix identity that we will exploit later

$$|\mathbf{B}^{-1}| |\mathbf{A}^{-1}| |\mathbf{A} + \mathbf{K}^H \mathbf{B} \mathbf{K}| = |\mathbf{B}^{-1} + \mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H|. \quad (\text{A.2})$$

Proof.

$$\begin{aligned} &|\mathbf{B}^{-1}| |\mathbf{A}^{-1}| |\mathbf{A} + \mathbf{K}^H \mathbf{B} \mathbf{K}| \\ &= |\mathbf{B}^{-1}| |\mathbf{A}^{-1}| |\mathbf{K}^H [(\mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H)^{-1} + \mathbf{B}] \mathbf{K}| \\ &= |\mathbf{B}^{-1}| |\mathbf{A}^{-1}| |\mathbf{K}| |(\mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H)^{-1} + \mathbf{B}| |\mathbf{K}^H| \\ &= |\mathbf{K}| |\mathbf{A}^{-1}| |\mathbf{K}^H| |[(\mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H)^{-1} + \mathbf{B}] \mathbf{B}^{-1}| \\ &= |\mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H [(\mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H)^{-1} \mathbf{B}^{-1} + \mathbf{I}]| \\ &= |\mathbf{B}^{-1} + \mathbf{K} \mathbf{A}^{-1} \mathbf{K}^H|. \end{aligned} \quad (\text{A.3})$$

□

Now, we can begin with the derivation of the update of the hyperparameters α_l . Let us define $\tilde{\mathbf{B}}^{-1} = \beta^{-1} \tilde{\mathbf{\Lambda}}$. According to (A.2) we see that

$$\begin{aligned} &|\tilde{\mathbf{B}}^{-1} + \tilde{\mathbf{K}} \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{K}}^H| \\ &= |\tilde{\mathbf{B}}^{-1}| |\tilde{\mathbf{\Lambda}}^{-1}| |\tilde{\mathbf{A}} + \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{K}}| = |\tilde{\mathbf{B}}^{-1}| |\tilde{\mathbf{\Lambda}}^{-1}| |\tilde{\boldsymbol{\Phi}}^{-1}|. \end{aligned} \quad (\text{A.4})$$

Making use of this result, we can write

$$\begin{aligned}
& \frac{\partial \mathcal{L}(\boldsymbol{\alpha}, \beta \mid \tilde{\mathbf{z}})}{\partial \log(\alpha_l)} \\
&= \frac{\partial}{\partial \log \alpha_l} \left\{ -\log |\tilde{\mathbf{B}}^{-1}| |\tilde{\mathbf{A}}^{-1}| |\tilde{\Phi}^{-1}| \right. \\
&\quad \left. - \tilde{\mathbf{z}}^H (\tilde{\mathbf{B}}^{-1} + \tilde{\mathbf{K}} \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{K}}^H)^{-1} \tilde{\mathbf{z}} \right. \\
&\quad \left. + \sum_{l=1}^L (\epsilon \log \alpha_l - \zeta \alpha_l) \right\} \quad (\text{A.5}) \\
&= \frac{\partial \log |\mathbf{A}|^P}{\partial \log \alpha_l} + \sum_{p=1}^P \frac{\partial \log |\Phi_p|}{\partial \log \alpha_l} + (\epsilon - \zeta \alpha_l) \\
&\quad - \tilde{\mathbf{z}}^H \frac{\partial (\tilde{\mathbf{B}} - \tilde{\mathbf{B}} \tilde{\mathbf{K}} (\tilde{\mathbf{A}} + \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{K}})^{-1} \tilde{\mathbf{K}}^H \tilde{\mathbf{B}})}{\partial \log \alpha_l} \tilde{\mathbf{z}},
\end{aligned}$$

where in the latter expression the Woodbury inversion identity [28] was used to expand the term $(\tilde{\mathbf{B}}^{-1} + \tilde{\mathbf{K}} \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{K}}^H)^{-1}$. After taking the derivative we arrive at

