

# Bayesian Regularization and Nonnegative Deconvolution for Room Impulse Response Estimation

Yuanqing Lin and Daniel D. Lee

**Abstract**—This paper proposes *Bayesian Regularization And Nonnegative Deconvolution* (BRAND) for accurately and robustly estimating acoustic room impulse responses for applications such as time-delay estimation and echo cancellation. Similar to conventional deconvolution methods, BRAND estimates the coefficients of convolutive finite-impulse-response (FIR) filters using least-square optimization. However, BRAND exploits the nonnegative, sparse structure of acoustic room impulse responses with nonnegativity constraints and  $L_1$ -norm sparsity regularization on the filter coefficients. The optimization problem is modeled within the context of a probabilistic Bayesian framework, and expectation-maximization (EM) is used to derive efficient update rules for estimating the optimal regularization parameters. BRAND is demonstrated on two representative examples, subsample time-delay estimation in reverberant environments and acoustic echo cancellation. The results presented in this paper show the advantages of BRAND in high temporal resolution and robustness to ambient noise compared with other conventional techniques.

**Index Terms**—Bayesian regularization, echo cancellation, nonnegative deconvolution, time-delay estimation.

## I. INTRODUCTION

**S**IGNAL processing algorithms for a wide variety of acoustic applications, including source localization and echo cancellation, fundamentally rely upon estimating room impulse responses. These algorithms are typically based upon an analysis of the following linear time-invariant system:

$$x(t) = \int_{\tau} h(\tau)s(t - \tau) d\tau + n(t) \quad (1)$$

where a measured signal  $x(t)$  is given by the convolution of the acoustic source signal  $s(t)$  and the room impulse response  $h(t)$ , corrupted by additive noise  $n(t)$ . In this paper, we describe a novel algorithm for identifying the system  $h(t)$ , given a known source  $s(t)$  and an observation  $x(t)$ . Our algorithm is based upon solving the deconvolution problem

$$\min_{h(t)} \int dt \frac{1}{2} \left\| x(t) - \int dt' h(t')s(t - t') \right\|^2. \quad (2)$$

By incorporating nonnegativity constraints and Bayesian  $L_1$ -norm sparsity regularization on the filter coefficients, we

show how the estimation of  $h(t)$  is both highly time-resolved and robust to noise. We contrast the results of our *Bayesian Regularization And Nonnegative Deconvolution* (BRAND) algorithm with other algorithms based upon cross correlation or linear deconvolution on two representative acoustic signal processing problems: subsample time-delay estimation for source localization and room impulse response estimation for echo cancellation. In these examples, we demonstrate how the acoustic room impulse response can be accurately and robustly estimated using BRAND, and how it differs from other techniques.

Conventional cross correlation can be viewed as an optimization of (2) under the assumption that the room impulse response is a delta function, namely,  $h(t) = \alpha_{\tau}\delta(t - \tau)$ . With this assumption, the optimal estimates of  $\tau$  and  $\alpha_{\tau}$  given the source signal  $s(t)$  and the measured signal  $x(t)$  are [1]

$$\hat{\tau} = \arg \max_{\tau} \int x(t)s(t - \tau) dt \quad (3)$$

$$\hat{\alpha}_{\tau} = \frac{\int x(t)s(t - \hat{\tau}) dt}{\int s(t)^2 dt}. \quad (4)$$

Equations (3) and (4) show that the optimal estimates are related to the maximal value of the cross correlation between  $s(t)$  and  $x(t)$ . Because of its computational simplicity, cross correlation has been widely adopted time-delay estimation. However, the underlying assumption of a delta-function impulse response causes cross-correlation estimates to degrade in reverberant environments where multipath reflections are not negligible. Generalized cross-correlation techniques such as the phase alignment transform prewhiten the signals before performing cross correlation to help alleviate some of these difficulties [1], but their effectiveness is still limited by the underlying simple delta-function assumption on the room impulse response.

On the other hand, the least-square optimization of (2) without any constraints on the impulse response is equivalent to conventional linear deconvolution [2], [3]. With discrete-time signals, (2) can be written in matrix form as

$$\min_{\alpha} \frac{1}{2} \|\mathbf{x} - \mathbf{S}\alpha\|^2 \quad (5)$$

where  $\mathbf{x} = [x(t_1) \ x(t_2) \ \cdots \ x(t_N)]^T$  is an  $N \times 1$  vector,  $\mathbf{S} = [s(t - \Delta t_1) \ s(t - \Delta t_2) \ \cdots \ s(t - \Delta t_M)]$  is an  $N \times M$  matrix consisting of time-shifted column vectors  $\mathbf{s}(t) = [s(t_1) \ s(t_2) \ \cdots \ s(t_N)]^T$ .  $\alpha$  is a vector of  $M$  discrete samples of the impulse response  $h(t)$  at the time delays  $\{\Delta t_i\}_{i=1}^M$ . When the number of measurements  $N$  is larger than the number

Manuscript received October 16, 2004; revised April 25, 2005. This work has been supported by the U.S. National Science Foundation and Army Research Office. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Simon J. Godsill.

The authors are with the GRASP Laboratory, Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104-6228 USA (e-mail: linyuanq@seas.upenn.edu; ddlee@seas.upenn.edu).

Digital Object Identifier 10.1109/TSP.2005.863030

of time lags  $M$ , the resulting matrix optimization can be solved by taking the pseudo-inverse

$$\boldsymbol{\alpha} = (\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T \mathbf{x}. \quad (6)$$

This optimal solution can also be approximated in an online fashion through stochastic gradient descent

$$\boldsymbol{\alpha}_i \leftarrow \boldsymbol{\alpha}_{i-1} + \mu \mathbf{s}_i e_i \quad (7)$$

where the source vector  $\mathbf{s}_i = [s_i, s_{i+1}, \dots, s_{i-M+1}]^T$ , reconstruction error  $e_i = x_i - \boldsymbol{\alpha}_{i-1}^T \mathbf{s}_i$ , and  $\mu$  is learning rate. The normalized least-mean-square (NLMS) algorithm [3] used in adaptive echo cancellation is related to (7) by adaptively choosing the learning rate in the following manner:

$$\boldsymbol{\alpha}_i \leftarrow \boldsymbol{\alpha}_{i-1} + \frac{\mu'}{\mathbf{s}_i^T \mathbf{s}_i + \eta} \mathbf{s}_i e_i \quad (8)$$

where  $\mu' \approx 1$  and  $\eta \ll 1$  are dimensionless constants. Stochastic gradient descent algorithms such as NLMS distribute the estimation of filter coefficients  $\boldsymbol{\alpha}$  over time and have been widely adopted in real-time implementations.

