

Fast communication

Mixed linear system estimation and identification

Argyrios Zymnis, Stephen Boyd*, Dimitry Gorinevsky

Department of Electrical Engineering, Stanford University, Packard 264, Stanford, CA 94305, USA

ARTICLE INFO

Article history:

Received 5 June 2009
 Received in revised form
 22 September 2009
 Accepted 22 September 2009
 Available online 26 September 2009

Keywords:

Statistical estimation
 Convex relaxation
 Interior-point methods

ABSTRACT

We consider a mixed linear system model, with both continuous and discrete inputs and outputs, described by a coefficient matrix and a set of noise variances. When the discrete inputs and outputs are absent, the model reduces to the usual noise-corrupted linear system. With discrete inputs only, the model has been used in fault estimation, and with discrete outputs only, the system reduces to a probit model. We consider two fundamental problems: estimating the model input, given the model parameters and the model output; and identifying the model parameters, given a training set of input–output pairs. The estimation problem leads to a mixed Boolean-convex optimization problem, which can be solved exactly when the number of discrete variables is small enough. In other cases, when the number of discrete variables is large, the estimation problem can be solved approximately, by solving a convex relaxation, rounding, and possibly, carrying out a local optimization step. The identification problem is convex and so can be exactly solved. Adding ℓ_1 regularization to the identification problem allows us to trade off model fit and model parsimony. We illustrate the identification and estimation methods with a numerical example.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

1.1. System model

In this paper we introduce a model with multiple continuous (i.e., real valued) and discrete (i.e., Boolean) inputs and outputs. The continuous outputs are linearly related to the continuous and discrete inputs and are corrupted by zero mean Gaussian noise of known variance. The discrete outputs indicate whether a linear combination of the continuous and discrete inputs, corrupted by zero mean Gaussian noise, is above or below a known threshold. As such, the model that we consider is a hybrid generalized linear model (GLM) [1–3].

The model has the form

$$y = A^{cc}x + A^{dc}d + b^c + v_c, \quad (1)$$

$$z = \text{pos}(A^{cd}x + A^{dd}d + b^d + v_d), \quad (2)$$

where

- $x \in \mathbf{R}^{n_c}$ is the continuous input,
- $d \in \{0, 1\}^{n_d}$ is the discrete input,
- $y \in \mathbf{R}^{m_c}$ is the continuous measurement or output,
- $z \in \{0, 1\}^{m_d}$ is the discrete measurement or output,
- $v_c \in \mathbf{R}^{m_c}$ is the continuous noise term, and
- $v_d \in \mathbf{R}^{m_d}$ is the discrete noise term.

The function $\text{pos} : \mathbf{R}^{m_d} \rightarrow \{0, 1\}^{m_d}$ is defined by

$$\text{pos}(u)_i = \begin{cases} 1, & u_i > 0, \\ 0, & u_i \leq 0. \end{cases}$$

Note that for any diagonal positive matrix D , we have $\text{pos}(Du) = \text{pos}(u)$.

The noises are Gaussian, with all components independent, with

$$(v_c)_i \sim \mathcal{N}(0, \sigma_i^2), \quad i = 1, \dots, n_c, \quad (v_d)_i \sim \mathcal{N}(0, 1), \quad i = 1, \dots, n_d.$$

* Corresponding author.

E-mail addresses: azymnis@stanford.edu (A. Zymnis), boyd@stanford.edu (S. Boyd), gorin@stanford.edu (D. Gorinevsky).

(We can assume the discrete noise components have unit variance without loss of generality, using a positive diagonal scaling.)

The model is defined by the matrices A^{cc} , A^{dc} , A^{cd} , and A^{dd} , the intercept terms b^c and b^d , and the continuous noise variances σ_i^2 , $i = 1, \dots, n_c$.

In this paper we use the terms discrete and Boolean to mean the same thing. But discrete variables, that take on a finite number of values (more than two) are also readily handled in our model, by representing each discrete variable as a set of Boolean variables.

