

A PROBABILISTIC LEAST-MEAN-SQUARES FILTER

Jesús Fernández-Bes, Víctor Elvira

Department of Signal Theory and Communications
Universidad Carlos III de Madrid
Leganés (Spain)

Steven Van Vaerenbergh

Department of Communications Engineering
University of Cantabria
Santander (Spain)

ABSTRACT

We introduce a probabilistic approach to the LMS filter. By means of an efficient approximation, this approach provides an adaptable step-size LMS algorithm together with a measure of uncertainty about the estimation. In addition, the proposed approximation preserves the linear complexity of the standard LMS. Numerical results show the improved performance of the algorithm with respect to standard LMS and state-of-the-art algorithms with similar complexity. The goal of this work, therefore, is to open the door to bring some more Bayesian machine learning techniques to adaptive filtering.

Index Terms— probabilistic models, least-mean-squares, adaptive filtering, state-space models

1. INTRODUCTION

Probabilistic models have proven to be very useful in a lot of applications in signal processing where signal estimation is needed [1, 2, 3]. Some of their advantages are that 1) they force the designer to specify all the assumptions of the model, 2) they provide a clear separation between the model and the algorithm used to solve it, and 3) they usually provide some measure of uncertainty about the estimation.

On the other hand, adaptive filtering is a standard approach in estimation problems when the input is received as a stream of data that is potentially non-stationary. This approach is widely understood and applied to several problems such as echo cancellation [4], noise cancellation [5], and channel equalization [6].

Although these two approaches share some underlying relations, there are very few connections in the literature. The first important attempt in the signal processing community to relate these two fields was the connection between a linear

Gaussian state-space model (i.e. Kalman filter) and the RLS filter, by Sayed and Kailath [7] and then by Haykin *et al.* [8]. The RLS adaptive filtering algorithm emerges naturally when one defines a particular state-space model (SSM) and then performs exact inference in that model. This approach was later exploited in [9] to design a kernel RLS algorithm based on Gaussian processes.

A first attempt to approximate the LMS filter from a probabilistic perspective was presented in [10], focusing on a kernel-based implementation. The algorithm of [10] makes use of a Maximum a Posteriori (MAP) estimate as an approximation for the predictive step. However, this approximation does not preserve the estimate of the uncertainty in each step, therefore degrading the performance of the algorithm.

In this work, we provide a similar connection between state-space models and least-mean-squares (LMS). Our approach is based on approximating the posterior distribution with an isotropic Gaussian distribution. We show how the computation of this approximated posterior leads to a linear-complexity algorithm, comparable to the standard LMS. Similar approaches have already been developed for a variety of problems such as channel equalization using recurrent RBF neural networks [11], or Bayesian forecasting [12]. Here, we show the usefulness of this probabilistic approach for adaptive filtering.

The probabilistic perspective we adopt throughout this work presents two main advantages. Firstly, a novel LMS algorithm with adaptable step size emerges naturally with this approach, making it suitable for both stationary and non-stationary environments. The proposed algorithm has less free parameters than previous LMS algorithms with variable step size [13, 14, 15], and its parameters are easier to be tuned w.r.t. these algorithms and standard LMS. Secondly, the use of a probabilistic model provides us with an estimate of the error variance, which is useful in many applications.

Experiments with simulated and real data show the advantages of the presented approach with respect to previous works. However, we remark that the main contribution of this paper is that it opens the door to introduce more Bayesian machine learning techniques, such as variational inference and Monte Carlo sampling methods [16], to adaptive filtering.

The work of Jesús Fernández-Bes was supported by the Spanish MEC FPU program and MICINN project PRI-PIBIN-2011-1266. The work of Víctor Elvira was supported by the Spanish Government (MINECO) under project OTOSiS (TEC2013-41718-R), the Community of Madrid Government under project AGES (S2010/BMD-2422), and the European Union's 7th FP through the Marie Curie ITN MLPM2012 (Grant No. 316861).

