Magnitude Least–Squares Fitting via Semidefinite Programming with Applications to Beamforming and Multidimensional Filter Design

Peter Kassakian

Dept. of Electrical Engineering and Computer Science
The Center for New Music and Audio Technologies
University of California, Berkeley
Berkeley, California 94709, USA

ABSTRACT

The standard least–squares problem seeks to find a linear combination of columns of a given matrix that best approximates a target vector in Euclidean norm. The problem of finding a linear combination of columns, the component-wise magnitude of which approximates a target, is not a convex problem, but can be well–approximated using semidefinite programming. High quality solutions can be found by reformulating the problem as a generalization of a graph partitioning problem, relaxing a rank constraint, and rounding back onto the feasible set. A bound on the gap between the objectives of the global optimum and the approximate solution can be calculated for instances of the problem, and for many practical problems can be quite small. The problem is shown to have application in array pattern synthesis, multidimensional filtering, and spectral factorization.

1. INTRODUCTION

Motivated by filter design problems where the important characteristic of the filter is its magnitude frequency response, we investigate a related general optimization problem,

$$\min_{x\in\mathbb{C}^n} \|\, |Ax| - b\,\|_2^2, \quad A\in\mathbb{C}^{m\times n}, \quad b\in\mathbb{R}_+^m. \qquad (1)$$

The problem is not convex, and is difficult to solve directly as formulated, but by reformulating it as a type of two–way partitioning problem [1], and performing a semidefinite relaxation, we can nonetheless obtain high quality solutions. The problem naturally arises in multidimensional filtering and array pattern synthesis problems where targets are specified in terms of ideal magnitude responses (disregarding phase).

In the case of one–dimensional filters and uniformly spaced line arrays, the difficulty of the problem is greatly reduced because of the guaranteed existence of one–dimensional spectral factorizations of magnitude responses [2] [3] [4]. This allows the designer to formulate a variety of one–dimensional filter problems in terms of autocorrelation coefficients, which are linearly related to the magnitude response. For multidimensional filters and nonuniformly spaced arrays, the problem is, in general, much more difficult.

Related recent work on two–dimensional filtering and array pattern synthesis with semidefinite programming can be found in [5] and [6]. The author in [5] uses a semidefinite relaxation for the purposes of obtaining discrete coefficient solutions, and the authors in [6] identify an objective function similar to (1) arising in array pattern synthesis, and approximately solve the problem using iteration. Our strategy in overcoming the convexity difficulties associated with (1) is to reformulate, relax through duality, and round the solution onto the original feasible set.

The author would like to thank Meyer Sound Laboratories Inc. for continued collaboration and research support.

2. FORMULATION

In their breakthrough paper [7], Goemans and Williamson formulate a method for approximately solving an NP-complete graph partitioning problem called MAXCUT using semidefinite programming. The authors prove that the method produces suboptimal solutions that are provably close in performance to the true global optimum. This and other successful employments of the technique have motivated us to cast our problem into a form that resembles the two–way partitioning problem.

There are two key observations that lead to the formulation. The first is that we can eliminate the absolute value bars by introducing an $m\times 1$ vector of extra unity modulus variables, and rewriting (1) as

$$\begin{aligned} \underset{x\in\mathbb{C}^n,\, c\in\mathbb{C}^m}{\text{minimize}} \quad & \|Ax - Bc\|_2^2 \\ \text{subject to} \quad & |c| = \mathbf{1}, \\ \text{where} \quad & B = \mathbf{diag}(b) \end{aligned} \qquad (2)$$

For fixed x, the ith component c_i of the optimal c in (2) is $\exp(j\mathrm{Arg}(A_i x))$. The vector c represents freedom in the problem to choose a complex target Bc that matches Ax in componentwise polar angle. Formulation (2) is equivalent to (1), transferring difficulties associated with the objective function (the absolute value operator) into the constraint $|c| = \mathbf{1}$ (nonlinear equality constraint).

The second key observation is that we can first minimize over x and then minimize over c. In other words, problem (2) can be rewritten as

$$\min_{|c|=\mathbf{1}} \ (\min_{x\in\mathbb{C}^n} \ \|Ax - Bc\|_2^2) \tag{3}$$

The subproblem in parentheses has an analytic solution in terms of c. Assuming A has full rank, we have

$$\begin{aligned} \underset{c\in\mathbb{C}^m}{\text{minimize}} \quad & \|A(A^HA)^{-1}A^HBc - Bc\|_2^2 \\ \text{subject to} \quad & |c| = \mathbf{1}, \end{aligned} \tag{4}$$

Letting $U = A(A^HA)^{-1}A^H - I$, and $W = (UB)^H(UB)$, we can convert the problem into a "complex two–way partitioning problem",

$$\begin{aligned} \underset{c\in\mathbb{C}^m}{\text{minimize}} \quad & c^HWc \\ \text{subject to} \quad & |c| = \mathbf{1}, \end{aligned} \tag{5}$$

3. RELATION TO REAL VALUED PARTITIONING PROBLEM

In order to solve the nonconvex problem (5), we will perform a semidefinite relaxation and round the solution onto the feasible set of the original problem. First, though, it is illustrative to relate the complex problem (5) to the real–valued two–way partitioning problem as described in [1].