The model (1)–(2) includes several well known and widely used special cases. For $m_d = 0$ and $n_d = 0$ we have a simple linear model with additive Gaussian noise. For $n_c = 0$ and $n_d = 1$, we obtain a probit model [1]. We mention several applications in Section 1.4.

In this paper we address two basic problems associated with this model: estimation and identification.

1.2. Estimation

We first look at the problem of estimating the model inputs x and d , given one or more output samples. The prior distribution on the inputs is specified by a density $p(x)$ for x , which we assume is log-concave, and the probability p_j that $d_j = 1$. (We assume that x and all d_j are independent.) We will see that the maximum a posteriori (MAP) estimate of (x, d) is the solution of a mixed Boolean convex problem. If n_d is small enough, this problem can be solved exactly, by exhaustive enumeration of the possible values of d , or by a branch-and-bound or other global optimization method. For other cases, we propose to solve the optimization problem approximately, by solving a convex relaxation, and rounding the result, possibly followed by a local optimization step. We refer to the resulting estimate as the RMAP (‘relaxed MAP’) estimate. Numerical simulation suggests that the estimation performance of the RMAP estimator is quite similar to the MAP estimate; unlike the MAP estimate, however, it is computationally tractable even for very large problems [4].

1.3. Identification

We then look at the dual problem of fitting a model of the form (1)–(2), i.e., determining values for the model parameters, given a given set of (training) data samples

$$(x^{(1)}, d^{(1)}, y^{(1)}, z^{(1)}), \dots, (x^{(K)}, d^{(K)}, y^{(K)}, z^{(K)}).$$

We show that the associated log-likelihood function is a concave function of the model parameters, so maximum likelihood (ML) model fitting reduces to solving a convex optimization problem. To obtain a parsimonious model, i.e., one in which many of the entries of the parameter matrices are zero, we propose ℓ_1 -regularized ML fitting (which is also a convex problem). By varying a regularization parameter, we can trade off model fit and model parsimony [5–8].

1.4. Prior and related work

The model that we consider is in essence a hybrid generalized linear model (GLM). These models have been extensively used for explanatory modelling. Some good references on GLMs are [1–3]. There is a considerable amount of research that deals with special cases of our problem. For example, when $m_c = 0$ and $n_c = 0$, we get a model which is very similar to a digraph model commonly used in diagnostics, e.g., see [9,10]. The formulation in this paper is an extension of the earlier work of the authors [4,11,12]. In [4] we considered a special case of the current problem, in the context of fault identification. The paper [11] considers a dynamic system with continuous and discrete states and continuous outputs. The upcoming paper [12] considers a special case of sparse parametric inputs and discrete outputs, where the goal is to compute the sparsity pattern of the inputs.

2. Estimation

In this section we consider the problem of estimating the most probable values of (x, d) , given measurements (y, z) . We first derive the log posterior probability of (x, d) given (y, z) and show that it is jointly concave in x and d (with d relaxed to take values in $[0, 1]$). Thus the problem of estimating the maximum a posteriori (MAP) estimate of x and d is a mixed integer convex problem, i.e., a convex optimization problem with the additional constraint that some of the variables take values in $\{0, 1\}$. We then present an efficient heuristic for solving this problem approximately, based on a convex relaxation of the combinatorial MAP problem. Our analysis follows closely the previous work of the authors in fault detection [4].

The method that we present is readily extended to the case when we have multiple measurements $(y^{(1)}, z^{(1)}), \dots, (y^{(M)}, z^{(M)})$ for the same input (x, d) : we just have to stack all these measurements and augment the system model equations appropriately.

2.1. Maximum a posteriori estimation

2.1.1. Log posterior

From Bayes’ rule the posterior probability (density) of x and d given y and z is

$$p(x, d|y, z) \propto \prod_{i=1}^{m_c} p(y_i|x, d) \prod_{i=1}^{m_d} p(z_i|x, d)p(d)p(x).$$

Taking logarithms, we have $\log p(x, d|y, z) = l(x, d) + C$, where C is a constant and

$$l(x, d) = l_{mc}(x, d) + l_{md}(x, d) + l_{pd}(d) + l_{pc}(x), \tag{3}$$

with the terms described below. The posterior contribution due to the continuous measurements is

$$l_{mc}(x, d) = - \sum_{i=1}^{m_c} \frac{1}{2\sigma_i^2} (y - A^{cc}x - A^{dc}d - b^c)_i^2.$$




Fig. 1. Plot of $\log \Phi(x)$ versus x .