The work of Steven Van Vaerenbergh was supported by the Spanish Government (MINECO) under project RACHEL (TEC2013-47141-C4-3-R).

2. PROBABILISTIC MODEL

Throughout this work, we assume the observation model to be linear-Gaussian with the following distribution,

$$p(y_k | \mathbf{w}_k) = \mathcal{N}(y_k; \mathbf{x}_k^T \mathbf{w}_k, \sigma_n^2), \quad (1)$$

where σ_n^2 is the variance of the observation noise, \mathbf{x}_k is the regression vector and \mathbf{w}_k is the parameter vector to be sequentially estimated, both M -dimensional column vectors.

In a non-stationary scenario, \mathbf{w}_k follows a dynamic process. In particular, we consider a diffusion process (random-walk model) with variance σ_d^2 for this parameter vector:

$$p(\mathbf{w}_k | \mathbf{w}_{k-1}) = \mathcal{N}(\mathbf{w}_k; \mathbf{w}_{k-1}, \sigma_d^2 \mathbf{I}), \quad (2)$$

where \mathbf{I} denotes the identity matrix. In order to initiate the recursion, we assume the following prior distribution on \mathbf{w}_k

$$p(\mathbf{w}_0) = \mathcal{N}(\mathbf{w}_0; 0, \sigma_d^2 \mathbf{I}).$$

3. EXACT INFERENCE IN THIS MODEL: REVISITING THE RLS FILTER

Given the described probabilistic SSM, we would like to infer the posterior probability distribution $p(\mathbf{w}_k | y_{1:k})$. Since all involved distributions are Gaussian, one can perform exact inference, leveraging the probability rules in a straightforward manner. The resulting probability distribution is

$$p(\mathbf{w}_k | y_{1:k}) = \mathcal{N}(\mathbf{w}_k; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k),$$

in which the mean vector $\boldsymbol{\mu}_k$ is given by

$$\boldsymbol{\mu}_k = \boldsymbol{\mu}_{k-1} + \mathbf{K}_k (y_k - \mathbf{x}_k^T \boldsymbol{\mu}_{k-1}) \mathbf{x}_k,$$

where we have introduced the auxiliary variable

$$\mathbf{K}_k = \frac{(\boldsymbol{\Sigma}_{k-1} + \sigma_d^2 \mathbf{I})}{\mathbf{x}_k^T (\boldsymbol{\Sigma}_{k-1} + \sigma_d^2 \mathbf{I}) \mathbf{x}_k + \sigma_n^2},$$

and the covariance matrix $\boldsymbol{\Sigma}_k$ is obtained as

$$\boldsymbol{\Sigma}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{x}_k \mathbf{x}_k^T) (\boldsymbol{\Sigma}_{k-1} + \sigma_d^2 \mathbf{I}),$$

Note that the mode of $p(\mathbf{w}_k | y_{1:k})$, i.e. the maximum-a-posteriori estimate (MAP), coincides with the RLS adaptive rule

$$\mathbf{w}_k^{(RLS)} = \mathbf{w}_{k-1}^{(RLS)} + \mathbf{K}_k (y_k - \mathbf{x}_k^T \mathbf{w}_{k-1}^{(RLS)}) \mathbf{x}_k. \quad (3)$$

This rule is similar to the one introduced in [8].

Finally, note that the covariance matrix $\boldsymbol{\Sigma}_k$ is a measure of the uncertainty of the estimate \mathbf{w}_k conditioned on the observed data $y_{1:k}$. Nevertheless, for many applications a single

scalar summarizing the variance of the estimate could prove to be sufficiently useful. In the next section, we show how such a scalar is obtained naturally when $p(\mathbf{w}_k | y_{1:k})$ is approximated with an isotropic Gaussian distribution. We also show that this approximation leads to an LMS-like estimation.