$$\begin{aligned}
& \frac{\partial \mathcal{L}(\boldsymbol{\alpha}, \beta \mid \tilde{\mathbf{z}})}{\partial \log(\alpha_l)} = P \operatorname{tr} \left[\mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial \log \alpha_l} \right] + \sum_{p=1}^P \operatorname{tr} \left[\Phi_p^{-1} \frac{\partial \Phi_p}{\partial \log \alpha_l} \right] \\
&\quad + (\epsilon - \zeta \alpha_l) - \tilde{\mathbf{z}}^H \tilde{\mathbf{B}} \tilde{\mathbf{K}} \tilde{\Phi} \frac{\partial (\tilde{\mathbf{A}} + \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{K}})}{\partial \log \alpha_l} \tilde{\Phi} \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{z}} \\
&= P - \sum_{p=1}^P \operatorname{tr} [\alpha_l \mathbf{E}_{ll} \Phi_p] + (\epsilon - \zeta \alpha_l) \\
&\quad - \tilde{\mathbf{z}}^H \tilde{\mathbf{B}} \tilde{\mathbf{K}} \tilde{\Phi} \alpha_l \tilde{\mathbf{E}}_{ll} \tilde{\Phi} \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{z}}. \quad (\text{A.6})
\end{aligned}$$

Here \mathbf{E}_{ll} is a matrix with the l th element on the main diagonal equal to 1, and all other elements being zero. Similarly, $\tilde{\mathbf{E}}_{ll}$ is the P -times repetition of \mathbf{E}_{ll} on its main diagonal. By noting that $\tilde{\boldsymbol{\mu}} = \tilde{\Phi} \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{z}}$, we arrive at

$$\begin{aligned}
& \frac{\partial \mathcal{L}(\boldsymbol{\alpha}, \beta \mid \tilde{\mathbf{z}})}{\partial \log(\alpha_l)} = P - \sum_{p=1}^P \operatorname{tr} [\alpha_l \mathbf{E}_{ll} \Phi_p] \\
&\quad + (\epsilon - \zeta \alpha_l) - \tilde{\boldsymbol{\mu}}^H \alpha_l \tilde{\mathbf{E}}_{ll} \tilde{\boldsymbol{\mu}} = 0. \quad (\text{A.7})
\end{aligned}$$

Solving for α_l , we obtain the final expression for the hyperparameter update

$$\alpha_l = \frac{P + \epsilon}{\sum_{p=1}^P (\Phi_{p,ll} + |\mu_{p,l}|^2) + \zeta}. \quad (\text{A.8})$$

Note that by setting $\zeta = \epsilon = 0$ we effectively remove the influence of the prior $p(\boldsymbol{\alpha} \mid \zeta, \epsilon)$.