The posterior contribution due to the discrete measurements is

$$l_{\text{md}}(x, d) = \sum_{i=1}^{m_d} z_i \log \Phi(A^{\text{cd}}x + A^{\text{dd}}d + b^{\text{d}})_i + \sum_{i=1}^{m_d} (1 - z_i) \log \Phi(-A^{\text{cd}}x - A^{\text{dd}}d - b^{\text{d}})_i,$$

where Φ is the cumulative distribution of a standard Gaussian. The discrete variable prior term is

$$l_{\text{pd}}(d) = \lambda^T d,$$

where $\lambda_j = \log(p_j/(1 - p_j))$. The continuous variable prior term is

$$l_{\text{pc}}(x) = \log p(x).$$

The log posterior (3) is jointly concave in x and d , when d is relaxed to take values in $[0, 1]^{n_d}$. Indeed, each of the terms is concave in x and d : l_{mc} is a concave quadratic in x and d , l_{pd} is a linear function of d , and $\log p(x)$ is concave by assumption. Concavity of l_{md} follows from log-concavity of Φ ; see, e.g., [13, Section 3.5.2]. Fig. 1 shows a plot of $\log \Phi(x)$ versus x for x ranging between -5 and 5 .

2.1.2. MAP estimation

The problem of estimating the most probable input variables x and d , given the outputs y and z , can be cast as the following optimization problem:

$$\begin{aligned} & \text{maximize} && l(x, d) \\ & \text{subject to} && d \in \{0, 1\}^{n_d}, \end{aligned} \quad (4)$$

with variables $x \in \mathbf{R}^{n_c}$ and $d \in \{0, 1\}^{n_d}$. This is a mixed integer convex problem. One straightforward method for solving it is to enumerate all 2^{n_d} possible values for d , and to find the optimum x in each case, which is tractable, since for fixed d the problem is convex in x [13]. This approach is not practical for n_d larger than around 15 or so, or smaller, if the other dimensions are large. A branch and bound method, or other global optimization technique, can be used to (possibly) speed up computation of the globally optimal solution [14,15]. But the worst case

complexity of any (currently known) method that computes the global solution is exponential in n_d . For this reason we need to consider heuristics for solving problem (4) approximately.

2.2. Relaxed MAP estimation

In this section we describe a heuristic for approximately solving the MAP problem (4). Our heuristic is based on replacing the hard constraints $d_j \in \{0, 1\}$ with soft constraints $d_j \in [0, 1]$. This results in a convex optimization problem that we can solve efficiently. We follow this by rounding and (possibly) a simple local optimization method to improve our estimate.

2.2.1. Linear relaxation

We relax problem (4) to the following optimization problem:

$$\begin{aligned} & \text{maximize} && l(x, d) \\ & \text{subject to} && 0 \leq d \leq 1, \end{aligned} \quad (5)$$

with variables $x \in \mathbf{R}^{n_c}$ and $d \in \mathbf{R}^{n_d}$. This is a convex optimization problem, since it involves maximizing a concave function over a convex set. We can efficiently solve this problem in many ways, e.g., via interior-point methods [16,13]. The complexity of such methods can be shown to be cubic in $n_c + n_d$ (assuming a fixed number of iterations of an interior-point method). Since the feasible set for the relaxed MAP problem (5) contains the feasible set for the MAP problem (4), the optimal value of the relaxed MAP problem, which we denote l_{ub} , gives an upper bound on the optimal value of the MAP problem.