4. APPROXIMATING THE POSTERIOR DISTRIBUTION: LMS FILTER

The proposed approach consists in approximating the posterior distribution $p(\mathbf{w}_k | y_{1:k})$, in general a multivariate Gaussian distribution with a full covariance matrix, by an isotropic spherical Gaussian distribution

$$\hat{p}(\mathbf{w}_k | y_{1:k}) = \mathcal{N}(\mathbf{w}_k; \hat{\boldsymbol{\mu}}_k, \hat{\sigma}_k^2 \mathbf{I}). \quad (4)$$

In order to estimate the mean and covariance of the approximate distribution $\hat{p}(\mathbf{w}_k | y_{1:k})$, we propose to select those that minimize the Kullback-Leibler divergence with respect to the original distribution, i.e.,

$$\{\hat{\boldsymbol{\mu}}_k, \hat{\sigma}_k\} = \arg \min_{\boldsymbol{\mu}_k, \sigma_k} \{D_{KL}(p(\mathbf{w}_k | y_{1:k}) || \hat{p}(\mathbf{w}_k | y_{1:k}))\}.$$

The derivation of the corresponding minimization problem can be found in Appendix A. In particular, the optimal mean and the covariance are found as

$$\hat{\boldsymbol{\mu}}_k = \boldsymbol{\mu}_k; \quad \hat{\sigma}_k^2 = \frac{\text{Tr}\{\boldsymbol{\Sigma}_k\}}{M}. \quad (5)$$

We now show that by using (4) in the recursive predictive and filtering expressions we obtain an LMS-like adaptive rule. First, let us assume that we have an approximate posterior distribution at $k-1$, $\hat{p}(\mathbf{w}_{k-1} | y_{1:k-1}) = \mathcal{N}(\mathbf{w}_{k-1}; \hat{\boldsymbol{\mu}}_{k-1}, \hat{\sigma}_{k-1}^2 \mathbf{I})$. Since all involved distributions are Gaussian, the predictive distribution is obtained as

$$\begin{aligned} \hat{p}(\mathbf{w}_k | y_{1:k-1}) &= \int p(\mathbf{w}_k | \mathbf{w}_{k-1}) \hat{p}(\mathbf{w}_{k-1} | y_{1:k-1}) d\mathbf{w}_{k-1} \\ &= \mathcal{N}(\mathbf{w}_k; \boldsymbol{\mu}_k |_{k-1}, \boldsymbol{\Sigma}_k |_{k-1}), \end{aligned} \quad (6)$$

where the mean vector and covariance matrix are given by

$$\begin{aligned} \hat{\boldsymbol{\mu}}_{k|k-1} &= \hat{\boldsymbol{\mu}}_{k-1} \\ \hat{\boldsymbol{\Sigma}}_{k|k-1} &= (\hat{\sigma}_{k-1}^2 + \sigma_d^2) \mathbf{I}. \end{aligned}$$

From (6), the posterior distribution at time k can be computed using Bayes' Theorem and standard Gaussian manipulations (see for instance [17, Ch. 4]). Then, we approximate the posterior $p(\mathbf{w}_k | y_{1:k})$ with an isotropic Gaussian,

$$\hat{p}(\mathbf{w}_k | y_{1:k}) = \mathcal{N}(\mathbf{w}_k; \hat{\boldsymbol{\mu}}_k, \hat{\sigma}_k^2 \mathbf{I}),$$

where

$$\begin{aligned} \hat{\boldsymbol{\mu}}_k &= \hat{\boldsymbol{\mu}}_{k-1} + \frac{(\hat{\sigma}_{k-1}^2 + \sigma_d^2)}{(\hat{\sigma}_{k-1}^2 + \sigma_d^2) \|\mathbf{x}_k\|^2 + \sigma_n^2} (y_k - \mathbf{x}_k^T \hat{\boldsymbol{\mu}}_{k-1}) \mathbf{x}_k \\ &= \hat{\boldsymbol{\mu}}_{k-1} + \eta_k (y_k - \mathbf{x}_k^T \hat{\boldsymbol{\mu}}_{k-1}) \mathbf{x}_k. \end{aligned} \quad (7)$$

Note that, instead of a gain matrix \mathbf{K}_k as in Eq. (3), we now have a scalar gain η_k that operates as a variable step size.

