

OPTIMAL POSITIONING IN LOW-DIMENSIONAL CONTROL SPACES USING CONVEX OPTIMIZATION

Peter Kassakian & David Wessel

Center for New Music and Audio Technologies (CNMAT)

University of California Berkeley

{kassak, wessel}@cnmat.berkeley.edu

ABSTRACT

A music analysis and control system for use in live performance is presented and demonstrated. Musical features, such as harmony, rhythmic patterning, or melodic structure, are extracted and automatically placed at appropriate locations in a control space. The control space is of low dimensionality, usually only two dimensions, wherein perceptual dissimilarity is represented as distance. The system can be viewed as a listening assistant that aids a performing computer musician with usable information about the evolving musical context so that the material under control can be adjusted to respond in interesting ways to that context. The method developed to position the points is based on a semidefinite programming relaxation of a nonconvex positioning problem and is both novel and elegant. Another novel feature of our work is an inter-process communication implementation that provides for two-way traffic between Max/MSP and MATLAB in real-time. We demonstrate the system with a harmony space and a melodic process space. The method is robust, rapid, and not plagued by problems with local minima. This successful use of convex optimization suggests that semidefinite programming may have broad application in the computer music field.

1. INTRODUCTION

When performing with others it would seem important to know what they are doing; what rhythms they are playing, what harmonies they are employing, what melodic processes they are gesturing upon, what timbres they use, and what auditory environments they inhabit. We offer a real-time method for mapping machine generated perceptions of such aspects of a musical scene to low dimensional gestural control structures that afford a variety of musical reactions to past and present contexts. Our goal is to develop a rudimentary real-time music analysis system that aids in the mapping of a low-dimensional control structure to both stable and evolving features of the musical scene. A rich variety of control spaces is presented in [8].

The general setup for our analysis problem is as follows. We assume the control space contains a set of fixed points, called anchors, with which we can compare new

musical data. We specify a way to evaluate similarity between the new element and the anchors, and position the element in a manner that optimally reflects the dissimilarities as geometric distances. The result is a representation of the new data as a positioned point in the control space, useful both in providing information, and for subsequent control.

The optimization problem associated with this positioning problem is solved by an interior point method – a type of method for solving convex optimization problems that since the early 1990’s has been an area of intense research interest. The interior point methods have favorable computational complexity, and are very efficient in practice. We present the optimization problem in detail, showing how it’s possible to form a convex approximation to the nonconvex localization problem that reliably finds the global minimizer of our error function. We emphasize the value of convex optimization in computer music applications, where nonconvex and peculiar optimization problems abound.

To demonstrate the analysis-to-controller link we present a two straight-forward examples, one in harmonic analysis and another melodic analysis.

Finally we discuss an implementation of the complete system that links Cycling74’s Max/MSP[1] and the Mathworks MATLAB[7]. We have built a Max/MSP external object that directly links, via dynamic C libraries, to the MATLAB Engine, an interface designed to allow programmers call MATLAB from within other programs written in C. This project represents one of many computer music projects that involve complicated numerical routines that are conveniently handled in MATLAB, but are difficult to implement directly in the Max/MSP environment or in C. In our case, we call an interior point solver SeDuMi[11] from MATLAB, called from within Max/MSP.

2. ALGORITHM

2.1. Introduction and Motivation

In this section we describe the method for positioning points representing newly gathered musical elements. We assume that the musician already has constructed an abstract space, defined by existing points, now held fixed for the purposes of locating an incoming point. The dissimilarity-

ties of the musical element in relation to the fixed elements are evaluated, and thought of as a set of desired Euclidean distances between the new point and the existing points. It is unlikely that the evaluated dissimilarities are consistent in the sense that they unambiguously and perfectly determine the coordinates of the point in the control space. This is especially true if the number of fixed points is large, and the new point represents a unique type of musical element not encountered previously.

It is because of this likely inconsistency that we seek the coordinates that minimize the sum of the squared errors between desired dissimilarities and achieved distances. The important attribute of the evaluated dissimilarities is how they are related relative to each other so we calculate an optimal global rescaling factor, in effect working with distance ratios instead of absolute distances.

2.2. Formulation

We now turn to the mathematical details of the optimization problem and show how a convex relaxation of the problem can be obtained and solved with interior point methods. Interior point methods differ from gradient and Newton methods in that they incorporate global information about the geometric structure of the optimization problem, and find solutions by progressing through the interior of the feasible region. Gradient and Newton methods are local in the sense that they only query the description of the problem to obtain derivatives of the objective function. They require a suitable starting point, can find only local minima, and can be inconsistent in terms of number of iterations required to obtain a precise solution. In contrast, the formulation we present is mathematically informative, centered on duality and convexity. The use of interior point methods in numerically calculating the solution is efficient and consistent. In all problem instances we have encountered and checked, the method has yielded the globally optimal solution, even in cases where there are many local minima.