Let (x^*, d^*) be an optimal point for the relaxed MAP problem (5), so we have $l_{\text{ub}} = l(x^*, d^*) \geq l(x^*, d)$ for any Boolean d . If d^* is also Boolean, i.e., $d_j^* \in \{0, 1\}$ for all j , we conclude that d^* is in fact optimal for the MAP problem. In other words, when a solution to the relaxed MAP problem turns out to have Boolean entries, it is optimal for the MAP problem. In general, of course, this does not happen; at least some values of d_j^* will lie between 0 and 1.

2.2.2. Rounding

Let (x^*, d^*) denote the optimal point of the problem (5). We refer to d^* as a *soft decision*, since its components can be strictly between 0 and 1. The next step is to round the soft decision d^* to obtain a valid Boolean solution (or *hard decision*) for d . Let $\theta \in (0, 1)$ and set

$$\hat{d} = \text{pos}(d^* - \theta).$$

To create \hat{d} , we simply round all entries of d_j^* smaller than the threshold θ to zero. Thus θ is a threshold for guessing that a discrete input variable is 1, based on the relaxed MAP solution d^* . As θ varies from 0 to 1, this method generates up to n different estimates \hat{d} , as each entry in d falls below the threshold. We can efficiently find them all by sorting the entries of d^* , and setting the values of \hat{d}_j to one in the order of increasing d_j^* .

We evaluate the log-posterior for each of these (or a subset) by solving the optimization problem

$$\text{maximize } l(x, \hat{d}), \quad (6)$$

with variables $x \in \mathbf{R}^{n_c}$. This is an unconstrained convex problem that can also be solved efficiently, again with cubic complexity in n_c . The RMAP continuous variable estimate \hat{x} is obtained as the minimizer of (6) corresponding to the best obtained estimate \hat{d} .

2.2.3. Local optimization

Further improvement in our estimate can sometimes be obtained by a local optimization method. We describe here the simplest possible such method. We initialize \hat{d} as the one which results in the largest value of $l(x, d)$ after rounding. We then cycle through $j = 1, \dots, n$, at step j replacing \hat{d}_j with $1 - \hat{d}_j$. If this leads to an increase in the optimal value of problem (6), we accept the change and continue. If (as usually is the case) flipping the j th bit results in a decrease in l , we go on to the next index. We continue until we have rejected changes in all entries in \hat{d} . (At this point we can be sure that \hat{d} is 1-OPT, which means that no change in one entry will improve the loss function.) Numerical experiments show that this local optimization method often has no effect, which means that the rounded solution is 1-OPT. In some cases, however, it can lead to a modest increase in l .

2.2.4. Performance of RMAP

The performance of our estimate of (x, d) should be judged by (for example) the mean-square error in estimating x , and the probability of making errors in estimating d (which could be further broken down into false positive and false negative error rates). Numerical examples show that RMAP has very similar performance as MAP, but has the advantage of tractability. This can be partially explained as follows. When the estimation problem is ‘hard’, for example, when the noise levels are high, no estimation method (and in particular, neither MAP nor RMAP) can do a good job at estimating x and d . When the estimation problem is ‘easy’, for example, when the noise levels are low, even simple estimation methods (including RMAP) can do a good job at estimating x and d . So it is only problems in between ‘hard’ and ‘easy’ where we could possibly see a significant difference in estimation performance between MAP and RMAP. In this region, however, we observe from numerical experiments that MAP and RMAP achieve very similar performance.

3. Identification

In Section 2 we addressed the problem of estimating (x, d) given measurements (y, z) when the system model is known. In this section we look at the dual problem of fitting a model of the form (1)–(2) to given data, assuming that the continuous noise variances are known. We show that the log likelihood of the model parameters given the measurements is a concave function, so we can solve the maximum likelihood (ML) problem efficiently using convex optimization.

We first observe that since all measurements are independent of each other, we can separately identify the model parameters that correspond to each measurement. The complexity of the resulting ML identification

technique is thus linear in the total number of measurements. (It is also trivially parallelizable, since the parameters corresponding to each component can be estimated separately.)

Finally, we present a simple technique, variously known as compressed sensing [5–7], the Lasso [17,18], sparse signal recovery [19], and basis pursuit [20] that can be used to identify parsimonious models that fit the data well. This involves using the ℓ_1 -norm of the parameter vector of a given linear model as a surrogate for the model sparsity.