Finally, to obtain the posterior variance, which is our measure of uncertainty, we apply (5) and the trick $\text{Tr}\{\mathbf{x}_k \mathbf{x}_k^T\} = \mathbf{x}_k^T \mathbf{x}_k = \|\mathbf{x}_k\|^2$,

$$\hat{\sigma}_k^2 = \frac{\text{Tr}(\boldsymbol{\Sigma}_k)}{M} \quad (8)$$

$$= \frac{1}{M} \text{Tr} \{ (\mathbf{I} - \eta_k \mathbf{x}_k \mathbf{x}_k^T) (\hat{\sigma}_{k-1}^2 + \sigma_d^2) \} \quad (9)$$

$$= \left(1 - \frac{\eta_k \|\mathbf{x}_k\|^2}{M} \right) (\hat{\sigma}_{k-1}^2 + \sigma_d^2). \quad (10)$$

If MAP estimation is performed, we obtain an adaptable step-size LMS estimation

$$\mathbf{w}_k^{(LMS)} = \mathbf{w}_{k-1}^{(LMS)} + \eta_k (y_k - \mathbf{x}_k^T \mathbf{w}_{k-1}^{(LMS)}) \mathbf{x}_k, \quad (11)$$

with

$$\eta_k = \frac{(\hat{\sigma}_{k-1}^2 + \sigma_d^2)}{(\hat{\sigma}_{k-1}^2 + \sigma_d^2) \|\mathbf{x}_k\|^2 + \sigma_n^2}.$$

At this point, several interesting remarks can be made:

- The adaptive rule (11) has linear complexity since it does not require us to compute the full matrix $\boldsymbol{\Sigma}_k$.
- For a stationary model, we have $\sigma_d^2 = 0$ in (7) and (10). In this case, the algorithm remains valid and both the step size and the error variance, $\hat{\sigma}_k$, vanish over time k .
- Finally, the proposed adaptable step-size LMS has only two parameters, σ_d^2 and σ_n^2 , (and only one, σ_n^2 , in stationary scenarios) in contrast to other variable step-size algorithms [13, 14, 15]. More interestingly, both σ_d^2 and σ_n^2 have a clear underlying physical meaning, and they can be estimated in many cases. We will comment more about this in the next section.

5. EXPERIMENTS

We evaluate the performance of the proposed algorithm in both stationary and tracking experiments. In the first experiment, we estimate a fixed vector \mathbf{w}^o of dimension $M = 50$. The entries of the vector are independently and uniformly chosen in the range $[-1, 1]$. Then, the vector is normalized so that $\|\mathbf{w}^o\| = 1$. Regressors \mathbf{x}_k are zero-mean Gaussian vectors with identity covariance matrix. The additive noise variance is such that the SNR is 20 dB. We compare our algorithm with standard RLS and three other LMS-based algorithms: LMS, NLMS [18], VSS-LMS [15].¹ The probabilistic LMS algorithm in [10] is not simulated because it is not suitable for stationary environments.

¹The used parameters for each algorithm are: for RLS $\lambda = 1$, $\epsilon^{-1} = 0.01$; for LMS $\mu = 0.01$; for NLMS $\mu = 0.5$; and for VSS-LMS $\mu_{max} = 1$, $\alpha = 0.95$, $C = 1e - 4$.

In stationary environments, the proposed algorithm has only one parameter, σ_n^2 . We simulate both the scenario where we have perfectly knowledge of the amount of noise (probLMS1) and the case where the value σ_n^2 is 100 times smaller than the actual value (probLMS2). The Mean-Square Deviation (MSD = $\mathbb{E}\|\mathbf{w}_0 - \mathbf{w}_k\|^2$), averaged out over 50 independent simulations, is presented in Fig. 1.

