Characteristic Shape Sequences for Measures on Images

Rachael L. Pingel ^{*} Mark A. Abramson [†] Thomas J. Asaki [‡]
J. E. Dennis, Jr. [§]

November 22, 2006

Abstract

Researchers in many fields often need to quantify the similarity between images using metrics that measure qualities of interest in a robust quantitative manner. We present here the concept of image dimension reduction through characteristic shape sequences. We formulate the problem as a nonlinear optimization program and demonstrate the solution on a test problem of extracting maximal area ellipses from two-dimensional image data. To solve the problem numerically, we augment the class of mesh adaptive direct search (MADS) algorithms with a filter, so as to allow infeasible starting points and to achieve better local solutions. Results here show that the MADS filter algorithm is successful in the test problem of finding good characteristic ellipse solutions from simple but noisy images.

Key words: Optimization, mesh adaptive direct search algorithms (MADS), filter methods, nonsmooth optimization, image metrics

1 Introduction

In this paper we present a method for property preserving dimension reduction in images that can be used to construct metrics for quantifying similarity between two images in a

^{*}Brigham Young University, Department of Mathematics, 292 TMCB, Provo, Utah 84602 USA, rachaelp@byu.net

[†]Air Force Institute of Technology, Department of Mathematics and Statistics, 2950 Hobson Way, Bldg 641, Wright Patterson AFB, Ohio 45433 USA, Mark.Abramson@afit.edu, <http://www.afit.edu/en/ENC/-Faculty/MAbramson/NOMADm.html>. Support for this author was provided by LANL.

[‡]Los Alamos National Laboratory, MS D413, Los Alamos, New Mexico 87545 USA asaki@lanl.gov

[§]Rice University, Department of Computational and Applied Mathematics, 8419 42nd Avenue SW, Seattle, WA 98136-2360 USA, dennis@rice.edu, <http://www.rice.edu/~dennis>. Support for this author was provided by LANL, AFORS F49620-01-1-0013, The Boeing Company, and ExxonMobil Upstream Research Company.

high-dimensional space. This should not be confused with the related concept of determining image measure properties, which is a function of only one image. Examples of image metrics include L_p distances, correlation functions, and image warping distances. The latter capitalizes on image registration techniques common in medical imaging applications [20]. Image metrics are also essential for data retrieval applications [21].

The most useful and versatile image comparison methods will share significant properties. A key idea is that a metric must quantitatively measure qualities of interest. Equally important is that a metric must ignore image differences that do not matter in the opinion of an expert. One such difference is noise, a common and understandable theme. However, robustness in the presence of non-noise features, such as distortions, occlusions, and missing data, is also important. Metrics are especially useful if they are automated and rely as little as possible on expert intervention.

Our consideration is the characterization of an image through extraction of geometric properties of image intensity level sets and a subsequent metric comparison. The problem that inspired our task is parameter estimation and hydrodynamic code validation studies using fluid-flow experimental data. The essential question is: How well does a given choice of hydrodynamic parameters and code implementation predict what is recorded in experimental images? The related question is: Can we define a metric in the space of images that can help us make better parameter choices and measure the quality of the simulation code? Experimental fluid images are typically characterized by an evolving object or region of interest which may have a simple geometry such as a wavefront, or a more complex character such as the onset of turbulent flow. In either case we can expect to capitalize on geometric property extraction. For other methods that focus on geometric object properties, see [21].

In this paper, we consider a method of characteristic shape sequences as a reduced dimension representation of images. The idea is to optimally map each intensity level set of the image onto a parameterized family of shapes. If the shape family is characteristic of image features of interest, then the shape parameter sequences define a low-dimensional representation of the original image. Image similarity can then be measured by sequence similarity metrics. This simple idea may be useful for certain classes of fluid dynamics experiments and medical imaging applications for which characteristic shapes can be of primary importance, background details are relatively simple, and level sets are not topologically complicated.