Let $r_i \in \mathbb{R}^n$, $i = 1, 2, \dots, m$ be the locations of m fixed points in the space. Let $d_i \geq 0$, $i = 1, 2, \dots, m$ be the evaluated dissimilarities between the new element and the fixed elements. We want to find $x^* \in \mathbb{R}^n$, and $\alpha^* \in \mathbb{R}_+$ that solve

$$p(r, d) = \min_{x, \alpha} \sum_{i=1}^m (\|x - r_i\| - \alpha d_i)^2. \quad (1)$$

It will be convenient to express the objective function in a different form. Referring to Figure 1, we see that each term in (1) can be expressed as the optimal value of a subproblem:

$$\begin{aligned} (\|x - r_i\| - \alpha d_i)^2 &= \min_{b_i \in \mathbb{R}^n} \|x - b_i\|^2 \\ \text{s.t. } \|r_i - b_i\|^2 &= \alpha^2 d_i^2. \end{aligned} \quad (2)$$

Combining (1) and (2), and writing in matrix form, we

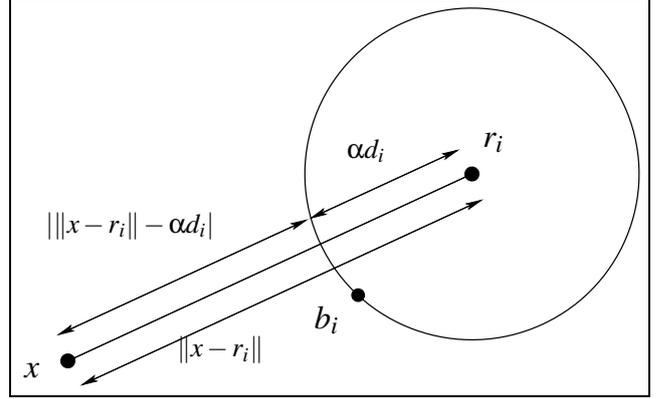


Figure 1. Geometric intuition for the subproblem (2).

have

$$p(r, d) = \min_{x, \alpha, b} \|Jx - b\|^2 \quad (3)$$

$$\text{s.t. } \|r_i - b_i\|^2 = \alpha^2 d_i^2, \quad i = 1, 2, \dots, m,$$

$$\text{where } b \in \mathbb{R}^{mn} \text{ is defined } b \equiv [b_1^T, b_2^T, \dots, b_m^T]^T$$

$$\text{and } J \in \mathbb{R}^{mn \times n} \text{ is defined } J \equiv [I, I, \dots, I]^T.$$

The strategy for solving this nonconvex problem is to relax the constraints via duality. The dual problem that arises through Lagrange multipliers is always convex, and because of the algebraic form of the problem (quadratic), it can be analytically derived, and solved using interior point methods. We will solve the “dual of the dual” which is also convex, and has an interpretation as a relaxation of the primal problem (3). This problem is

$$\min_{Z, z, v} \text{Trace}(WZ) \quad \text{such that} \quad (4a)$$

$$\|r_i\|^2 - v d_i^2 + \sum_{j=1}^n 2(r_i)_j z_{k_{ij}} + Z_{k_{ij}, k_{ij}} = 0, \quad (4b)$$

$$k_{ij} = (i-1)m + j, \quad i = 1, 2, \dots, m,$$

$$\begin{bmatrix} 1 & z^T \\ z & Z \end{bmatrix} \text{ is positive semidefinite,} \quad (4c)$$

$$v \geq 0, \quad (4d)$$

$$\text{where } Z \in \mathcal{S}^{mn \times mn}, z \in \mathbb{R}^{mn}, v \in \mathbb{R}$$

$$\text{and } W \equiv I - J(J^T J)^{-1} J^T.$$

The problem can be solved with semidefinite program solvers, for example SeDuMi, developed by Joe F. Sturm [11]. Because the convex program is a relaxation of (3), the variables do not directly correspond to the variables in our original problem (3), though they are related. In fact, including an additional (nonconvex) constraint $zz^T = Z$ in (4) yields an alternative formulation of the original primal problem (3). Hence the dual relaxation can be directly derived by reformulating (3), and ignoring a constraint of the form $zz^T = Z$. To obtain a primal feasible solution from the relaxation, i.e., a solution that indeed satisfies $zz^T = Z$, we use a randomized method first reported by Goemans and Williamson in [3] and generalized by Y. Nesterov in [10].