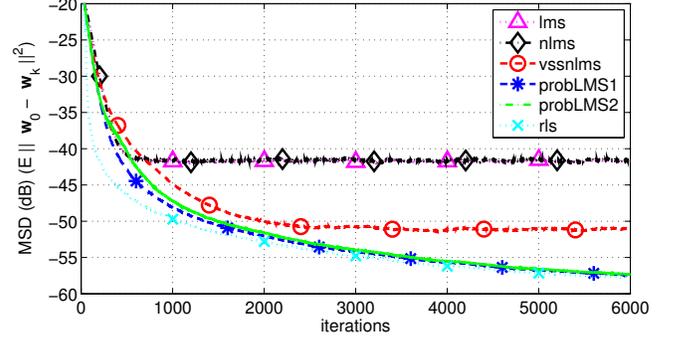


Fig. 1. Performance in terms of MSD of probabilistic LMS with both optimal (probLMS1) and suboptimal (probLMS2) compared to LMS, NLMS, VS-LMS, and RLS.

The performance of probabilistic LMS is close to RLS (obviously at a much lower computational cost) and largely outperforms previous variable step-size LMS algorithms proposed in the literature. Note that, when the model is stationary, i.e. $\sigma_d^2 = 0$ in (2), both the uncertainty $\hat{\sigma}_k^2$, and the adaptive step size η_k , vanish over time. This implies that the error tends to zero when k goes to infinity. Fig. 1 also shows that the proposed approach is not very sensitive to a bad choice of its only parameter, as demonstrated by the good results of probLMS2, which uses a σ_n^2 that is 100 times smaller than the optimal value.

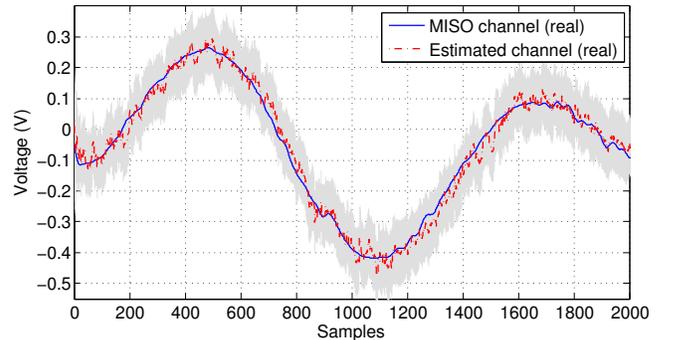


Fig. 2. Real part of one coefficient of the measured and estimated channel in experiment two. The shaded area represents two standard deviations from the prediction (the mean of the posterior distribution).

Method	LMS	NLMS	LMS-2013	VSSNLMS	probLMS	RLS
MSD (dB)	-28.45	-21.07	-14.36	-26.90	-28.36	-25.97

Table 1. Steady-state MSD of the different algorithms for the tracking of a real MISO channel.

In a second experiment, we test the tracking capabilities of the proposed algorithm with real data of a wireless MISO channel acquired in a realistic indoor scenario. More details on the setup can be found in [19]. Fig. 2 shows the real part of one of the channels, and the estimate of the proposed algorithm. The shaded area represents the estimated uncertainty for each prediction, i.e. $\hat{\mu}_k \pm 2\hat{\sigma}_k$. Since the experimental setup does not allow us to obtain the optimal values for the parameters, we fix these parameters to their values that optimize the steady-state mean square deviation (MSD). Table 1 shows this steady-state MSD of the estimate of the MISO channel with different methods. As can be seen, the best tracking performance is obtained by standard LMS and the proposed method.

6. CONCLUSIONS AND OPENED EXTENSIONS

We have presented a probabilistic interpretation of the least-mean-square filter. The resulting algorithm is an adaptable step-size LMS that performs well both in stationary and tracking scenarios. Moreover, it has fewer free parameters than previous approaches and these parameters have a clear physical meaning. Finally, as stated in the introduction, one of the advantages of having a probabilistic model is that it is easily extensible:

- If, instead of using an isotropic Gaussian distribution in the approximation, we used a Gaussian with diagonal covariance matrix, we would obtain a similar algorithm with different step sizes and measures of uncertainty, for each component of \mathbf{w}_k . Although this model can be more descriptive, it needs more parameters to be tuned, and the parallelism with LMS vanishes.
- Similarly, if we substitute the transition model of (2) by an Ornstein-Uhlenbeck process,

$$p(\mathbf{w}_k | \mathbf{w}_{k-1}) = \mathcal{N}(\mathbf{w}_k; \lambda \mathbf{w}_{k-1}, \sigma_d^2),$$

a similar algorithm is obtained but with a forgetting factor λ multiplying $\mathbf{w}_{k-1}^{(LMS)}$ in (11). This algorithm may have improved performance under such a kind of autoregressive dynamics of \mathbf{w}_k , though, again, the connection with standard LMS becomes dimmer.

- As in [10], the measurement model (1) can be changed to obtain similar adaptive algorithms for classification, ordinal regression, and Dirichlet regression for compositional data.

- A similar approximation technique could be applied to more complex dynamical models, i.e. switching dynamical models [20]. The derivation of efficient adaptive algorithms that explicitly take into account a switch in the dynamics of the parameters of interest is a non-trivial and open problem, though the proposed approach could be useful.
- Finally, like standard LMS, this algorithm can be kernelized for its application in estimation under non-linear scenarios.

A. KL DIVERGENCE BETWEEN A GENERAL GAUSSIAN DISTRIBUTION AND AN ISOTROPIC GAUSSIAN

We want to approximate $p_{\mathbf{x}_1}(x) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ by $p_{\mathbf{x}_2}(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_2, \sigma_2^2 \mathbf{I})$. In order to do so, we have to compute the parameters of $p_{\mathbf{x}_2}(\mathbf{x})$, $\boldsymbol{\mu}_2$ and σ_2^2 , that minimize the following Kullback-Leibler divergence,

$$\begin{aligned} D_{KL}(p_{\mathbf{x}_1} \| p_{\mathbf{x}_2}) &= \int_{-\infty}^{\infty} p_{\mathbf{x}_1}(\mathbf{x}) \ln \frac{p_{\mathbf{x}_1}(\mathbf{x})}{p_{\mathbf{x}_2}(\mathbf{x})} d\mathbf{x} \\ &= \frac{1}{2} \{-M + \text{Tr}(\sigma_2^{-2} \mathbf{I} \cdot \boldsymbol{\Sigma}_1^{-1}) \\ &\quad + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \sigma_2^{-2} \mathbf{I} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) \\ &\quad + \ln \frac{\sigma_2^{2M}}{\det \boldsymbol{\Sigma}_1}\}. \end{aligned} \quad (12)$$

Using symmetry arguments, we obtain

$$\boldsymbol{\mu}_2^* = \arg \min_{\boldsymbol{\mu}_2} \{D_{KL}(p_{\mathbf{x}_1} \| p_{\mathbf{x}_2})\} = \boldsymbol{\mu}_1. \quad (13)$$

Then, (12) gets simplified into

$$D_{KL}(p_{\mathbf{x}_1} \| p_{\mathbf{x}_2}) = \frac{1}{2} \{-M + \text{Tr}(\frac{\boldsymbol{\Sigma}_1}{\sigma_2^2}) + \ln \frac{\sigma_2^{2M}}{\det \boldsymbol{\Sigma}_1}\}. \quad (14)$$

The variance σ_2^2 is computed in order to minimize this Kullback-Leibler divergence as

$$\begin{aligned} \sigma_2^{2*} &= \arg \min_{\sigma_2^2} D_{KL}(P_{x_1} \| P_{x_2}) \\ &= \arg \min_{\sigma_2^2} \{\sigma_2^{-2} \text{Tr}\{\boldsymbol{\Sigma}_1\} + M \ln \sigma_2^2\}. \end{aligned} \quad (15)$$

Deriving and making it equal zero leads to

$$\frac{\partial}{\partial \sigma_2^2} \left[\frac{\text{Tr}\{\boldsymbol{\Sigma}_1\}}{\sigma_2^2} + M \ln \sigma_2^2 \right] = \frac{M}{\sigma_2^2} - \frac{\text{Tr}\{\boldsymbol{\Sigma}_1\}}{(\sigma_2^2)^2} \Big|_{\sigma_2^2 = \sigma_2^{2*}} = 0.$$