We proceed similarly to calculate the update of β

$$\begin{aligned}
& \frac{\partial \mathcal{L}(\boldsymbol{\alpha}, \beta \mid \tilde{\mathbf{z}})}{\partial \log(\beta)} = \sum_{p=1}^P \frac{\partial \log |\mathbf{B}_p|}{\partial \log \beta} + \sum_{p=1}^P \frac{\partial \log |\Phi_p|}{\partial \log \beta} + (v - \kappa \beta) \\
&\quad - \tilde{\mathbf{z}}^H \frac{\partial (\tilde{\mathbf{B}} - \tilde{\mathbf{B}} \tilde{\mathbf{K}} (\tilde{\mathbf{A}} + \tilde{\mathbf{K}}^H \tilde{\mathbf{B}} \tilde{\mathbf{K}})^{-1} \tilde{\mathbf{K}}^H \tilde{\mathbf{B}})}{\partial \log \beta} \tilde{\mathbf{z}} \\
&= \sum_{p=1}^P \frac{\partial \log \beta^N |\Lambda_p^{-1}|}{\partial \log \beta} + \sum_{p=1}^P \operatorname{tr} \left[\Phi_p^{-1} \frac{\partial \Phi_p}{\partial \log \beta} \right] \\
&\quad + (v - \kappa \beta) - \tilde{\mathbf{z}}^H \frac{\partial \beta \tilde{\Lambda}^{-1}}{\partial \log \beta} \tilde{\mathbf{z}} \\
&\quad + \tilde{\mathbf{z}}^H \frac{\partial (\beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}} (\tilde{\mathbf{A}} + \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}})^{-1} \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1})}{\partial \log \beta} \tilde{\mathbf{z}} \\
&= PN - \sum_{p=1}^P \operatorname{tr} \left[\Phi_p^{-1} \Phi_p \frac{\partial (\mathbf{A} + \mathbf{K}_p^H \beta \Lambda_p^{-1} \mathbf{K}_p)}{\partial \log \beta} \Phi_p \right] \\
&\quad + (v - \kappa \beta) - \tilde{\mathbf{z}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{z}} \\
&\quad + \tilde{\mathbf{z}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}} \tilde{\Phi} \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{z}} \\
&\quad + \tilde{\mathbf{z}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}} \frac{\partial (\tilde{\mathbf{A}} + \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}})}{\partial \log \beta} \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{z}} \\
&\quad + \tilde{\mathbf{z}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}} \tilde{\Phi} \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{z}} \\
&= PN - \sum_{p=1}^P \operatorname{tr} [\mathbf{K}_p^H \beta \Lambda_p^{-1} \mathbf{K}_p \Phi_p] \\
&\quad + (v - \kappa \beta) - \tilde{\mathbf{z}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{z}} + \tilde{\mathbf{z}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}} \tilde{\boldsymbol{\mu}} \\
&\quad + \tilde{\boldsymbol{\mu}}^H \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{K}} \tilde{\boldsymbol{\mu}} + \tilde{\boldsymbol{\mu}}^H \tilde{\mathbf{K}}^H \beta \tilde{\Lambda}^{-1} \tilde{\mathbf{z}}. \quad (\text{A.9})
\end{aligned}$$

Thus we arrive at the final expression:

$$\begin{aligned}
& \frac{\partial \mathcal{L}(\boldsymbol{\alpha}, \beta \mid \tilde{\mathbf{z}})}{\partial \log(\beta)} = PN - \sum_{p=1}^P \operatorname{tr} [\mathbf{K}_p^H \beta \Lambda_p^{-1} \mathbf{K}_p \Phi_p] \\
&\quad + (v - \kappa \beta) \\
&\quad - \sum_{p=1}^P (\mathbf{z}_p - \mathbf{K}_p \boldsymbol{\mu}_p)^H \beta \Lambda_p^{-1} (\mathbf{z}_p - \mathbf{K}_p \boldsymbol{\mu}_p) = 0. \quad (\text{A.10})
\end{aligned}$$

Solving for β we finally obtain

$$\begin{aligned}
& \beta = (PN + v) \left(\sum_{p=1}^P \operatorname{tr} [\mathbf{K}_p^H \Lambda_p^{-1} \mathbf{K}_p \Phi_p] \right. \\
&\quad \left. + \sum_{p=1}^P (\mathbf{z}_p - \mathbf{K}_p \boldsymbol{\mu}_p)^H \Lambda_p^{-1} (\mathbf{z}_p - \mathbf{K}_p \boldsymbol{\mu}_p) + \kappa \right)^{-1}. \quad (\text{A.11})
\end{aligned}$$

Here again the choice $\kappa = \nu = 0$ removes the influence of the prior $p(\beta \mid \kappa, \nu)$ on the evidence maximization.