The concept of metric comparisons on reduced dimension measures is not new. The idea is to use an appropriate mapping from the high-dimensional image space to a low-dimensional measure space that preserves, as much as possible, the image qualities of interest, so that metrics can then be applied to the reduced image data. Recent applications include speech summarization [13], hyperspectral analysis [15], and experimental design [10]. There are two good reasons for this approach. First, it is possible, in principle, to formulate a property-preserving mapping to a reduced dimension space [9, 14, 16]. Second, we can take advantage of existing sequence similarity algorithms for metric comparison (see [24] and included references).

As a test case, we consider the sequence of maximal area ellipses interior to image level sets. Each ellipse is defined by five parameters in the image space, and each optimal ellipse is computed by solving a constrained nonlinear optimization problem. Robustness of the solution to noise and experimental uncertainty is enforced by allowing the ellipses to contain a specified small percentage of pixels that lie outside the specified image level set. The optimization problem is solved numerically by applying a mesh adaptive direct search (MADS) algorithm [6] with a filter [5]. The filter allows intermediate solutions that violate constraints in order to provide a more robust global search of the parameter space.

The paper is outlined as follows. In Section 2 we describe and formulate the problem in detail. In Section 3 we introduce a hybrid mesh adaptive direct search filter (filter-MADS) algorithm. Computational results obtained by the algorithm are presented and analyzed in Section 4, followed by some concluding remarks in Section 5.

2 Problem Formulation

In formalizing the mathematical description of the problem, we begin with a pixilated rectangular image $\Pi = \{z_{ij} = z(x_i, y_j) : i = 1, 2, \dots, M, j = 1, 2, \dots, N\}$, where $\{(x_i, y_j) : i = 1, 2, \dots, M, j = 1, 2, \dots, N\} \subset \mathbb{R}^2$ are the coordinates of the pixels and $z = z(x, y)$ is a real-valued function that measures intensity at location (x, y) . An example is given in Figure 1, where intensity values are differentiated by color. Without loss of generality,

Figure 1: Test Image for the maximal area ellipse problem

we assume that the image lies in the first quadrant of a rectangular coordinate system and that the coordinate (x_1, y_1) coincides with the origin. Thus, for each pixel $(x_i, y_j) \in \Pi$,

$i = 1, 2, \dots, M, j = 1, 2, \dots, N$, the associated pixel indices (i, j) can be determined by dividing x_i and y_j by the pixel width p_w and height p_h , respectively.

An ellipse $E = E(x_c, y_c, a, b, \theta) \subset \mathbb{R}^2$ can be represented by five variables: the location of its center $(x_c, y_c) \in \mathbb{R}^2$, the lengths of its axes $a \geq 0$ and $b \geq 0$ (measured from the center), and the angle of its rotation $\theta \in [0, \frac{\pi}{2}]$, measured counterclockwise from the positive x -axis. Its area is given by $A = \pi ab$.

The objective is to find the maximum area ellipse, subject to reasonable bound constraints on each variable, for which a specified percentage $\rho \in [0, 100]$ of the z -values interior to the ellipse meet a specified intensity threshold z_0 and for which a specified $\bar{\rho} \in [0, 100]$ percent of the z -values on the ellipse boundary meet the threshold z_0 . Specifically, the optimization problem can be formulated by

$$\begin{aligned}
 & \max \quad A = \pi ab \\
 & \text{subject to} \quad z(x, y) \geq z_0 \quad \text{for } \rho\% \text{ of } (x, y) \in \Pi \cap E(x_c, y_c, a, b, \theta) \\
 & \quad \quad \quad z(x, y) \geq z_0 \quad \text{for } \bar{\rho}\% \text{ of } (x, y) \in \partial E(x_c, y_c, a, b, \theta) \\
 & \quad \quad \quad p_w \leq x_c \leq x_{\max} \\
 & \quad \quad \quad p_h \leq y_c \leq y_{\max} \\
 & \quad \quad \quad 0 \leq a \leq \frac{1}{2} \max\{x_{\max}, y_{\max}\} \\
 & \quad \quad \quad 0 \leq b \leq \frac{1}{2} \max\{x_{\max}, y_{\max}\} \\
 & \quad \quad \quad 0 \leq \theta \leq \frac{\pi}{2},
 \end{aligned} \tag{1}$$

where $x_{\max} = Mp_w$ and $y_{\max} = Np_h$. Since no data exists for a point lying outside the image boundary (defined by the bound constraints in (1)), objective function values there are assumed to be infinite. This is not a problem for any of the algorithms used here, provided that the initial point satisfies the bound constraints.