Linear deconvolution in (5) makes no assumption about the characteristics of the acoustic room impulse response. Unfortunately, these algorithms can suffer from both poor temporal resolution and sensitivity to noise. When the signals are bandwidth-limited, the temporal resolution of linear deconvolution algorithms is limited by the near-degeneracy of the columns of matrix  $\mathbf{S}$  in (5). In the extreme case when (6) is used to solve for filter parameters with subsample temporal precision, the ill-conditioning of the matrix  $\mathbf{S}^T \mathbf{S}$  can give rise to wildly fluctuating room impulse response estimates. Thus, any noise present in the system will be greatly amplified by the deconvolution algorithms. Employing longer data sequences may help to better condition the estimates for the impulse responses, but this will also limit the speed of convergence of the estimates. In the following sections, we demonstrate how the BRAND algorithm can overcome some of these difficulties by utilizing prior knowledge about the structure of the room impulse responses.

In particular, BRAND relies upon a theoretical image model for room acoustics which predicts room impulse responses are well described by sparse, nonnegative filter coefficients [4]. Accordingly, we first introduce nonnegativity constraints on the filter coefficients in the deconvolution problem:

$$\min_{\boldsymbol{\alpha} \geq 0} \frac{1}{2} \|\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}\|^2. \quad (9)$$

Recent work has shown that nonnegativity constraints can help to regularize solutions for estimation problems such as (9), [5], [6]. With no noise in measured signals, we show that nonnegative deconvolution is able to precisely resolve the room impulse response with high temporal resolution.

Second, we exploit the sparse structure of acoustic room impulse responses to improve noise robustness by introducing  $L_1$ -norm sparsity regularization in the deconvolution optimization

$$\min_{\boldsymbol{\alpha} \geq 0} \frac{1}{2} \|\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}\|^2 + \hat{\lambda}' \sum_i \alpha_i \quad (10)$$

where the regularization parameter  $\hat{\lambda}'$  controls the sparsity of the resulting estimate by giving preference to solutions with small  $L_1$ -norm [7], [8].  $L_1$ -norm regularization has been previously used to help regularize the linear deconvolution problem [9] and nonnegative matrix factorization [10], but the regularization parameter was heuristically determined. Since the choice of regularization parameter may be critical for good estimates, we formulate the nonnegative deconvolution in (10) within a probabilistic framework, leading to an expectation-maximization (EM) procedure that infers the optimal regularization parameters. This framework also allows us to efficiently generalize the uniform regularization in (10). Instead of a single regularization parameter, each of the individual filter coefficients can be associated with an independent regularization parameter  $\hat{\lambda}_i$

$$\min_{\boldsymbol{\alpha} \geq 0} \frac{1}{2} \|\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}\|^2 + \sum_i \hat{\lambda}_i \alpha_i. \quad (11)$$

Independent regularization has been used previously in the relevance vector machine [11], where independent  $L_2$ -norm regularization was employed to regularize nonlinear regression problems. We show that the independent prior results in stronger sparsity regularization than the uniform prior does.

The remainder of the paper is arranged as follows. In Section II, a Bayesian framework is presented for the  $L_1$ -norm regularized nonnegative deconvolution problem, and iterative update rules for computing the optimal regularization parameters are derived. Two computational procedures for solving the associated nonnegative quadratic programming problems are introduced, one based upon a modified simplex method and the other on parallel multiplicative updates. In Section III, BRAND with the modified simplex method is applied to subsample time-delay estimation in a simulated reverberant environment. Its performance is compared with cross-correlation-based methods as well as linear deconvolution. Different regularization strategies are also compared to illustrate the advantages of BRAND for estimating regularization parameters. In Section IV, BRAND with multiplicative update rules is employed to estimate the room impulse response for echo cancellation, and its performance is compared to the other adaptive filtering algorithms. Finally, we conclude with a discussion of these results in Section V.

## II. BRAND PROBABILISTIC MODEL

In this section, the probabilistic generative model for BRAND and Bayesian inference is used to derive optimal regularization parameters, leading to estimates of filter coefficients with appropriate sparseness.

The underlying probabilistic model assumes the measured signal  $x(t)$  is generated by convolving the source signal  $s(t)$  with a sparse, nonnegative room impulse response. The measurement  $x(t)$  is corrupted by additive Gaussian noise with zero mean and covariance  $\sigma^2$  so that the conditional likelihood is

$$P(\mathbf{x} | \mathbf{S}, \boldsymbol{\alpha}, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}\|^2\right). \quad (12)$$

Each of the filter coefficients in  $\boldsymbol{\alpha}$  is assumed to have been drawn from its own independent exponential distribution. This exponential prior only has support in the nonnegative orthant and the sharpness of the distribution is controlled by the regularization parameters  $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_M]^T$

$$P(\boldsymbol{\alpha} | \boldsymbol{\lambda}) = \prod_i \lambda_i \exp\{-\lambda_i \alpha_i\}, \quad \boldsymbol{\alpha} \geq 0. \quad (13)$$

To infer the optimal settings of the regularization parameters  $\sigma^2$  and  $\boldsymbol{\lambda}$ , we need to maximize the posterior distribution

$$P(\boldsymbol{\lambda}, \sigma^2 | \mathbf{x}, \mathbf{S}) = \frac{P(\mathbf{x} | \boldsymbol{\lambda}, \sigma^2, \mathbf{S}) P(\boldsymbol{\lambda}, \sigma^2)}{P(\mathbf{x} | \mathbf{S})}. \quad (14)$$

Assuming that  $P(\boldsymbol{\lambda}, \sigma^2)$  is relatively flat [12],  $\sigma^2$  and  $\boldsymbol{\lambda}$  are estimated by maximizing the marginal likelihood [13]

$$\begin{aligned} P(\mathbf{x} | \boldsymbol{\lambda}, \sigma^2, \mathbf{S}) &= \int_{\boldsymbol{\alpha} \geq 0} d\boldsymbol{\alpha} P(\mathbf{x} | \mathbf{S}, \boldsymbol{\alpha}, \sigma^2) P(\boldsymbol{\alpha} | \boldsymbol{\lambda}) \\ &= \frac{\prod_i \lambda_i}{(2\pi\sigma^2)^{N/2}} \int_{\boldsymbol{\alpha} \geq 0} d\boldsymbol{\alpha} \exp[-F(\boldsymbol{\alpha})] \end{aligned} \quad (15)$$

where

$$F(\boldsymbol{\alpha}) = \frac{1}{2\sigma^2} (\mathbf{x} - \mathbf{S}\boldsymbol{\alpha})^T (\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}) + \boldsymbol{\lambda}^T \boldsymbol{\alpha}. \quad (16)$$