First consider the objective function. The matrix W being conjugate symmetric, we can express c^HWc as

$$c^HWc = \mathbf{Tr}\left((c_Rc_R^T + c_Ic_I^T)W_R + (c_Rc_I^T - c_Ic_R^T)W_I\right), \tag{6}$$

where c_R and c_I denote real and imaginary parts of c, W_R and W_I denote the real and imaginary parts of W, and $\mathbf{Tr}()$ is the trace operator.

Regarding $|c| = \mathbf{1}$, the following holds:

$$|c| = \mathbf{1} \iff \mathbf{diag}(c_Rc_R^T + c_Ic_I^T) = \mathbf{1}, \tag{7}$$

where, in this context, the function $\mathbf{diag}()$ vectorizes the diagonal elements. Combining, we can express the original complex problem as a real SDP with a rank constraint.

$$\begin{aligned} \underset{C\in\mathbb{R}^{2m\times 2m}}{\text{minimize}} \quad & \mathbf{Tr}(C\tilde{W}) \\ \text{subject to} \quad & C_{ii} + C_{i+m,i+m} = 1, \quad \forall i = 1, ..., m \\ & C \succeq 0 \\ & \mathbf{rank}(C) = 1 \\ \text{where} \quad & \tilde{W} = \begin{bmatrix} W_R & -W_I \\ W_I & W_R \end{bmatrix} \end{aligned} \tag{8}$$

This is very similar to the real–valued two–way partitioning problem statement in [1], differing only in the form of the diagonal equality constraints. In (8), the first constraint states that pairs of diagonal elements of C must sum to one, whereas in the standard two–way partitioning problem, all the diagonal elements are forced to unity.

4. RELAXATION

The semidefinite relaxation is found by ignoring the rank constraint. The rank–relaxed problem is exactly the dual of the associated (always convex) dual problem. The following relaxations are equivalent:

Complex SDP relaxation:

$$\begin{aligned} \underset{\mathscr{C}\in\mathbb{C}^{m\times m}}{\text{minimize}} \quad & \mathbf{Tr}(\mathscr{C}W) \\ \text{subject to} \quad & \mathscr{C}_{ii} = 1, \quad \forall i = 1, ..., m \\ & \mathscr{C} \succeq 0 \end{aligned} \tag{9}$$

Real SDP Relaxation:

$$\begin{aligned} \underset{C\in\mathbb{R}^{2m\times 2m}}{\text{minimize}} \quad & \mathbf{Tr}(C\tilde{W}) \\ \text{subject to} \quad & C_{ii} + C_{i+m,i+m} = 1, \quad \forall i = 1, ..., m \\ & C \succeq 0 \end{aligned} \tag{10}$$

Depending on the algorithm used to solve the semidefinite programs above, one of the formulations may be advantageous. For example, the complex formulation may be more efficient for solvers that can natively handle complex numbers.

5. FEASIBLE SOLUTION

The optimal solution to the relaxation is a positive semidefinite matrix. Often it is of low rank, but it's unlikely to be rank 1, which would be required for the solution to the relaxation to correspond exactly to the solution of the original nonconvex problem (5). We must find a rank 1 matrix that approximates the optimal matrix for the relaxation in order to obtain a useful solution. This can be done in a variety of ways, as described in [8] [9]. In the examples that follow, we draw samples $\{[\bar{s}_i^R; \bar{s}_i{}^I]\}_{i=1,...,l}$ from a random multivariate distribution with covariance C_{opt}, the matrix optimizing (10). Candidate feasible vectors are then created as $s_i = (\bar{s}_i{}^R + j\bar{s}_i{}^I)/|\bar{s}_i{}^R + j\bar{s}_i{}^I|$. The rank 1 surrogate for C_{opt} is then taken to be $s_{\text{opt}}s_{\text{opt}}^H$, where s_{opt} minimizes $s_i^HWs_i$ over all the samples $\{s_i\}_{i=1,...,l}$.

An alternative scheme is to take s_{opt} to be the eigenvector corresponding to the largest eigenvalue of the matrix C_{opt}. Finally, the approximate solution to Problem (1) can be calculated as $x_{\text{opt}} = (A^HA)^{-1}A^HBs_{\text{opt}}$.