3 Algorithm

3.1 Mesh Adaptive Direct Search (MADS)

The problem given in (1) is not solvable by traditional gradient-based methods because no derivative information for the nonlinear constraint functions is available. Thus we turn to derivative-free methods, and in particular, the class of mesh adaptive direct search (MADS) algorithms. MADS is an extension of the class of generalized pattern search (GPS) methods [17, 18, 23] to general nonlinear programming (NLP) problems that overcomes limitations of GPS when dealing with nonlinear constraints and nonsmoothness.

Lewis and Torczon [19] extended GPS to nonlinear constraints by applying GPS to an augmented Lagrangian subproblem based on the method proposed in [7]. Their approach ensures convergence for twice continuously differentiable objectives and constraints to a KKT

point. This is not a realistic assumption for the class of problems considered in this paper, and, in practice, the augmented Lagrangian approach may suffer from many of the problems associated with penalty function methods, such as the effective choices of penalty parameters and Lagrange multipliers.

Audet and Dennis [5] introduced a GPS filter method (filter-GPS), in which new iterates are generated whenever simple decrease in the objective or an aggregate constraint violation function is achieved. Filter methods were first introduced by Fletcher and Leyffer [12] as a way to globalize sequential linear programming (SLP) and sequential quadratic programming (SQP) without using a penalty function. The Audet-Dennis GPS filter approach avoids the pitfalls of penalty parameters, but its convergence to standard first-order stationary points is not guaranteed [5] due to GPS limitations.

Because of the weaknesses inherent to filter-GPS, Audet and Dennis [6] more recently introduced the class of MADS algorithms. Instead of limiting local exploration to a finite number of directions (as GPS does), MADS systematically generates an asymptotically dense set of directions in the limit. Because of this property, convergence to both first-order [6] and second-order [2] stationary points is guaranteed under reasonable assumptions without the need for a filter but with the requirement that the starting point be feasible.

One primary advantage of MADS is that it can handle problems with general “Yes/No” or set constraints. However, MADS requires the initial point to be feasible, while filter-GPS does not. In this paper, we describe a filter-MADS algorithm, which we show in Section 4 to be more advantageous than MADS and filter-GPS in solving (1).

The target class of problems is of the form,

$$\begin{aligned} \min_{x \in X} \quad & f(x) \\ \text{subject to} \quad & C(x) \leq 0, \end{aligned}$$

where $X \subset \mathbb{R}^n$, $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $C = (C_1, C_2, \dots, C_m) : \mathbb{R}^n \rightarrow \mathbb{R}^m$. When using filter-GPS, we must assume that the set X is a linearly constrained region; otherwise, convergence to stationary points is not guaranteed. However, no such restriction is needed for MADS. The algorithms described here are actually applied, not to the function f , but to the barrier function $f_X = f + \psi_X$, where ψ_X is an indicator function defined as zero inside X and $+\infty$ otherwise. Convergence of the algorithm is based on the smoothness properties of f , not f_X .

MADS is a class of algorithms that generates a sequence of iterates with nonincreasing objective function values. Associated with this sequence is a set of n_D directions, $D \subset \mathbb{R}^n$, that positively span $\mathbb{R}^n$; i.e., any vector in $\mathbb{R}^n$ can be represented as a nonnegative linear combination of vectors in D . Furthermore, each direction $d_j \in D$, $j = 1, 2, \dots, n_D$, must be constructed such that $d_j = Gz_j$, where $G \in \mathbb{R}^{n \times n}$ is a fixed nonsingular generating matrix and $z_j \in \mathbb{Z}^n$ is an integer vector. For convenience, the set D is also viewed as a real $n \times n_D$ matrix whose columns are the elements of the set.