Unfortunately, the integral in (15) cannot be directly maximized. An alternative approach would be to heuristically derive self-consistent fixed-point equations and to solve those equations iteratively [11], [12], [14]. In BRAND, the EM procedure [15] is used to derive iterative update rules for  $\boldsymbol{\lambda}$  and  $\sigma^2$  as

$$\frac{1}{\lambda_i} \leftarrow \int_{\boldsymbol{\alpha} \geq 0} d\boldsymbol{\alpha} \alpha_i Q(\boldsymbol{\alpha}) \quad (17)$$

$$\sigma^2 \leftarrow \frac{1}{N} \int_{\boldsymbol{\alpha} \geq 0} d\boldsymbol{\alpha} (\mathbf{x} - \mathbf{S}\boldsymbol{\alpha})^T (\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}) Q(\boldsymbol{\alpha}) \quad (18)$$

where the required expectations are taken over the distribution

$$Q(\boldsymbol{\alpha}) = \frac{\exp[-F(\boldsymbol{\alpha})]}{\mathcal{Z}_\alpha} \quad (19)$$

with normalization constant  $\mathcal{Z}_\alpha = \int_{\boldsymbol{\alpha} \geq 0} d\boldsymbol{\alpha} \exp[-F(\boldsymbol{\alpha})]$ . In the EM updates,  $\mathbf{x}$  is treated as the observed variable,  $\boldsymbol{\alpha}$  as the hidden variables, and both  $\boldsymbol{\lambda}$  and  $\sigma^2$  as the parameters. These updates ((17) and (18)) have guaranteed convergence properties and can be intuitively understood as iteratively reestimating the optimal  $\boldsymbol{\lambda}$  and  $\sigma^2$  from the statistics of the current estimate for  $Q(\boldsymbol{\alpha})$ . Unfortunately, the integrals in (17) and (18) are difficult to compute analytically; thus, we seek an appropriate approximation of  $Q(\boldsymbol{\alpha})$  with respect to the mode of the distribution,  $\boldsymbol{\alpha}^{\text{ML}}$ .

#### A. Estimation of $\boldsymbol{\alpha}^{\text{ML}}$

The integrals in (17) and (18) are dominated by  $\boldsymbol{\alpha} \approx \boldsymbol{\alpha}^{\text{ML}}$  where the most likely  $\boldsymbol{\alpha}^{\text{ML}}$  is given by

$$\boldsymbol{\alpha}^{\text{ML}} = \arg \min_{\boldsymbol{\alpha} \geq 0} \frac{1}{2\sigma^2} (\mathbf{x} - \mathbf{S}\boldsymbol{\alpha})^T (\mathbf{x} - \mathbf{S}\boldsymbol{\alpha}) + \boldsymbol{\lambda}^T \boldsymbol{\alpha}. \quad (20)$$

This optimization is equivalent to the nonnegative quadratic programming problem in (11) with  $\hat{\lambda}_i = \lambda_i \sigma^2$ . This can be written in standard quadratic form

$$\min_{\boldsymbol{\alpha} \geq 0} \frac{1}{2} \boldsymbol{\alpha}^T \mathbf{A} \boldsymbol{\alpha} + \mathbf{b}^T \boldsymbol{\alpha} \quad (21)$$

with  $\mathbf{A} = (1/\sigma^2) \mathbf{S}^T \mathbf{S}$ , and  $\mathbf{b} = -(1/\sigma^2) \mathbf{S}^T \mathbf{x} + \boldsymbol{\lambda}$ . These quadratic terms can be efficiently computed in the frequency domain [16] when edge effects are not significant, as follows:

$$A_{ik} = \Re \left[ \frac{1}{\sigma^2} \sum_f |S(f)|^2 e^{-j2\pi f(\Delta t_i - \Delta t_k)} \right] \quad (22)$$

$$b_i = \Re \left[ -\frac{1}{\sigma^2} \sum_f X(f) S^*(f) e^{j2\pi f \Delta t_i} \right] + \lambda_i \quad (23)$$

where  $\Delta \mathbf{t}$  is the time-delay set,  $S(f)$  and  $X(f)$  are the respective discrete Fourier transforms of  $\mathbf{s}$  and  $\mathbf{x}$ ,  $S^*(f)$  is the complex conjugate of  $S(f)$ , and  $\Re(\cdot)$  denotes taking the real component of a complex number. A frequency domain implementation is desirable when subsample time resolution is required. In particular, the approximate form of  $\mathbf{A}$  in (22) has Toeplitz structure, which further improves the efficiency in computing  $\mathbf{A}$ .

In order to solve the nonnegative quadratic programming problem in (21), two different but complementary methods have been recently developed. The first method is based upon finding a solution  $\boldsymbol{\alpha}^{\text{ML}}$  that satisfies the Karush–Kuhn–Tucker (KKT) conditions [17] for (21)

$$\begin{aligned} \mathbf{A} \boldsymbol{\alpha} + \mathbf{b} &= \boldsymbol{\beta}, \quad \boldsymbol{\alpha} \geq 0, \quad \boldsymbol{\beta} \geq 0, \quad \alpha_i \beta_i = 0, \\ & \quad i = 1, 2, \dots, M. \end{aligned} \quad (24)$$

By introducing additional artificial variables  $\mathbf{a}$ , the KKT conditions can be transformed into the linear optimization of  $\min \sum_i a_i$  subject to the constraints

$$\begin{aligned} \mathbf{A} \boldsymbol{\alpha} - \boldsymbol{\beta} + \text{sign}(-\mathbf{b}) \mathbf{a} &= -\mathbf{b} \\ \mathbf{a} \geq 0, \quad \boldsymbol{\alpha} \geq 0, \quad \boldsymbol{\beta} \geq 0, \quad \alpha_i \beta_i &= 0, \quad i = 1, 2, \dots, M. \end{aligned} \quad (25)$$

The only nonlinear constraints in this optimization are the products  $\alpha_i \beta_i = 0$ . These constraints can be handled in the simplex method by modifying the simplex variable selection procedure to ensure that  $\alpha_i$  and  $\beta_i$  are never simultaneously in the basic variable set. With this simple modification, all the artificial variables can be driven to zero and thus the optimal  $\boldsymbol{\alpha}$  satisfying the KKT conditions (24) is found. One beneficial feature of the modified simplex method is that its convergence is relatively insensitive to the conditioning of  $\mathbf{A}$ , which may be large when high temporal resolution is desired. Unfortunately, the modified simplex method cannot generically handle very large scale problems, such as when the dimensionality  $M$  of  $\boldsymbol{\alpha}$  becomes larger than 1000.