A more conventional method for optimizing this objective function, for example a gradient method, could result optimal solutions, especially if time is allocated for repeated solution using several starting points. In our application though, we are pleased with the interior point approach because of its consistent speed and high-quality solutions. Convex optimization and interior point methods form a powerful framework both in terms of mathematical modeling and numerical complexity; for a reference on the subject see [2]. They are promising methods for many computer music applications.

3. EXAMPLES

In this section we consider two simple examples of the idea and solution method. The first uses the key space discussed in [5] as a way to track implied tonality of a soloing musician. We analyze four choruses of Charlie Parker’s “Blues for Alice”, transcribed by Jamie Aebersold and Ken Slone, published in [4]. The second example involves a melodic contour space, where similarity between melodies is evaluated as a function of the type of intervallic motion encountered. This example serves to show that the optimization method finds globally optimal points, even when the objective function possesses several stationary points.

3.1. Key Space Example

Consider the four dimensional key space discussed in [5]. Each of the 24 keys is represented by a point, and in such a way that the distance between any two points reflects the dissimilarity between the keys, as measured by subjective evaluation. The technique used to calculate the locations of these points is called multidimensional scaling. Our problem is related, but simpler. We wish to use the locations of the 24 keys in the space as anchors, and position an arbitrary histogram of tones within this structured space. This would allow us to track tonality in a systematic way.

A normalized histogram of notes is called a tone profile. For each of the 24 keys there exists a canonical tone profile, also described in [5]. We will use the natural Euclidean metric to measure similarity, i.e., if $r_1, r_2 \in \mathbb{R}^{12}$ are tone profiles, then the dissimilarity between them is $d = \|r_1 - r_2\|_2$. Hence, presented with an unfamiliar tone profile, we can calculate its dissimilarity relative to each of the 24 keys, and position it appropriately using the outlined method.

A representation of the space in two dimensions can be found by projecting the points onto a torus, and unwrapping the torus onto a square in \mathbb{R}^2 .

Here we have compiled tone profiles for each of 12 measures in Charlie Parker’s “Blues for Alice”, as transcribed in [4]. The tone profiles were gathered over four choruses of the song including the melody. By plotting their optimal locations in the key space, we can gain insight into the song’s harmonic structure as interpreted by

Parker. See Figure 2.

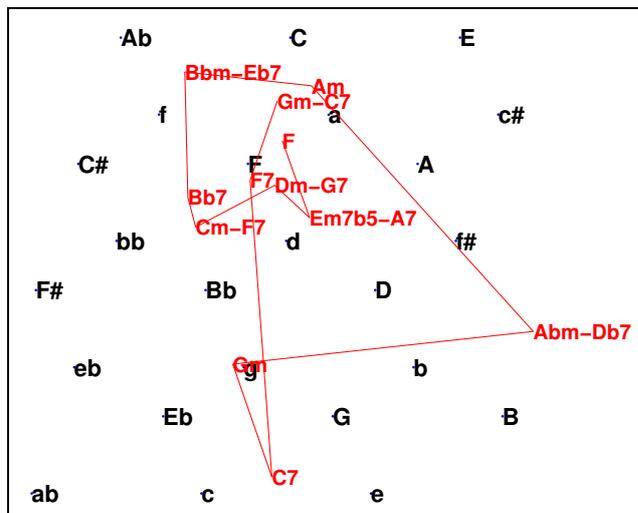


Figure 2. Charlie Parker’s solo on Blues for Alice

3.2. Melodic Process Example

In this example, we form a two-dimensional melodic control and analysis space. Anchors in the space are melodies with motion consisting of only small steps, only large leaps, leaps up with stepwise motion down, leaps down with stepwise motion up, and several other variants. Dissimilarities are computed in the following way. Building on Krumhansl’s interpretation [6] of Narmour’s melodic process [9], we first associate a Markov transition matrix with each melody that explains the empirical frequencies of occurrence of steps downward followed by steps upward, steps downward followed by steps downward, leaps upward followed by steps downward, etc. The dissimilarity between two melodies in this space is evaluated by computing the sum of the squared differences in the entries of the respective Markov matrices. Figure 3 shows the layout of the space, with graphical depictions of the melodies associated with the points.

The contours in the figure are contours of the objective function for the optimization problem. We see that the solution computed by the semidefinite relaxation method finds the global minimum, even though there is more than one stationary point of the function. This is a very desirable characteristic of the method.