Finally, since the divergence has a single extremum in R_+ ,

$$\sigma_2^{2*} = \frac{\text{Tr}\{\boldsymbol{\Sigma}_1\}}{M}. \quad (16)$$

7. REFERENCES

- [1] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [2] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for on-line nonlinear/non-gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, 2002.
- [3] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2346–2356, 2008.
- [4] A. Gilloire and M. Vetterli, "Adaptive filtering in subbands with critical sampling: analysis, experiments, and application to acoustic echo cancellation," *IEEE Trans. Signal Process.*, vol. 40, no. 8, pp. 1862–1875, 1992.
- [5] P. A. Nelson and S. J. Elliott, *Active control of sound*, Academic press, 1991.
- [6] D. Falconer, S. L. Ariyavisitakul, A. Benyamin-Seeyar, and B. Eidson, "Frequency domain equalization for single-carrier broadband wireless systems," *IEEE Commun. Mag.*, vol. 40, no. 4, pp. 58–66, 2002.
- [7] A. H. Sayed and T. Kailath, "A state-space approach to adaptive RLS filtering," *IEEE Signal Process. Mag.*, vol. 11, no. 3, pp. 18–60, 1994.
- [8] S. Haykin, A. H. Sayed, J. R. Zeidler, P. Yee, and P. C. Wei, "Adaptive tracking of linear time-variant systems by extended RLS algorithms," *IEEE Trans. Signal Process.*, vol. 45, no. 5, pp. 1118–1128, 1997.
- [9] S. Van Vaerenbergh, M. Lázaro-Gredilla, and I. Santamaría, "Kernel recursive least-squares tracker for time-varying regression," *IEEE Trans. Neural Netw. and Learn. Syst.*, vol. 23, no. 8, pp. 1313–1326, 2012.
- [10] I. M. Park, S. Seth, and S. Van Vaerenbergh, "Probabilistic kernel least mean squares algorithms," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Florence (Italy), May 2014.
- [11] J. Cid-Sueiro, A. Artés-Rodríguez, and A. R. Figueiras-Vidal, "Recurrent radial basis function networks for optimal symbol-by-symbol equalization," *Signal Process.*, vol. 40, no. 1, pp. 53–63, 1994.
- [12] J. Harrison, *Bayesian Forecasting & Dynamic Models*, Springer, 1999.
- [13] R. H. Kwong and E. W. Johnston, "A variable step size LMS algorithm," *IEEE Trans. Signal Process.*, vol. 40, no. 7, pp. 1633–1642, 1992.
- [14] T. Aboulnasr and K. Mayyas, "A robust variable step-size LMS-type algorithm: analysis and simulations," *IEEE Trans. Signal Process.*, vol. 45, no. 3, pp. 631–639, 1997.
- [15] H.-C. Shin, A. H. Sayed, and W.-J. Song, "Variable step-size NLMS and affine projection algorithms," *IEEE Signal processing Letters*, vol. 11, no. 2, pp. 132–135, 2004.
- [16] D. Barber, *Bayesian reasoning and machine learning*, Cambridge University Press, 2012.
- [17] K. P. Murphy, *Machine learning: a probabilistic perspective*, MIT press, 2012.
- [18] A. H. Sayed, *Adaptive filters*, John Wiley & Sons, 2008.
- [19] J. Gutiérrez, Ó. González, J. Pérez, D. Ramírez, L. Vielva, J. Ibáñez, and I. Santamaría, "Frequency-domain methodology for measuring MIMO channels using a generic test bed," *IEEE Trans. on Instrum. Meas.*, vol. 60, no. 3, pp. 827–838, Mar. 2011.
- [20] D. Barber and A. T. Cemgil, "Graphical models for time-series," *IEEE Signal Process. Mag.*, vol. 27, no. 6, pp. 18–28, 2010.