For these high-dimensional problems, we recently developed an algorithm with parallel, multiplicative updates [18]. These multiplicative updates have the advantage of guaranteed convergence and no adjustable convergence parameters to tune. We first decompose the matrix  $\mathbf{A} = \mathbf{A}^+ - \mathbf{A}^-$  in (21) into its positive and negative components such that

$$A_{ij}^+ = \begin{cases} A_{ij} & \text{if } A_{ij} > 0 \\ 0 & \text{if } A_{ij} \leq 0 \end{cases} \quad A_{ij}^- = \begin{cases} 0 & \text{if } A_{ij} \geq 0 \\ -A_{ij} & \text{if } A_{ij} < 0 \end{cases}. \quad (26)$$

In terms of the nonnegative matrices  $\mathbf{A}^+$  and  $\mathbf{A}^-$ , the following is an auxiliary function that upper bounds (21) for all  $\tilde{\boldsymbol{\alpha}} > 0$ :

$$G(\boldsymbol{\alpha}, \tilde{\boldsymbol{\alpha}}) = \mathbf{b}^T \boldsymbol{\alpha} + \frac{1}{2} \sum_i \frac{(\mathbf{A}^+ \tilde{\boldsymbol{\alpha}})_i}{\tilde{\alpha}_i} \alpha_i^2 - \frac{1}{2} \sum_{i,j} A_{ij}^- \tilde{\alpha}_i \tilde{\alpha}_j \left( 1 + \ln \frac{\alpha_i \alpha_j}{\tilde{\alpha}_i \tilde{\alpha}_j} \right). \quad (27)$$

Minimizing (27) gives the following iterative rule for updating the current estimate for  $\boldsymbol{\alpha}$ :

$$\alpha_i \leftarrow \alpha_i \frac{-b_i + \sqrt{b_i^2 + 4(\mathbf{A}^+ \boldsymbol{\alpha})_i (\mathbf{A}^- \boldsymbol{\alpha})_i}}{2(\mathbf{A}^+ \boldsymbol{\alpha})_i}. \quad (28)$$

Equation (28) can be interpreted as an interior point method which explicitly preserves the nonnegativity constraints on  $\boldsymbol{\alpha}$ , with guaranteed convergence to the optimal  $\boldsymbol{\alpha}^{\text{ML}}$ . Note that both matrices  $\mathbf{A}^+$  and  $\mathbf{A}^-$  are Toeplitz if  $\mathbf{A}$  is Toeplitz so the matrix-vector products in (28) can be efficiently computed using fast Fourier transforms (FFTs). Thus, these computations can be efficiently implemented in real time for even fairly large optimization problems.

### B. Approximation of $Q(\boldsymbol{\alpha})$

After  $\boldsymbol{\alpha}^{\text{ML}}$  has been determined, one simple approach for re-estimating the parameters  $\boldsymbol{\lambda}$  and  $\sigma^2$  in (17) and (18) is to approximate the distribution  $Q(\boldsymbol{\alpha}) \approx \delta(\boldsymbol{\alpha} - \boldsymbol{\alpha}^{\text{ML}})$  in the EM updates. Unfortunately, this simple approximation can cause estimates for  $\boldsymbol{\lambda}$  and  $\sigma$  to diverge with bad initial estimates. To overcome these difficulties, we require a better approximation to properly account for variability in the distribution  $Q(\boldsymbol{\alpha})$ .

We first note that the solution  $\boldsymbol{\alpha}^{\text{ML}}$  for the nonnegative quadratic optimization in (21) naturally partitions its elements into two distinct subsets  $\boldsymbol{\alpha}_I$  and  $\boldsymbol{\alpha}_J$ , consisting of components  $i \in I$  such that  $(\boldsymbol{\alpha}^{\text{ML}})_i = 0$ , and components  $j \in J$  such that  $(\boldsymbol{\alpha}^{\text{ML}})_j > 0$ , respectively. The joint distribution  $Q(\boldsymbol{\alpha})$  is then approximated by the factorized distribution

$$Q(\boldsymbol{\alpha}) \approx Q_I(\boldsymbol{\alpha}_I) Q_J(\boldsymbol{\alpha}_J). \quad (29)$$

For the components  $\boldsymbol{\alpha}_J$ , none of the nonnegativity constraints are active, so it is reasonable to approximate the distribution  $Q_J(\boldsymbol{\alpha}_J)$  by the unconstrained Gaussian

$$Q_J(\boldsymbol{\alpha}_J) \propto \exp[-F(\boldsymbol{\alpha}_J | \boldsymbol{\alpha}_I = 0)]. \quad (30)$$

This Gaussian distribution has mean  $\boldsymbol{\alpha}_J^{\text{ML}}$  and inverse covariance given by the submatrix  $\mathbf{A}_{JJ}$  of  $\mathbf{A} = (1/\sigma^2) \mathbf{S}^T \mathbf{S}$ .

However, for the other components  $\boldsymbol{\alpha}_I$ , since  $\boldsymbol{\alpha}_I^{\text{ML}} = 0$  is on the boundary of the distribution, it is important to consider the effects of nonnegativity constraints. This distribution is represented by first considering the Taylor expansion about  $\boldsymbol{\alpha}^{\text{ML}}$

$$\begin{aligned} Q_I(\boldsymbol{\alpha}_I \geq 0) &\propto \exp \left\{ - \left[ \left( \frac{\partial F}{\partial \boldsymbol{\alpha}} \right) \Big|_{\boldsymbol{\alpha}^{\text{ML}}} \right]^T \boldsymbol{\alpha}_I - \frac{1}{2} \boldsymbol{\alpha}_I^T \mathbf{A}_{II} \boldsymbol{\alpha}_I \right\} \\ &\propto \exp \left[ -(\mathbf{A} \boldsymbol{\alpha}^{\text{ML}} + \mathbf{b})_I^T \boldsymbol{\alpha}_I - \frac{1}{2} \boldsymbol{\alpha}_I^T \mathbf{A}_{II} \boldsymbol{\alpha}_I \right]. \end{aligned} \quad (31)$$

Since the first-order derivative is nonzero at  $\boldsymbol{\alpha}_I^{\text{ML}}$ , we use a variational approximation for  $Q_I(\boldsymbol{\alpha}_I)$ , representing it by the exponential distribution

$$\hat{Q}_I(\boldsymbol{\alpha}_I \geq 0) = \prod_{i \in I} \frac{1}{\mu_i} e^{-\alpha_i / \mu_i}. \quad (32)$$