At each iteration k , points are evaluated on a mesh constructed using vectors in D in an attempt to find a point with a function value lower than any previously evaluated points,

which we term an *improved mesh point*. The current mesh is defined as

$$M_k = \bigcup_{x \in S_k} \{x + \Delta_k^m Dz : z \in \mathbb{N}^{n_D}\},$$

where S_k is the set of points where the objective function f had already been evaluated by the start of iteration k (S_0 is the set of initial points), and Δ_k is the current mesh size parameter.

Each iteration of MADS has two main steps: an optional SEARCH and a local POLL. In the SEARCH step, any finite number of mesh points (including none) are evaluated. The strategy for picking such points is left entirely to the user. Common choices are discussed in [4] or [5], for example.

If the SEARCH fails to find an improved mesh point, then the more rigidly defined POLL step is invoked. In this step, a subset of mesh points near the incumbent are evaluated, where proximity is based on the *poll size parameter* Δ_k^p . This set of trial points is called a *frame*. Less general than the frames of Coope and Price [8], the MADS frame uses the current incumbent solution x_k (the *frame center*) and the poll and mesh size parameters, $\Delta_k^p > 0$ and $\Delta_k^m > 0$, respectively, and a positive spanning set of n_{D_k} directions $D_k \subset D$. The MADS frame at iteration k is more formally defined [6] to be the set:

$$P_k = \{x_k + \Delta_k^m d : d \in D_k\} \subset M_k,$$

where D_k is a positive spanning set such that $0 \notin D_k$, and for each $d \in D_k$,

- d can be written as a nonnegative integer combination of the directions in $D : d = Du$ for some vector $u \in \mathbb{N}^{n_D}$ that may depend on the iteration number k ;
- the distance from the frame center x_k to a frame point $x_k + \Delta_k^m d \in P_k$ is bounded by a constant times the poll size parameter: $\Delta_k^m \|d\| \leq \Delta_k^p \max\{\|d'\| : d' \in D\}$;
- limits (as defined in Coope and Price [8]) of the normalized sets D_k are positive spanning sets.

If either the SEARCH or POLL step is successful in finding an improved mesh point, then that point becomes the incumbent, and the mesh is retained or coarsened. If both the SEARCH and POLL steps fail to find an improved mesh point, then the current iterate is declared a *mesh local optimizer*, the frame is called a *minimal frame*, and the frame center x_k is a *minimal frame center*. In this case, the current iterate $x_{k+1} = x_k$ is retained, and the mesh is refined. Coarsening and refining of the mesh are performed according to the rule,

$$\Delta_{k+1}^m = \tau^{w_k} \Delta_k^m \tag{2}$$

$$\text{where } w_k \in \begin{cases} \{0, 1, \dots, w^+\} & \text{if an improved mesh point is found} \\ \{-1, -2, \dots, w^-\} & \text{otherwise.} \end{cases} \tag{3}$$

where $\tau > 1$ is rational, and $w^- \leq -1$ and $w^+ \geq 0$ are integers. The poll size parameter is also updated such that $\Delta_k^m \leq \Delta_k^p$ for all k , but also such that

$$\lim_{k \in K} \Delta_k^m = 0 \text{ if and only if } \lim_{k \in K} \Delta_k^p = 0, \quad (4)$$

where K denotes any subsequence of mesh local optimizers. Under this construction, GPS is the specific implementation of MADS for which $\Delta_k^m = \Delta_k^p$ for all k .

Figures 2 and 3 give examples of a GPS and MADS frame, respectively [6]. As shown

Figure 2: Example of GPS frames $P_k = \{x_k + \Delta_k^m d : d \in D_k\} = \{p^1, p^2, p^3\}$ for different values of $\Delta_k^m = \Delta_k^p$. In all three figures, the mesh M_k is the intersection of all lines.