Since the vector s_{opt} is feasible for the unrelaxed problem (5), the objective $s_{\text{opt}}^H W s_{\text{opt}}$ is an upper bound for the globally optimal objective of (5). The ϵ–optimal objective of the relaxation provides us with a lower bound. The gap thus bounds the suboptimality of the solution, and can be evaluated as a byproduct of the method for specific problem instances.

6. APPLICATIONS

The problem, as stated in its original form, can easily be seen to have many uses in engineering and elsewhere. For example, the complex system Ax could represent samples of superimposed waves created by n sources. Often it is important to manage the amplitude of the resultant waveforms $|Ax|$. Linear filter design is a primary example, including multidimensional filtering and beamforming. It is sometimes the case that the magnitude frequency response of the filter is the most important attribute of the design.

6.1. Magnitude beamforming with arbitrary array

Here we consider optimally choosing the weights for a beamformer to achieve a response close in magnitude to a target. In general we do not have to assume anything special about the geometry of the array or the uniformity of the array element responses – for the example, we will consider a twelve element linear array, but with random spacing between the elements. By sampling the far–field responses of each array element individually, and the desired response pattern across angle, we can construct a complex matrix A and positive–valued real vector b for use in our formulation. In this specific problem instance, the globally optimal objective can be no less than 99.6% of the objective achieved for the approximate solution. Solving the corresponding standard least–squares problem $\min \|Ax - b\|$ with hopes that the solution will be close to the optimal with respect to magnitude difference yields a solution that is 75 times worse. Access to guarantees on the degree to which the solution is suboptimal, generated numerically in the course of solving, is an important feature of the method.

6.2. Two-dimensional filtering

We wish to design a two-dimensional filter that has magnitude response close to a target. As in beamforming, there exists a variety of methods available for solving variants of the problem. One advantage of the outlined framework is that it is very general, requiring no special symmetries in the design specifications nor any essential properties of the dimension of the filter. Figure 2 shows the magnitude of a complex two–dimensional filter of size 15×15, designed to have a magnitude response that resembles a pyramid. The target response was sampled on 31×31 equally spaced grid,

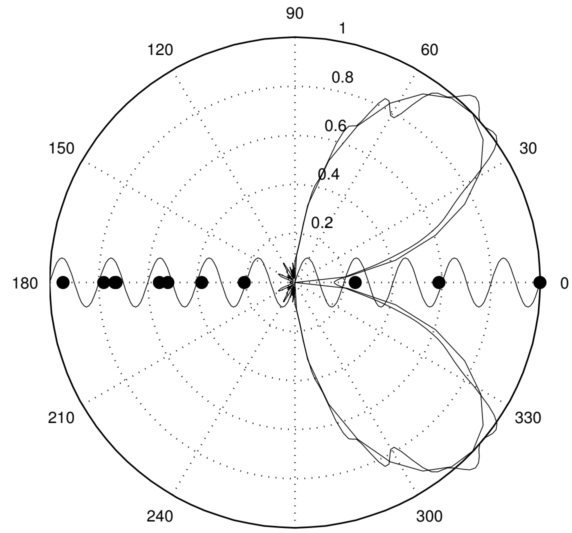

Fig. 1. Magnitude beamforming with arbitrary array geometry. The array element locations are shown as small circles along with a wave (indicating relative wavelength).

making the dimensions of the A matrix 961×225. For this example, the globally optimal objective can be no less than 66% of the objective of the achieved approximation. In this case the bound is weaker, but nevertheless, the achieved pattern matches the target closely. The optimal objective could indeed be very close to the achieved value.

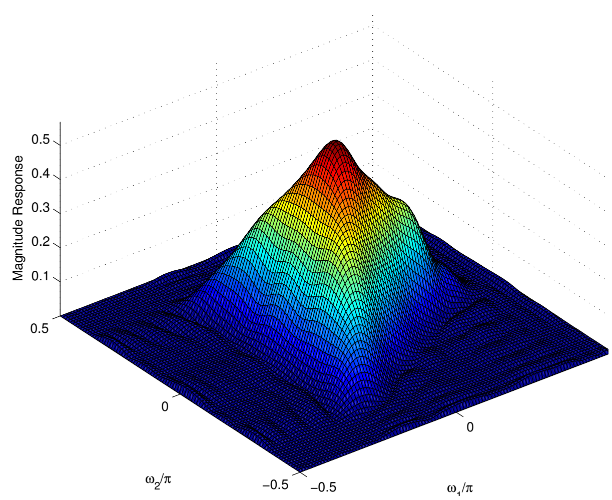

Fig. 2. Magnitude frequency response of 15×15 two–dimensional filter.