The variational parameters  $\boldsymbol{\mu}$  are defined by minimizing the Kullback-Leibler divergence between  $Q_I$  and  $\hat{Q}_I$

$$\min_{\boldsymbol{\mu} \geq 0} \int_{\boldsymbol{\alpha}_I \geq 0} d\boldsymbol{\alpha}_I \hat{Q}_I(\boldsymbol{\alpha}_I) \ln \frac{\hat{Q}_I(\boldsymbol{\alpha}_I)}{Q_I(\boldsymbol{\alpha}_I)}. \quad (33)$$

The integral in (33) can easily be computed and yields the following optimization for  $\boldsymbol{\mu}$ :

$$\min_{\boldsymbol{\mu} \geq 0} \hat{\mathbf{b}}_I^T \boldsymbol{\mu} + \frac{1}{2} \boldsymbol{\mu}^T \hat{\mathbf{A}} \boldsymbol{\mu} - \sum_{i \in I} \ln \mu_i \quad (34)$$

where  $\hat{\mathbf{b}}_I = (\mathbf{A} \boldsymbol{\alpha}^{\text{ML}} + \mathbf{b})_I$ ,  $\hat{\mathbf{A}} = \mathbf{A}_{II} + \text{diag}(\mathbf{A}_{II})$ . To solve this minimization problem, we again use an auxiliary function for (34) that is similar to (27) used for nonnegative quadratic programming

$$\begin{aligned} G'(\boldsymbol{\mu}, \tilde{\boldsymbol{\mu}}) &= \hat{\mathbf{b}}_I^T \boldsymbol{\mu} + \frac{1}{2} \sum_{i \in I} \frac{(\hat{\mathbf{A}}^+ \tilde{\boldsymbol{\mu}})_i}{\tilde{\mu}_i} \mu_i^2 \\ &\quad - \frac{1}{2} \sum_{i,j \in I} \hat{A}_{ij}^- \tilde{\mu}_i \tilde{\mu}_j \left( 1 + \ln \frac{\mu_i \mu_j}{\tilde{\mu}_i \tilde{\mu}_j} \right) - \sum_{i \in I} \ln \mu_i \end{aligned} \quad (35)$$

where  $\hat{\mathbf{A}} = \hat{\mathbf{A}}^+ - \hat{\mathbf{A}}^-$  is the decomposition of  $\hat{\mathbf{A}}$  into its positive and negative components.

The parameters  $\mu_i$  in (32) can then be solved by iterating

$$\mu_i \leftarrow \mu_i \frac{-\hat{b}_i + \sqrt{\hat{b}_i^2 + 4(\hat{\mathbf{A}}^+ \boldsymbol{\mu})_i \left[ (\hat{\mathbf{A}}^- \boldsymbol{\mu})_i + \frac{1}{\mu_i} \right]}}{2(\hat{\mathbf{A}}^+ \boldsymbol{\mu})_i}. \quad (36)$$

These iterations again have guaranteed convergence to the optimal approximate variational distribution  $\hat{Q}_I(\boldsymbol{\alpha}_I)$ .

Using the factorized approximation  $Q(\boldsymbol{\alpha}) = \hat{Q}_I(\boldsymbol{\alpha}_I) Q_J(\boldsymbol{\alpha}_J)$ , the expectations in (17) and (18) can be analytically calculated. The mean value of  $\boldsymbol{\alpha}$  under this distribution is given by

$$\bar{\alpha}_i = \begin{cases} \alpha_i^{\text{ML}} & \text{if } i \in J \\ \mu_i & \text{if } i \in I \end{cases} \quad (37)$$

and its covariance  $\mathbf{C}$  is

$$C_{ij} = \begin{cases} (\mathbf{A}_{JJ}^{-1})_{ij} & \text{if } i, j \in J \\ \mu_i^2 \delta_{ij} & \text{otherwise} \end{cases}. \quad (38)$$

The update rules for  $\boldsymbol{\lambda}$  and  $\sigma^2$  are then given by

$$\lambda_i \leftarrow \frac{1}{\bar{\alpha}_i} \quad (39)$$

$$\sigma^2 \leftarrow \frac{1}{N} [(\mathbf{x} - \mathbf{S} \bar{\boldsymbol{\alpha}})^T (\mathbf{x} - \mathbf{S} \bar{\boldsymbol{\alpha}}) + \text{Tr}(\mathbf{S}^T \mathbf{S} \mathbf{C})] \quad (40)$$

To summarize, the complete BRAND algorithm for estimating the nonnegative filter coefficients and independent regularization parameters consists of the following steps.

- 1) Initialize  $\boldsymbol{\lambda}$  and  $\sigma^2$ , and choose a discrete set of possible time delays  $\{\Delta t_i\}$ .

- 2) Determine  $\alpha^{\text{ML}}$  by solving the nonnegative quadratic programming in (20).
- 3) Approximate the distribution  $Q(\alpha) \approx \hat{Q}_I(\alpha_I)Q_J(\alpha_J)$  by solving the variational equations for  $\mu$  using (36).
- 4) Calculate the mean  $\bar{\alpha}$  and covariance  $\mathbf{C}$  for this distribution.
- 5) Reestimate regularization parameters  $\lambda$  and  $\sigma^2$  using (39) and (40).
- 6) Go back to Step 2) until convergence.

Under this model, a different prior on  $\alpha$  results in a uniform regularization as described in (10). This uniform model is associated with the distribution

$$P(\alpha | \lambda') = (\lambda')^M \exp\left\{-\lambda' \sum_i \alpha_i\right\}, \quad \alpha \geq 0. \quad (41)$$

In this case, the iterative estimates are similar to before except that (16), (17), and (39) become

$$F(\alpha) = \frac{1}{2\sigma^2}(\mathbf{x} - \mathbf{S}\alpha)^T(\mathbf{x} - \mathbf{S}\alpha) + \lambda' \mathbf{e}^T \alpha \quad (42)$$

$$\frac{1}{\lambda'} \leftarrow \frac{1}{M} \int_{\alpha \geq 0} d\alpha \mathbf{e}^T \alpha Q(\alpha) \quad (43)$$

$$\lambda' \leftarrow \frac{M}{\sum_i \bar{\alpha}_i} \quad (44)$$

respectively, where  $\mathbf{e} = [1, 1, \dots, 1]^T$  is the uniform vector [19].

Since there are fewer parameters to estimate, a uniform regularization prior is convenient to use at the beginning of the algorithm. Using an independent prior results in stronger sparsity regularization, but with more parameters to estimate. Thus, in our implementation of BRAND, a uniform regularization is initially used in the beginning iterations and then the solution is optimally refined using independent regularization parameters.