REFERENCES

- [1] H. Krim and M. Viberg, "Two decades of array signal processing research: the parametric approach," *IEEE Signal Processing Magazine*, vol. 13, no. 4, pp. 67–94, 1996.
- [2] B. H. Fleury, M. Tschudin, R. Heddergott, D. Dahlhaus, and K. I. Pedersen, "Channel parameter estimation in mobile radio environments using the SAGE algorithm," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 3, pp. 434–450, 1999.
- [3] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, John Wiley & Sons, New York, NY, USA, 2nd edition, 2000.
- [4] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 2, pp. 387–392, 1985.
- [5] J. Rissanen, "Modelling by the shortest data description," *Automatica*, vol. 14, no. 5, pp. 465–471, 1978.
- [6] S. Haykin, Ed., *Kalman Filtering and Neural Networks*, John Wiley & Sons, New York, NY, USA, 2001.
- [7] M. Feder and E. Weinstein, "Parameter estimation of superimposed signals using the EM algorithm," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 36, no. 4, pp. 477–489, 1988.
- [8] D. J. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, Cambridge, UK, 2003.
- [9] W. J. Fitzgerald, "The Bayesian approach to signal modelling," in *IEE Colloquium on Non-Linear Signal and Image Processing (Ref. No. 1998/284)*, pp. 9/1–9/5, London, UK, May 1998.
- [10] G. Schwarz, "Estimating the dimension of a model," *Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.
- [11] J. J. Rissanen, "Fisher information and stochastic complexity," *IEEE Transactions on Information Theory*, vol. 42, no. 1, pp. 40–47, 1996.
- [12] A. D. Lanterman, "Schwarz, Wallace, and Rissanen: intertwining themes in theories of model selection," *International Statistical Review*, vol. 69, no. 2, pp. 185–212, 2001.
- [13] D. J. C. MacKay, "Bayesian interpolation," *Neural Computation*, vol. 4, no. 3, pp. 415–447, 1992.
- [14] D. J. C. MacKay, "Bayesian methods for backpropagation networks," in *Models of Neural Networks III*, E. Domany, J. L. van Hemmen, and K. Schulten, Eds., chapter 6, pp. 211–254, Springer, New York, NY, USA, 1994.
- [15] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, vol. 1, no. 3, pp. 211–244, 2001.
- [16] T. S. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall PTR, Saddle River, NJ, USA, 2002.
- [17] D. Heckerman, "A tutorial on learning with Bayesian networks," Tech. Rep. MSR-TR-95-06, Microsoft Research, Advanced Technology Division, One Microsoft Way, Redmond, Wash, USA, March 1995.
- [18] R. Neal, *Bayesian Learning for Neural Networks*, vol. 118 of *Lecture Notes in Statistics*, Springer, New York, NY, USA, 1996.
- [19] O. Berger, *Statistical Decision Theory and Bayesian Analysis*, Springer, New York, NY, USA, 2nd edition, 1985.
- [20] P. Grünwald, "A tutorial introduction to the minimum description length principle," in *Advances in Minimum Description Length: Theory and Applications*, P. Grünwald, I. Myung, and M. Pitt, Eds., pp. 80 pages, MIT Press, Cambridge, Mass, USA, 2005.
- [21] A. Barron, J. Rissanen, and B. Yu, "The minimum description length principle in coding and modeling," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2743–2760, 1998.
- [22] A. C. Faul and M. E. Tipping, "Analysis of sparse Bayesian learning," in *Advances in Neural Information Processing Systems (NIPS '01)*, T. G. Dietterich, S. Becker, and Z. Ghahramani, Eds., vol. 14, pp. 383–389, MIT Press, Vancouver, British Columbia, Canada, December 2002.
- [23] K. Conradsen, A. Nielsen, J. Schou, and H. Skriver, "A test statistic in the complex Wishart distribution and its application to change detection in polarimetric SAR data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 1, pp. 4–19, 2003.
- [24] N. R. Goodman, "Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction)," *The Annals of Mathematical Statistics*, vol. 34, no. 1, pp. 152–177, 1963.
- [25] M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions*, John Wiley & Sons, New York, NY, USA, 3rd edition, 2000.
- [26] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Dover, New York, NY, USA, 1972.
- [27] T. I. Laakso, V. Välimäki, M. Karjalainen, and U. K. Laine, "Splitting the unit delay [FIR/all pass filters design]," *IEEE Signal Processing Magazine*, vol. 13, no. 1, pp. 30–60, 1996.
- [28] G. H. Golub and C. F. van Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, Md, USA, 1996.