3.1. Continuous parameter identification

Suppose that we are given samples of the form $(x^{(j)}, d^{(j)}, y^{(j)})$ for $j = 1, \dots, K$. Let a_i^{cc}, a_i^{dc} denote the i th row of A^{cc}, A^{dc} , respectively. Since the continuous noise terms v_c are Gaussian, the ML estimate of $(a_i^{cc}, a_i^{dc}, b_i^c)$ given the data, is the one that maximizes

$$l_c(a_i^{cc}, a_i^{dc}, b_i^c) = - \sum_{j=1}^K (y_i^{(j)} - a_i^{ccT} x^{(j)} - a_i^{dcT} d^{(j)} - b_i^c)^2.$$

We can evaluate the maximum likelihood estimates of these model parameters efficiently (via least squares), if σ_i is known. In order to estimate a parsimonious model, as well as σ_i , we propose solving the following problem:

$$\text{maximize } l_c(a_i^{cc}, a_i^{dc}, b_i^c) - \mu_i^c (\|a_i^{cc}\|_1 + \|a_i^{dc}\|_1), \quad (7)$$

which is a convex optimization problem, for a fixed parameter $\mu_i^c > 0$. Having solved problem (7), we can then estimate the noise variance σ_i^2 as the variance of the measurement residuals $(y_i^{(j)} - a_i^{ccT} x^{(j)} - a_i^{dcT} d^{(j)} - b_i^c)$.

3.2. Discrete parameter identification

Now suppose that we are given samples of the form $(x^{(j)}, d^{(j)}, z^{(j)})$ for $j = 1, \dots, K$. Let a_i^{cd}, a_i^{dd} denote the i th row of A^{cd}, A^{dd} , respectively. We have that the log likelihood of $(a_i^{cd}, a_i^{dd}, b_i^d)$ given the data is

$$l_d(a_i^{cd}, a_i^{dd}, b_i^d) = \sum_{j=1}^K z_i^{(j)} \log \Phi(a_i^{cdT} x^{(j)} + a_i^{ddT} d^{(j)} + b_i^d) + (1 - z_i^{(j)}) \log \Phi(-a_i^{cdT} x^{(j)} - a_i^{ddT} d^{(j)} - b_i^d),$$

which is a concave function of $(a_i^{cd}, a_i^{dd}, b_i^d)$. We can thus efficiently evaluate the ML estimates of the discrete model parameters efficiently, for example using Newton’s method. This is equivalent to solving a probit regression problem. In order to estimate a parsimonious model we propose solving the following problem:

$$\text{maximize } l_d(a_i^{cd}, a_i^{dd}, b_i^d) - \mu_i^d (\|a_i^{cd}\|_1 + \|a_i^{dd}\|_1), \quad (8)$$

which is a convex optimization problem, for a fixed parameter $\mu_i^d > 0$.

3.3. Computational complexity

Problems (7) and (8) can be solved efficiently in a variety of methods such as interior-point methods (in a way similar to [21,8]) and first order methods, e.g.,

[22]. In any case the complexity of solving this problem is cubic in $n_c + n_d$ and linear in K . Thus the overall complexity of ℓ_1 -regularized ML system identification is $O(K(m_c + m_d)(n_c + n_d)^3)$.

3.4. Choice of regularization parameter

The regularization parameters μ_i^c and μ_i^d control the tradeoff between data fit, as measured by $l(A, b)$, where

$$l(A, b) = \sum_{i=1}^{n_c} l_c(a_i^{cc}, a_i^{dc}, b_i^c) + \sum_{i=1}^{n_d} l_d(a_i^{cd}, a_i^{dd}, b_i^d) \quad (9)$$

and model sparsity, as measured by the ℓ_1 norm. In order to keep our modelling procedure as flexible as possible, we use a different regularization parameter for each continuous and discrete sensor.