### III. BRAND FOR SUBSAMPLE TIME-DELAY ESTIMATION IN REVERBERANT ENVIRONMENTS

In this section, we demonstrate BRAND for estimating the time delays present in a simulated reverberant environment. For the source signal, a short segment of human speech (512 samples with sampling time  $T_s = 62.5 \mu\text{s}$ ) was used. This speech signal was convolved with a simulated room impulse response to model the measured signal. The signals are shown in Fig. 1. With no added noise, we first show how nonnegative deconvolution is able to precisely resolve the time-delay structure of the room impulse response, and compare its results with cross correlation and linear deconvolution methods. In addition, with added noise in the measured signal, we show how BRAND is able to robustly recover the room impulse response even at 10-dB signal-to-noise ratio (SNR) and demonstrate the differences between various regularization strategies.

#### A. Time-Delay Estimation With Zero Ambient Noise

With no noise ( $n(t) = 0$ ) added to the measured signal, Fig. 2 shows the results of estimating the room impulse response using several different methods. All the estimates were performed over a range from  $-10T_s$  to  $+10T_s$ , with discrete time increments of  $0.25T_s$ . Because the speech source signal




Fig. 1. Signals used for the simulation: (a) source signal  $s(t)$ , (b) source spectrum,  $|S(f)|$ , (c) the simulated measurement  $x(t) = s(t - T_s) + 0.5s(t - 8.75T_s) + n(t)$ , where  $T_s = 62.5 \mu\text{s}$  is the sample interval and  $n(t)$  is varying levels of ambient noise.

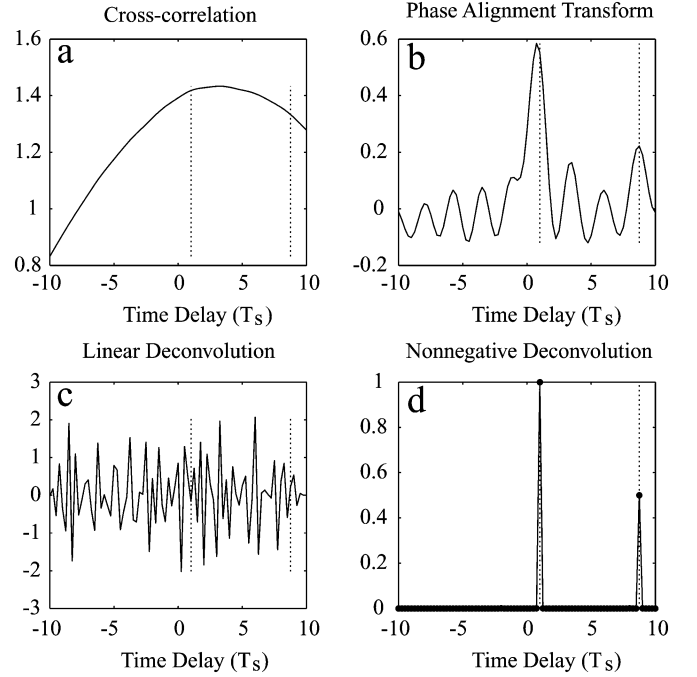


Fig. 2. Time-delay estimation of the room impulse response  $h(t) = \delta(t - T_s) + 0.5\delta(t - 8.75T_s)$  by (a) cross correlation, (b) phase alignment transform, (c) linear deconvolution, and (d) nonnegative deconvolution. The vertical dotted lines in each plot indicate the true positions of the time delays  $\Delta t = T_s$  and  $\Delta t = 8.75T_s$ , respectively.

has limited bandwidth, the cross correlation in Fig. 2(a) shows only a broad main lobe, resulting in poor time-delay resolution. Due to the multipath reflection, the peak of the cross correlation function estimates neither of the time delays present in the room impulse response. The phase alignment transform (PHAT) performs better than simple cross correlation, as shown in Fig. 2(b).




Fig. 3. Estimate of  $\sigma^2$  at each iteration of the BRAND algorithm for signals with varying levels of noise. Uniform regularization was employed for the first ten iterations, followed by another ten iterations of independent regularization.

PHAT whitens the signals before cross correlation [1], but again since the signals are not ideally white, there are still some errors in the time-delay estimation. PHAT also significantly degrades in performance with the presence of any ambient noise. Even worse results arise from linear deconvolution as shown in Fig. 2(c). The ill-conditioning of the source correlation matrix causes the resulting estimates of the room impulse response to fluctuate wildly.

The dramatic effect of nonnegativity constraints in regularizing the deconvolution problem is displayed in Fig. 2(d). Nonnegative deconvolution is able to precisely resolve the room impulse response, including the multipath time delays and amplitudes of the filter coefficients. In this example, the modified simplex algorithm was used to efficiently compute the estimated filter coefficients, resulting in rapid convergence.

In the above example, the true time delays were in the initial discrete set from  $-10T_s$  to  $+10T_s$  in increments of  $0.25T_s$ . A more general scheme is to start with a relatively coarse time-delay set and adaptively refine the time delays around the nonzero filter coefficients until the desired temporal resolution is achieved.

### B. Time-Delay Estimation With Ambient Noise

Here, we show the robustness of the BRAND algorithm to varying levels of additive noise in the measured signal. In addition to the coefficients of the room impulse response, the regularization parameters ( $\lambda$  and  $\sigma^2$ ) are iteratively estimated by (39) (or (44) for uniform regularization) and (40) as described in the previous section.

The same source signal  $s(t)$  and measured signal  $x(t)$  as shown in Fig. 1 were used in these simulations. The measured signal  $x(t)$  was corrupted with varying levels of added ambient noise, and Fig. 3 illustrates that BRAND is able to quickly and consistently estimate the true noise level  $\sigma^2$  even with bad initial estimates. The associated nonnegative deconvolutions were solved using the modified simplex method.

Fig. 4 illustrates the need for the BRAND algorithm to infer the optimal setting of the regularization parameters. When the measured signal is contaminated with  $-10$  dB ambient

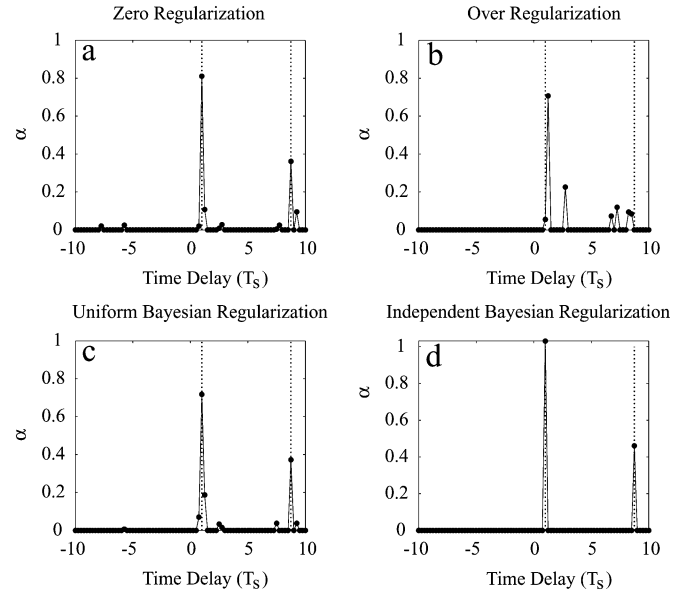


Fig. 4. Nonnegative deconvolution results under different  $L_1$ -norm regularizations, when the measured signal is contaminated by  $-10$  dB noise: (a) zero regularization; (b) manually set overregularization; (c) uniform Bayesian regularization; and (d) independent Bayesian regularization.