Dmitriy Shutin received his M.S. degree in computer science in 2000 from the Dnepropetrovsk State University, Ukraine. From 1998 to 2000, he was with the Faculty of Radiophysics, where he actively conducted research in signal and images processing with application to biomedical systems. In 2006, he received Dr. techn degree in electrical engineering from Graz University of Technology where he is currently a Teaching Assistant in the Signal Processing and Speech Communication Laboratory. His research interests are in nonlinear signal processing, machine learning, statistical pattern recognition, and adaptive systems.



Gernot Kubin was born in Vienna, Austria, on June 24, 1960. He received his Dipl.-Ing. and Dr. techn. (sub auspiciis praesidentis) degrees in electrical engineering from Vienna University of Technology, Vienna, Austria, in 1982 and 1990, respectively. Since 2000, he has been a Professor of Nonlinear Signal Processing and head of the Signal Processing and Speech Communication Laboratory (SPSC) at Graz University of Technology, Graz, Austria. Earlier international appointments include: CERN Geneva, Switzerland, in 1980, Vienna University of Technology from 1983 to 2000, Erwin Schroedinger Fellow at Philips Natuurkundig Laboratorium Eindhoven, The Netherlands, in 1985, AT&T Bell Labs Murray Hill, NJ, from 1992 to 1993, and in 1995, KTH Stockholm, Sweden, in 1998, and Global IP Sound, Sweden and USA, in 2000 and 2001. He is engaged in several national research centres for academia-industry collaboration such as the Vienna Telecommunications Research Centre FTW,



1999-present (Key Researcher and Member of the Board), the Christian Doppler Laboratory for Nonlinear Signal Processing 2002-present (Founding Director, main partner Infineon Technologies), and the Competence Network for Advanced Speech Technology COAST 2006-present (Scientific Director, main partner Philips Speech Recognition Systems). Dr. Kubin is a Member of the Board of Austrian Acoustics Association. His research interests are in nonlinear signals and systems, digital communications, computational intelligence, and speech communication. He has authored or co-authored over one hundred peer-reviewed publications and four patents.

Bernard H. Fleury received the Diploma degree in electrical engineering and mathematics in 1978 and 1990, respectively, and the doctoral degree in electrical engineering in 1990 from the Swiss Federal Institute of Technology Zurich (ETHZ), Switzerland. Since 1997, Bernard H. Fleury has been with the Department of Communication Technology, Aalborg University, Denmark, where he is Professor in Digital Communications. He has also been affiliated with the Telecommunication Research Center, Vienna (ftw.) since April 2006. Bernard H. Fleury is presently Chairman of Department 2 “Radio Channel Modelling for Design Optimisation and Performance Assessment of Next Generation Communication Systems” of the on-going FP6 network of excellence NEWCOM (Network of Excellence in Communications). During 1978–1985 and 1988–1992, he was Teaching Assistant and Research Assistant, respectively, at the Communication Technology Laboratory and at the Statistical Seminar at ETHZ. In 1992, he joined again the former laboratory as Senior Research Associate. In 1999, he was elected IEEE Senior Member. Bernard H. Fleury’s general fields of interest cover numerous aspects within Communication Theory and Signal Processing mainly for Wireless Communications. His current areas of research include stochastic modelling and estimation of the radio channel, characterization of multiple-input multiple-output (MIMO) channels, and iterative processing algorithms for joint channel estimation and data detection/decoding in multiuser communication systems.