Figure 3: Example of MADS frames $P_k = \{x_k + \Delta_k^m d : d \in D_k\} = \{p^1, p^2, p^3\}$ for different values of Δ_k^m and Δ_k^p . In all three figures, the mesh M_k is the intersection of all lines.

in Figure 3, the number points eligible to be chosen for a MADS poll set grows unbounded as the iteration sequence progresses. In fact, the total of all poll directions can be chosen to become asymptotically dense. A specific implementation of MADS was devised in [6] in which the sets D_k are limited in size but chosen randomly at each iteration. This preserves the denseness of poll directions, but does so only with probability one.

3.2 Filters

Filter-based algorithms attempt to minimize both the objective function f and a nonnegative aggregate constraint violation function h , where h satisfies the condition that $h(x) \geq 0$ with $h(x) = 0$ if and only if x is feasible. Consistent with [5], we specifically choose

$$h(x) = \|C(x)_+\|_2^2,$$

where $C(x)_+$ is the vector of constraint violations at x ; i.e., for $i = 1, 2, \dots, m$, $C_i(x)_+ = C_i(x)$ if $C_i(x) > 0$; otherwise, $C_i(x)_+ = 0$. This choice of h inherits whatever smoothness properties C possesses [5].

Also consistent with [5], we define a second constraint violation function, $h_X = h + \psi_X$, where the indicator function ψ_X is defined exactly as before. Thus, for $x \notin X$, $h_X(x) = \infty$, and the function h is not evaluated. Convergence properties with respect to h at the limit point of the algorithm depend on the local smoothness of h and not of h_X [5].

A filter is based on the idea of dominance, which is defined with respect to f and h below [3]. A formal filter definition follows.

Definition 3.1 *A point $x \in \mathbb{R}^n$ is said to dominate $y \in \mathbb{R}^n$, written $x \prec y$, if $f(x) \leq f(y)$ and $h_X(x) \leq h_X(y)$ with either $f(x) < f(y)$ or $h_X(x) < h_X(y)$.*

Definition 3.2 *A filter $\mathcal{F}$ is a finite set of points in the domain of f and h such that no pair of points x and y in the set have the relation $x \prec y$.*

Two additional restrictions are placed on the filter $\mathcal{F}$. First, a bound h_{max} is set on aggregate constraint violation so that each point $y \in \mathcal{F}$ satisfies $h_X(y) < h_{max}$. Secondly, to be consistent with [11, 12], only infeasible points are included in the filter, and feasible points are tracked separately. Using these two restrictions, we include the following terminology [3]:

Definition 3.3 *A point x is said to be filtered by a filter $\mathcal{F}$ if any of the following properties hold:*

1. $y \preceq x$ for some $y \in \mathcal{F}$,

Figure 4: A general MADS filter algorithm

$$2. \ h_X(x) \geq h_{\max},$$

$$3. \ h_X(x) = 0 \text{ and } f(x) \geq f^F, \text{ where } f^F = \min\{f(w) : w \in \mathcal{F}, h(w) = 0\}.$$

The point x is said to be unfiltered by $\mathcal{F}$ if it is not filtered by $\mathcal{F}$.

Note that one consequence of this definition is that, if $y \in \mathcal{F}$, then any other point $x \in \mathbb{R}^n$ satisfying $f(x) = f(y), h(x) = h(y)$ cannot be added to the filter.

3.3 The MADS Filter Algorithm

In adding the filter to the MADS algorithm, we change notation slightly. Instead of polling around the current iterate x_k , we poll around the frame center $p_k \in \{p_k^F, p_k^I\}$, where $p_k^F \in X$ is the best feasible point found so far (i.e., the feasible point with the lowest objective function value), and $p_k^I \in X$ is the least infeasible point (i.e., the infeasible point with the lowest value of h). If no feasible point has been found, then $p_k = p_k^I$ is chosen. Note that, by construction, both p_k^F and p_k^I , and thus p_k , will always satisfy the constraints that define X . A successful iteration is one that finds an unfiltered point is found, in which case, it is added to the filter.

The filter-MADS algorithm is described in detail in Figure 4. Its convergence properties are not presented here, but follow directly from those of the filter-GPS [5] and MADS [6] algorithms.