Gaussian white noise, different regularization strategies can lead to different filter estimates. With no regularization, the added noise causes the deconvolution solution to exhibit several small spurious peaks. However, manually setting too large of a regularization causes the time-delay estimates to deviate from the true room impulse response structure. The uniform Bayesian regularization strategy is much better with regards to estimating the true sparse structure of the filter, but the absolute magnitudes of the filter coefficients are underestimated. In contrast, the independent Bayesian regularization in BRAND leads to an almost perfect reconstruction of the room impulse response with the appropriate sparse structure.

## IV. BRAND FOR ADAPTIVE ECHO CANCELLATION

In this section, we demonstrate the utility of BRAND for echo cancellation. In acoustic echo cancellation, adaptive algorithms are typically used to estimate the transfer function between a speaker and a microphone present in a room. The adaptive filter estimates the combination of the speaker and microphone characteristics along with the room impulse response. To apply BRAND to this problem, we assume that the speaker and microphone characteristics are already known, and only the possibly changing room impulse response needs to be estimated.

For optimal echo cancellation performance, the estimated room impulse responses need to be quite long in duration (on the order of 10–100 ms). Thus, we illustrate the application of BRAND for estimating long room impulse responses using parallel, multiplicative update rules rather than the modified simplex algorithm to solve the nonnegative deconvolution with many filter coefficients. We compare the performance of BRAND with the conventional NLMS and batch least-mean-square (LMS) algorithms, and contrast the different algorithms in terms of accuracy and convergence in the presence of varying levels of noise.




Fig. 5. Source signal  $s(t)$  for echo cancellation: (a) time domain  $s(t)$  and (b) spectrum  $|S(f)|$ .

TABLE I  
ROOM PARAMETERS USED TO CALCULATE THE  
ACOUSTIC ROOM IMPULSE RESPONSE

Parameters	Value
Room dimensions (m)	(5,5,5)
Speaker position (m)	(2,3,1)
Microphone position (m)	(4,3,1)
Reflection coefficient	0.9
Number of reflections	3
Sampling rate (Hz)	16k

For these echo cancellation simulations, a short segment of 8192 samples of human speech with sampling rate 16 kHz was used as the source signal as shown in Fig. 5. An acoustic room impulse response was calculated using the theoretical image model of a room with the parameters shown in Table I [4]. The resulting room impulse response  $h(t)$  is represented by 1024 discrete time samples and plotted in Fig. 6.

The speech source signal was convolved with the room impulse response to calculate 8192 samples of a signal  $x(t)$ . This signal was then corrupted with additive noise before BRAND and NLMS was applied. In order to perform a fair comparison between the two algorithms, we attempted to roughly match their computational complexity. For this simulation, the signal data length was  $N = 8192$ , and the unknown filter length was  $M = 1024$ . An iteration of BRAND consisted of two multiplicative updates in solving the nonnegative quadratic programming problem for  $\alpha^{\text{ML}}$ , and a re-estimation of the regularization parameters  $\lambda$  (or  $\lambda'$  for uniform regularization) and  $\sigma^2$ . In each of these computations, the Toeplitz structure of the matrices was utilized to reduce memory storage and computational requirements. As a result, each of the iterations required on the order of  $M \log M$  multiplications and additions. The sparse structure of the filter coefficients in  $\alpha^{\text{ML}}$  also helped to speed the computations, and the matrix covariance computation in (40) was approximated. On the other hand, NLMS is quite efficient in needing  $2M$  multiplications for each stochastic iteration. The number of iterations needed for convergence in the two algorithms were scaled to reflect the same amount of computation in the simulations.




Fig. 6. Room impulse response derived from image model with parameters shown in Table I. The first 1024 samples were used in the simulation.

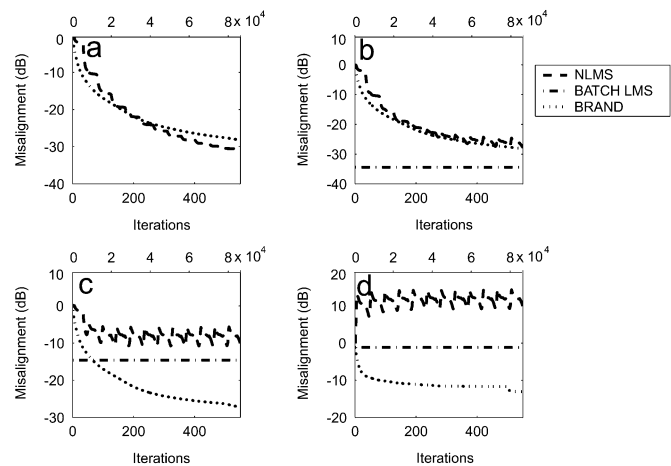


Fig. 7. Normalized misalignment of NLMS (dashed line), batch LMS (dashed-dotted line), and BRAND (dotted line) with varying levels of added noise: (a) no noise; (b)  $-40$  dB; (c)  $-20$  dB; and (d)  $0$  dB. The horizontal axes indicate the number of iterations for BRAND (lower axis) and NLMS (upper axis). Batch LMS led to zero misalignment when the signal had no noise.

We performed BRAND with uniform regularization for several hundred iterations, followed by independent regularization. Since NLMS is an online algorithm, the signals were repeatedly input to the algorithm to monitor convergence. Fig. 7 shows the performance of the algorithms when the signals are contaminated with varying levels of noise. The normalized misalignment measure

$$E = 10 \log_{10}(\|\mathbf{h} - \boldsymbol{\alpha}\|^2 / \|\mathbf{h}\|^2) \quad (45)$$

was used to evaluate the error of the estimated room impulse response  $\boldsymbol{\alpha}$  from the true impulse response  $\mathbf{h}$ . The resulting misalignment using batch LMS is also presented in the figure, indicating the performance limit of the NLMS algorithm.