We use cross validation to choose each of the regularization parameters μ_i^c and μ_i^d . We divide our available data into a training set and a test set. For each continuous sensor $i = 1, \dots, m_c$ and for each value of μ_i^c we solve problem (7) for the training set and measure the average square residual of measurement i on the test set. We then choose as μ_i^c the value of μ that gives the smallest residual. We repeat this process for the discrete sensors, where instead of square residual we use the average error rate in the discrete sensor measurement.

This is the simplest possible way of fitting the regularization parameters. There are various other methods, such as minimizing the Akaike information criterion (AIC), or minimizing the Bayesian information criterion (BIC). For a detailed description of these methods see [3, Chapter 7].

4. Numerical example

In this section we present the results of applying our estimation and identification methods to a small artificially generated example. Our experiment proceeds as follows. We first generate three different data sets, a training set, a validation set, and a test set, according to a pre-defined mixed linear system. We use the training and validation sets to perform system identification using the method described in Section 3. We then use this identified model to predict the inputs in the test set, given the outputs. The prediction performance on the test set shows the real value of the model identification and estimation methods, since here we are making predictions based on previously unseen data.

4.1. Data specification

We consider a system with dimensions $n_c = n_d = 10$ and $m_c = 18$ and $m_d = 16$. We first choose a random sparsity pattern for each of the matrices A^{cc} , A^{dc} , A^{cd} , and A^{dd} , with 10% of the entries nonzero. We draw the nonzero entries of A^{cc} , A^{dc} , A^{cd} , and A^{dd} from an $\mathcal{N}(0, 1)$ distribution. The entries of b^c and b^d are drawn from an $\mathcal{N}(0, 1)$ and an $\mathcal{N}(0, 100)$ distribution, respectively. We set $\sigma_i = 0.01$ for all i . Each element of input x is generated

according to an $\mathcal{N}(0, 1)$ prior. The prior probability of d_j being equal to 1 is 0.2 for all j .

4.2. System identification

We generate a training set and validation set, each consisting of 500 data points.

We look at 30 candidate values of μ , logarithmically spaced in the range $(10^{-2}, 10^3)$. For each value of μ , we use the method described in Section 3 to fit a mixed linear system to the given training set. For each continuous sensor, we choose the value of μ which achieves the minimum relative root mean square error r_y on the validation set, defined as

$$r_y = \frac{\|y - \hat{y}\|_2}{\|y\|_2},$$

where \hat{y} is the output predicted by the identified model, given the true input. For each discrete sensor we choose the value of μ that achieves the minimum average error rate on e_z on the predicted value of z for the validation set. As an example, Fig. 2 shows the plot of e_z versus μ for discrete sensor 5.

4.3. Estimation performance

We then generate a test set consisting of another 500 data points. We use this set to judge the estimation performance of the mixed linear system identified in Section 4.2. To do this, we predict the inputs (x, d) for the outputs (y, z) for the test set, using the relaxed MAP method described in Section 2. We compute the error rate in our estimate of d (which we call e_d), as well as the average relative root-mean-square (RMS) error between our estimate and the true value of x , defined as

$$r_x = \frac{\|x - \hat{x}\|_2}{\|x\|_2}.$$

Our identified model yields an error rate e_d of about 2.2×10^{-3} on this data and a relative RMS error in x of about 4.7×10^{-2} . For comparison, if we use the true model for estimation, we obtain an error rate e_d of about 2.0×10^{-3}




Fig. 2. Plot of e_z as a function of μ on the validation set for discrete sensor 5.

Table 1

Modelling performance of identified model.

	Train	Validation	Test
r_y	0.0056	0.0058	0.0057
e_z	0.0063	0.0080	0.0110

Table 2

Estimation performance of identified model.

	Train	Validation	Test
r_x	0.0500	0.0470	0.0460
e_d	0.0052	0.0048	0.0022

and a relative RMS error in x of about 4.1×10^{-2} . We thus see that our system identification method does a reasonable job at fitting a model to the given data, at least for the purposes of estimating the input.