4 Numerical Implementation

To evaluate the effectiveness of the filter-MADS algorithm on the maximal area ellipse problem, we tested it, together with the filter-GPS and MADS algorithms, using the NO-MADm [1] software package. The test image (with random noise already added) is the one illustrated in Figure 1. It consists of a 256 by 260 matrix, where each entry corresponds to the z -value of a pixel, ranging in value from -0.0637 to 1.0573 . Each pixel has a height of $p_h = 0.0125$ and a width of $p_w = 0.01168$. The z -values are interpreted as the discrete pixel-center sampling of the unknown continuous image.

For the filter-GPS and filter-MADS algorithms, the domain X is defined by the five bound constraints in (1), while the other two nonlinear constraints are treated by the filter. For MADS, X is defined by all of the constraints. Recall that constraints that define X are treated by the barrier approach, in which points lying outside of X are discarded without being evaluated.

4.1 Initial Iterates

We developed two different methods of selecting the initial iterate. The algorithm is able to use either one, or a combination of the two. In the first approach, we set the initial iterate to be a modified version of the inscribed ellipse of the image. This ellipse, which usually violates the first two nonlinear constraints in (1), is given by

$$E_0 = E(x_c, y_c, a, b, \theta) = E\left(\frac{1}{2}x_{\max}, \frac{1}{2}y_{\max}, \frac{1}{2}x_{\max} - 2pw, \frac{1}{2}y_{\max} - 2ph, 0\right). \quad (5)$$

In the second approach, we compute local maxima of the signed distance function associated the level set $\{(x, y) : z(x, y) \geq z_0\}$. The signed distance function is computed using the fast marching method [22]. Each local maximizer then becomes the center of a circle whose radius extends 90% of the distance to the boundary of the level set. Each circle is then given an orientation so that the difference in intensity between the two diameter end-points is maximal. A list of these potential initial iterates is compiled, and any circle that significantly overlaps with another is removed from the list. The algorithm then evaluates each ellipse from the list as an initial iterate.

The first approach often starts infeasible, but with an ellipse of large area. This often leads to better solutions, but achieving feasibility can be a challenge, especially for high values of z_0 . On the other hand, the second approach is more robust in finding a feasible point,

since it begins with smaller ellipses, but it is also more likely to converge to a local solution. By combining the two methods, a third approach is available, in which the algorithm begins with the initial point E_0 and then uses the second approach if the algorithm fails to find a feasible iterate from E_0 .

4.2 Test Results

We solved the problem numerically using the two different initial points, as described in Section 4.1. When using the filter, polling at each iteration was done around *both* the best feasible point and the least infeasible point. Whenever either the MADS or filter-MADS algorithm was used, 10 replications were performed to handle the inherent random factor. Problem parameter values were set to $z_0 = 0.241$, $\rho = 95\%$, and $\bar{\rho} = 75\%$. Termination of the algorithm was set to occur when the mesh size parameter fell below 10^{-8} .

For both tests, we compared the performance of three algorithms; namely, filter-GPS (using the standard standard $2n$ directions), MADS, and filter-MADS. In the case of MADS, infeasible trial points were simply discarded. When using the MADS or filter-MADS algorithm, we recorded the best result from the 10 replications along with appropriate averages.

Tables 1 and 2 provide results for feasible and infeasible starting points, respectively.

Table 1: Results when starting from a feasible point.

	Evaluations	Area	x_c	y_c	a	b	θ
filter-GPS	1208	1.52018	1.3059	2.051	0.5763	0.8396	0.9132
MADS							
<i>best</i>	403	1.58541	0.8581	2.334	0.689	0.7325	0.5741
<i>average</i>	426	1.44494					
filter-MADS							
<i>best</i>	3161	1.72286	0.9192	2.2875	0.6981	0.7856	0.5509
<i>average</i>	1987	1.54562					

When starting from a feasible initial point, all three algorithms converged to a local minimizer, with filter-MADS finding a better solution than MADS or filter-GPS. However, when the infeasible initial point was used, there was a significant difference in results between the three methods. Although filter-GPS was able to find a feasible ellipse, it was far from optimal. This is not surprising because only two of the ten standard $2n$ directions used by filter-GPS (increasing a or b) improve the objective function value, since the other three variables, x_c , y_c , and θ , have no affect on the objective function value. In contrast, filter-MADS showed significant improvement. The expanded set of search directions in MADS allowed the algorithm to find a more eccentric ellipse (i.e., one with a higher eccentricity

Table 2: Results when starting from an infeasible point, E_0 .