At low noise conditions, Fig. 7(a) and (b) shows that the relative convergence of the two algorithms is quite comparable when computational complexity is taken into account. However, when the ambient noise level is larger than  $-20$  dB, BRAND is clearly more robust than NLMS in estimating the room impulse response. In particular, at  $0$  dB SNR, BRAND is still able to estimate filter coefficients with only  $-13$  dB normalized misalignment, whereas batch LMS yielded a limit of  $-1$  dB normalized




Fig. 8. BRAND estimation of the ambient noise level  $\sigma^2$  at each iteration. The letters in decibels indicate the true noise level.

misalignment and NLMS actually computed completely inaccurate estimates.

Not only is BRAND able to estimate the filter coefficients, but it also estimates the level of ambient noise. Fig. 8 shows the rapid convergence of the  $\sigma^2$  estimates in inferring the amount of added noise. Only when the noise level is quite small, BRAND slightly overestimates the noise variance due to incomplete convergence of the algorithm. Thus, BRAND is able to accurately estimate the ambient noise level when it becomes significant. These noise level estimates could easily serve as an integral double-talk detector in an echo cancellation application.

## V. DISCUSSION

By incorporating nonnegativity constraints and Bayesian  $L_1$ -norm sparsity regularization, BRAND is able to precisely and robustly estimate room impulse responses for both time-delay estimation and echo cancellation even in the presence of a significant amount of ambient noise. Furthermore, the estimated  $\sigma^2$  in BRAND is very useful for indicating the noise level of the signal. Depending upon the number of filter coefficients that are estimated, BRAND also exhibits good temporal resolution in its estimates.

We used simulations to enable quantitative tests of accuracy and convergence to a known room impulse response. A practical implementation would require accurate characterizations of the experimental speakers, microphones, as well as room acoustics in order to test the validity of the underlying theoretical assumptions of BRAND. How well the estimated nonnegative, sparse filter structure models true experimental room conditions remains to be seen. Nevertheless, we believe that incorporating knowledge about the physics of room acoustics in estimating their impulse responses is extremely valuable in noisy situations. Other types of constraints on room impulse response characteristics could be incorporated into BRAND by simply modifying some of the priors in the probabilistic model.

BRAND can also be extended to help solve the blind deconvolution problem. Preliminary work indicates that by alternating BRAND with a source estimation algorithm, both the sparse, nonnegative room impulse responses as well as unknown source signal can be simultaneously estimated from a set of convoluted observations [20]. Thus, besides time-delay estimation and echo cancellation applications, BRAND may have potential uses in other acoustic signal processing domains.

## ACKNOWLEDGMENT

The authors would like to thank S. Seung, L. Saul, M. Sondhi, J. Ham, and F. Sha for useful discussions and the U.S. NSF and ARO for financial support.

## REFERENCES

- [1] C. H. Knapp and G. C. Carter, "The generalized correlation method for estimation of time delay," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 24, no. 4, pp. 320–327, Aug. 1976.
- [2] H. Krim and M. Viberg, "Two decades of array signal processing research: The parametric approach," *IEEE Signal Process. Mag.*, vol. 13, no. 4, pp. 67–94, Jul. 1996.
- [3] G.-O. Glentis, K. Berberidis, and S. Theodoridis, "Efficient least squares adaptive algorithms for FIR transversal filtering," *IEEE Signal Process. Mag.*, vol. 16, no. 4, pp. 13–41, Jul. 1999.
- [4] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Amer.*, vol. 65, pp. 943–950, 1979.
- [5] D. D. Lee and H. S. Seung, "Learning the parts of objects by nonnegative matrix factorization," *Nature*, vol. 401, pp. 788–791, 1999.
- [6] L. Benvenuti and L. Farina, "A tutorial on the positive realization problem," *IEEE Trans. Autom. Control*, vol. 49, no. 5, pp. 651–664, May 2004.
- [7] B. Olshausen and D. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for nature images," *Nature*, vol. 381, pp. 607–609, 1996.
- [8] D. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via  $l^1$  minimization," in *Proc. Nat. Academy of Sciences U.S.A.*, vol. 100, 2003, pp. 2197–2202.
- [9] J. Fuchs, "Multipath time-delay detection and estimation," *IEEE Trans. Signal Process.*, vol. 47, no. 1, pp. 237–243, Jan. 1999.
- [10] P. O. Hoyer, "Non-negative sparse coding," in *Proc. IEEE Workshop Neural Networks for Signal Processing*, Sep. 2002, pp. 557–565.
- [11] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Machine Learn. Res.*, vol. 1, pp. 211–244, 2001.
- [12] D. MacKay, "Bayesian interpolation," *Neural Comput.*, vol. 4, pp. 415–447, 1992.
- [13] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*. New York: Springer, 1993.
- [14] D. Foresee and M. Hagan, "Gauss-Newton approximation to Bayesian regularization," in *Proc. 1997 Int. Joint Conf. Neural Networks*, 1997, pp. 1930–1935.
- [15] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice-Hall, 2002.
- [16] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1998.
- [17] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena Scientific, 2003.
- [18] F. Sha, L. K. Saul, and D. Lee, "Multiplicative updates for nonnegative quadratic programming in support vector machines," in *Advances in Neural Information Processing Systems*, S. T. S. Becker and K. Obermayer, Eds. Cambridge, MA: MIT Press, 2003, vol. 15.
- [19] Y. Lin and D. D. Lee, "Bayesian regularization and nonnegative deconvolution for time delay estimation," in *Advances in Neural Information Processing Systems*, L. K. Saul, Y. Weiss, and L. Bottou, Eds. Cambridge, MA: MIT Press, 2005, vol. 17.
- [20] —, "Relevant deconvolution for acoustic source estimation," in *IEEE Int. Conf. Acoustics, Speech, Signal Processing (ICASSP)*, vol. 5, Mar. 2005, pp. 529–532.





**Yuanqing Lin** received the M.S. degree in optical engineering from Tsinghua University, Beijing, China, in 1999. He is currently working toward the Ph.D. degree at the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, where he is working with Prof. D. D. Lee on machine learning techniques applied to signal processing problems.

He previously worked for two years on near-infrared spectroscopy for biomedical imaging with Dr. B. Chance at the University of Pennsylvania.



**Daniel D. Lee** received the A.B. degree in physics from Harvard University, Cambridge, MA, in 1990 and the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 1995.

He previously was on staff in the Theoretical Physics and Biological Computation departments of Bell Laboratories, Lucent Technologies. He is currently an Associate Professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia. He has been a faculty member at the University of Pennsylvania

since 2001, where his research focuses on applying knowledge of biological information processing to building better artificial sensorimotor systems.