4. MAX/MSP-MATLAB-SEDUMI IMPLEMENTATION

To use the proposed idea in a performance setting, we combine the live performance convenience of Max/MSP, and the numerical computation convenience of MATLAB and SeDuMi. We have installed all three software packages on a Macintosh running OSX, and have written a Max/MSP external C program that dynamically links to the MATLAB Engine. The MATLAB Engine is an interface designed to give programmers the ability to call

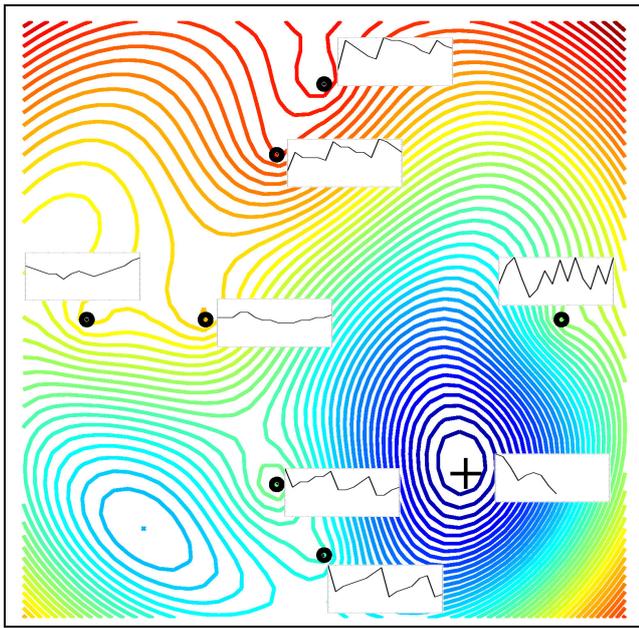


Figure 3. Melodic Process Space The contours of the objective function possess two local minima, yet we are able to locate the global minimizer.

functions and fully interact with a MATLAB session from another C program, in our case, our Max/MSP external. This allows the musician to perform complicated or specialized computation in a very convenient way. In our case, we call, via MATLAB, the semidefinite optimization software package SeDuMi, the numerical result of which appears in the Max/MSP environment. Work is being done to design a more general Max/MSP object to connect to MATLAB that would enable musicians to conveniently call complicated numerical routines in the midst of a live performance.

5. CONCLUSION

We described a method for the analysis of musical content in terms of placement within a defined control space. The method relies on numerical optimization by interior point methods, a relatively new methodology for solving convex optimization problems.

A detailed treatment of the underlying optimization problem was presented. The analysis of the problem in terms of duality and convexity is mathematically more informative than optimizing via traditional methods. This is because we are able to obtain an analytical formulation of a convex approximation to the nonconvex positioning problem. Having this representation allows the use of interior point methods which are efficient and well-behaved. We have applied the method to generate examples in harmonic and melodic analysis.

Finally, we mentioned an implementation of the system conveniently linking Max/MSP and MATLAB, through which we call the semidefinite program solver.

6. ACKNOWLEDGMENTS

The authors would like to thank Meyer Sound Laboratories, Rimantas Avizienis, Michael Zbyszynski, Adrian Freed, Laurent El-Ghaoui, and Lieven Vandenberghé for their help and advice.

7. REFERENCES

- [1] Cycling 74. Max/MSP. See www.cycling74.com for details.
- [2] Stephen Boyd and Lieven Vandenberghé. *Convex Optimization*. Cambridge University Press, 2004.
- [3] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of Association for Computing Machinery*, 42(6):1115–1145, June 1995.
- [4] Michael H. Goldsen, editor. *Charlie Parker Omnibook*. Atlantic Music Corp, 1978.
- [5] Carol L. Krumhansl. *Cognitive Foundations of Musical Pitch*. Oxford University Press, 1990.
- [6] Carol L. Krumhansl. Music psychology and music theory: Problems and prospects. *Music Theory Spectrum*, 17(1):53–80, 1995.
- [7] The Mathworks. MATLAB. See www.mathworks.com for details.
- [8] Ali Momeni and David Wessel. Characterizing and controlling musical material intuitively with geometric models. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression, Montreal, Canada, 2003*.
- [9] E. Narmour. *The Analysis and Cognition of Basic Melodic structures: The Implication-Realization Model*. University of Chicago Press, 1990.
- [10] Yurii Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optimization Methods and Software*, 9:141–160, 1998.
- [11] J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11:625–653, 1999.