7. FURTHER APPLICATIONS AND EXTENSIONS

In this section we will mention a simple extension of problem (1), and a more subtle application of the problem involving approximate spectral factorization.

7.1. Regularization and Weighting

An enabling factor approximating (1) using the semidefinite program (9) is the fact that we have an analytic expression for the solution to the least–squares subproblem in (3). This property also holds for a Tikhonov regularized version of (1):

$$\min_{x\in\mathbb{C}^n} \| \, |Ax| \, - \, b \, \|_2^2 + \delta\|x\|_2^2, \quad \delta \geq 0. \tag{11}$$

Problem (11) can be cast in the form of (5), but with

$$\begin{aligned} W &= P^H P + \delta Q^H Q, \quad \text{where} \\ Q &= (A^H A + \delta I)^{-1} A^H B \quad \text{and} \\ P &= AQ - B, \quad B = \mathbf{diag}(b). \end{aligned} \tag{12}$$

The formulation is useful in controlling the sensitivity of the objective with respect to the coefficients of the solution, and can be used to find locally robust solutions to (1).

Another straightforward extension is that of incorporating linear weightings in the objective function.

7.2. Spectral Factorization

The main problem (1) has close ties to general spectral factorization problems. As mentioned in the introduction, spectral factorization makes designing one dimensional filters with prescribed magnitude frequency responses a relatively easy problem. The ideas described can be leveraged to calculate approximate spectral factorizations of polynomials. Suppose

$$q(\omega_1, \omega_2, ...) = \sum_{\{i,k,...\}=1}^{Q} q[i,k,...] e^{ji\omega_1} e^{jk\omega_2} ... \tag{13}$$

is a multivariate trigonometric polynomial of fixed degree Q that is real and nonnegative. Then we can find a polynomial p of smaller degree $P < Q$, such that pp^* approximates q. This is done simply by sampling q and setting $\sqrt{q}$ as b in the formulation. Organizing the data appropriately, we can write

$$\begin{aligned} p(\omega_1, \omega_2, ...)p^*(\omega_1, \omega_2, ...) &= |p(\omega_1, \omega_2, ...)|^2 \\ &= \left| \sum_{\{i,k,...\}=1}^{P} p[i,k,...] e^{ji\omega_1} e^{jk\omega_2} ... \right|^2 . \end{aligned} \tag{14}$$

We can rewrite a sampled version of (14) as $|A\vec{p}|^2$, where $\vec{p}$ is a rearrangement of the elements of the multidimensional matrix p into a column vector. The data matrix A corresponds to samples of the trigonometric basis functions. Finding $\vec{p}$ such that $|A\vec{p}| \approx |b|$ thus generates an approximate spectral factor for the polynomial magnitude function.

8. CONCLUSION

We've formulated a semidefinite relaxation of a useful, but generally difficult problem arising in generalized filter design: finding filters with desired magnitude frequency responses. The method possesses features that it is flexible in its applicability and is understandable in terms of other successful approximation techniques arising in graph partitioning. Additionally, bounds on the degree to which the approximation is suboptimal can be calculated easily.

9. REFERENCES

[1] Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[2] S. Wu, S. Boyd, and L. Vandenberghe, "FIR filter design via spectral factorization and convex optimization," *Applied Computational Control, Signal and Communications*, pp. 1–33, 1997.

[3] T. Davidson, Z. Luo, and J. Sturm, "Linear matrix inequality formulation of spectral mask constraints," *IEEE Transactions on Signal Processing*, vol. 50, pp. 2702–2715, November 2002.

[4] Brien Alkire and Lieven Vandenberghe, "Convex optimization problems involving finite autocorrelation sequences," *Mathematical Programming, Series A*, vol. 93, no. 3, pp. 331–359, 2002.

[5] W.-S Lu, "A unified approach for the design of 2-D digital filters via semidefinite programming," *IEEE Transactions on Circuits and Systems*, 2002.

[6] Fan Wang, Venkataramanan Balakrishnan, Philip Yuanping Zhou, Jason Jiangnan Chen, Roy Yang, and Colin Frank, "Optimal array pattern synthesis using semidefinite programming," *IEEE Transactions on Signal Processing*, 2003.

[7] M. Goemans and D. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *Journal of Association for Computing Machinery*, vol. 42, no. 6, pp. 1115–1145, June 1995.

[8] Alexandre d' Aspremont and Stephen Boyd, "Relaxations and randomized methods for nonconvex QCQPs," Tech. Rep., Autumn 2003.

[9] Aharon Ben-Tal and Arkadi Nemirovski, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications (MPS-Siam Series on Optimization)*, Society for Industrial & Applied Math (SIAM), August 2001.