	Evaluations	Area	x_c	y_c	a	b	θ
filter-GPS	757	1.24320	1.0186	2.3461	.7098	.5575	0
MADS							
<i>best</i>	81	failed	1.5184	1.6	1.575	1.485	0
<i>average</i>	81	failed					
filter-MADS							
<i>best</i>	488	1.79746	1.2059	2.0375	0.575	0.995	0.5
<i>average</i>	561	1.13941					

$e = \sqrt{1 - b^2/a^2}$) to better fit the constraints. Without a filter, MADS did not move from the infeasible initial point. Since MADS sets the objective function value of all infeasible points to $-\infty$ (for a maximization problem), it moves as soon as it finds a feasible point. However, in this case, it did not find one. This is consistent with the theory, which requires a feasible initial point to ensure convergence to a stationary point [6].

In Figure 5, three images displaying the level set of all $z \geq z_0$ are shown. Superimposed over each image is the boundary (and axes) of the best ellipse found by the filter-GPS, MADS, and filter-MADS algorithms, respectively, when starting from the infeasible point E_0 .

Figure 5: Resulting ellipses starting from the infeasible point, E_0 : filter-GPS, MADS, and filter-MADS, respectively

Once we were able to solve this problem for a particular value of z_0 , we constructed shape sequences by computing optimal ellipses for a full range of z_0 values. In particular, we partitioned the range of z -values into 100 evenly spaced nodes, $\bar{z}_i, i = 1, 2, \dots, 100$, ran 10 replications of the filter-MADS algorithm for each node (i.e., setting $z_0 = \bar{z}_i$ in (1)), and recorded the best and average optimal ellipses for each run. For these runs, we used as starting points both E_0 and the initial points found using the signed distance function.

Figure 6 shows a plot the optimal values of x , y , a , b and θ for each value of z_0 , and Figure 7 shows the best and average optimal objective function values versus z_0 . The sequences in Figure 6 represent the characteristic signatures of the test image in the ellipse representation. The plots do not include the few cases in which the algorithm failed to find a feasible ellipse. In addition, other vector quantities, such as eccentricity, can be obtained from the basic sequences.

The high-frequency variation in the ellipse signatures can be attributed to two sources. First, and most significant, filter-MADS is finding a locally optimal solution for each z_0 on an objective function surface complicated by significant noise. Second, in the presence of noise, the parameters will show a significant sensitivity to z_0 , even if the global solution is attained for each value of z_0 .

The sequences of Figure 6 can be compared to identically obtained sequences on other images as an image metric, but these ideas are beyond the scope of this paper.

Figure 6: Evolution of the optimal ellipse as values of z_0 increase

5 Conclusions

We have demonstrated a method for quantifying reduced-dimension measures on images which can be used for further exploration as image comparison metrics. The maximal area ellipse problem explored in this paper contains the essential features of this class of problems: whole image characterization, low-dimensional characteristic property representation, and

Figure 7: Evolution of the objective function as values of z_0 increase

robustness to noise. Other potential interesting characteristic shapes include models of bones, library templates for image classification, simple closed curves (rectangles, ellipses, stars), and alphanumeric characters.

While filter-GPS and MADS are both appropriate methods for treating the general class of problems, the hybrid filter-MADS approach achieves a better result, primarily due to the superior convergence properties of MADS and the flexibility of filter methods in starting from a good but infeasible point.

The filter-MADS solutions do show significant solution variability. This shortcoming can be addressed, as it was in this study, by considering the best of several runs. Alternatively, some image preprocessing to reduce noise signature is a reasonable analysis step.