The identification and estimation performance of our identified model is shown in Tables 1 and 2, respectively. We judge the identification performance of the model, based on how well this model can predict the outputs from the inputs. As we can see from Table 1, the model does a good job at this task, given that the values of r_y and e_z for the test set are of the same order of magnitude as the same values for the training and validation sets. Furthermore, from Table 2, we see that the model also generalizes well for estimation. The values of r_x and e_d for the test set are of the same order as the ones for the training and validation sets.

5. Conclusions

We have introduced a class of mixed linear systems that includes the standard linear model and the probit model as special cases. We have presented a simple heuristic for estimating the input given the output of such a system based on a convex relaxation of a combinatorial MAP problem. We have also presented a simple heuristic that uses ℓ_1 -regularized ML to identify such a model given a set of data points. We have shown that this method performs well in practice through a numerical example.

References

- [1] P. McCullagh, J. Nelder, Generalized Linear Models, Chapman & Hall, London, 1989.
- [2] T. Hastie, R. Tibshirani, Generalized Additive Models, Chapman & Hall, 1990.
- [3] T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Springer, Berlin, 2001.
- [4] A. Zymnis, S. Boyd, D. Gorinevsky, Relaxed maximum a posteriori fault identification, Signal Processing 89 (6) (2009) 989–999.
- [5] D. Donoho, Compressed sensing, IEEE Transactions on Information Theory 52 (4) (2006) 1289–1306.
- [6] E. Candès, J. Romberg, T. Tao, Stable signal recovery from incomplete and inaccurate measurements, Communications on Pure and Applied Mathematics 59 (8) (2005) 1207–1223.
- [7] E. Candès, J. Romberg, T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, IEEE Transactions on Information Theory 52 (2) (2006) 489–509.
- [8] K. Koh, S.-J. Kim, S. Boyd, An interior point method for large-scale ℓ_1 -regularized logistic regression, Journal of Machine Learning Research 8 (2007) 1519–1555.
- [9] I. Sacks, Digraph matrix analysis, IEEE Transactions on Reliability 34 (5) (1985) 437–446.
- [10] S. Deb, K. Pattipati, V. Raghavan, M. Shakeri, R. Shrestha, Multi-signal flow graphs: a novel approach for system testability analysis and fault diagnosis, IEEE Aerospace and Electronic Systems Magazine 10 (5) (1995) 14–25.
- [11] A. Zymnis, S. Boyd, D. Gorinevsky, Mixed state estimation for a linear Gaussian Markov model, in: 47th IEEE Conference on Decision and Control, 2008, pp. 3219–3226.
- [12] A. Zymnis, S. Boyd, E. Candès, Compressed sensing with quantized measurements, 2009, available online at <www.stanford.edu/~boyd/quant_compr_sens.html>.
- [13] S. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, Cambridge, 2004.
- [14] E. Lawler, D. Wood, Branch-and-bound methods: a survey, Operations Research 14 (1966) 699–719.
- [15] R. Moore, Global optimization to prescribed accuracy, Computers and Mathematics with Applications 21 (6–7) (1991) 25–39.
- [16] J. Nocedal, S. Wright, Numerical Optimization, Springer, Berlin, 1999.
- [17] R. Tibshirani, Regression shrinkage and selection via the Lasso, Journal of the Royal Statistical Society 58 (1) (1996) 267–288.
- [18] B. Efron, T. Hastie, I. Johnstone, R. Tibshirani, Least angle regression, Annals of statistics (2004) 407–451.
- [19] J. Tropp, Just relax: convex programming methods for identifying sparse signals in noise, IEEE Transactions on Information Theory 52 (3) (2006) 1030–1051.
- [20] S. Chen, D. Donoho, M. Saunders, Atomic decomposition by basis pursuit, SIAM Review (2001) 129–159.
- [21] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, D. Gorinevsky, An interior-point method for large-scale ℓ_1 -regularized least squares, IEEE Journal of Selected Topics in Signal Processing 1 (4) (2007) 606–617.
- [22] E. Hale, W. Yin, Y. Zhang, Fixed-point continuation for ℓ_1 -minimization: methodology and convergence, SIAM Journal on Optimization 19 (2008) 1107–1130.