Acknowledgments

This work was done while the first author was a summer undergraduate intern at the Air Force Institute of Technology. The authors wish to thank Peter Schultz for useful discussions. The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, United States Government, or research sponsors.

References

- [1] M. A. Abramson. NOMADm optimization software. <http://www.afit.edu/en/ENC/Faculty/MAbramson/NOMADm.html>.
- [2] M. A. Abramson and C. Audet. Second-order convergence of mesh adaptive direct search. *SIAM J. Optim.*, To appear, 2005.
- [3] M. A. Abramson, C. Audet, and J. E. Dennis Jr. Filter pattern search methods for mixed variable constrained optimization problems. Technical Report TR04-09, Department of Computational and Applied Mathematics, Rice University, Houston Texas, 2006.
- [4] C. Audet and J. E. Dennis, Jr. Analysis of generalized pattern searches. *SIAM J. Optim.*, 13(3):889–903, 2003.
- [5] C. Audet and J. E. Dennis, Jr. A pattern search filter method for nonlinear programming without derivatives. *SIAM J. Optim.*, 14(4):980–1010, 2004.
- [6] C. Audet and J. E. Dennis, Jr. Mesh adaptive direct search algorithms for constrained optimization. *SIAM J. Optim.*, 17(2):188–217, 2006.
- [7] A. R. Conn, N. I. M. Gould, and P. L. Toint. A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Numerical Analysis*, 28(2):545–572, April 1991.
- [8] I. D. Coope and C. J. Price. Frame-based methods for unconstrained optimization. *J. Optim. Theory Appl.*, 107(2):261–274, 2000.
- [9] M. Deza and M. Laurent. *Geometry of Cuts and Metrics*. Springer-Verlag, New York, 1997.
- [10] V. V. Fedorov, A. M. Herzberg, and S. L. Leonov. Methods of selecting informative variables. *Biometric Journal*, 48(1), 2006.
- [11] R. Fletcher, N. I. M. Gould, S. Leyffer, P. L. Toint, and A. Wächter. On the global convergence of trust-region SQP-filter algorithms for general nonlinear programming. *SIAM J. Optim.*, 13(3):635–659, 2002.
- [12] R. Fletcher and S. Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming*, Series A, 91(2):239–269, 2002.
- [13] M. Hirohata, Y. Shinnaka, K. Iwano, and S. Furui. Sentence-extractive automatics speech summarization and evaluation techniques. *Speech Communication*, 48(9), 2006.
- [14] W. Johnson and J. Lindenstrauss. Extensions of Lipschitz mapping into Hilbert space. *Contemp. Math*, 26, 1984.
- [15] N. Keshava. Distance metrics and band selection in hyperspectral processing with applications to material identification and spectral libraries. *IEEE Transactions on Geoscience and Remote Sensing*, 42(7), July 2004.

- [16] J. A. Leea, A. Lendasseb, and M. Verleysena. Nonlinear projection with curvilinear distances: Isomap versus curvilinear distance analysis. *Neurocomputing*, 57, 2004.
- [17] R. M. Lewis and V. Torczon. Pattern search algorithms for bound constrained minimization. *SIAM Journal on Optimization*, 9(4):1082–1099, 1999.
- [18] R. M. Lewis and V. Torczon. Pattern search methods for linearly constrained minimization. *SIAM Journal on Optimization*, 10(3):917–941, 2000.
- [19] R. M. Lewis and V. Torczon. A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Optimization*, 12(4):1075–1089, 2002.
- [20] J. Modersitzki. *Numerical Methods for Image Registration*. Oxford university Press, 2004.
- [21] H. Muller, N. Michoux, D. Bandon, and A. Geissbuhler. A review of content-based image retrieval systems in medical applications – clinical benefits and future directions. *International Journal of Medical Informatics*, 73, 2004.
- [22] J. A. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, Cambridge, 2nd edition, 1999.
- [23] V. Torczon. On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25, February 1997.
- [24] R. Yaniv and D. Burshtein. An enhanced dynamic time warping model for improved estimation of DTW parameters. *IEEE Transactions on Speech and Audio Processing*, 11(3), May 2003.