CORRESPONDENCE OF PROJECTED 3-D POINTS AND LINES USING A CONTINUOUS GRASP

M. J. HIRSCH, P. M. PARDALOS, AND M. G. C. RESENDE

ABSTRACT. The field of computer vision has experienced rapid growth over the past fifty years. Many computer vision problems have been solved using theory and ideas from algebraic projective geometry. In this paper, we look at a previously unsolved problem from object recognition, namely object recognition when the correspondences between the object and image data are not known *a priori*. We formulate this problem as a mixed-integer nonlinear optimization problem in terms of the unknown projection relating the object and image, as well as the unknown assignments of object points and lines to those in the image. The global optimum of this problem recovers the relationship between the object points and lines with those in the image. When certain assumptions are enforced on the allowable projections mapping the object into the image, a proof is provided which permits one to solve the optimization problem via a simple decomposition. We illustrate this decomposition approach on some example scenarios.

1. INTRODUCTION

In the foreword of Hartley and Zisserman (2000, p. xi), Faugeras writes "making a computer see was something that leading experts in the field of artificial intelligence thought to be at the level of difficulty of a summer student's project back in the sixties. Forty years later the task is still unsolved and seems formidable." One of the main reasons for this is that the biological vision and recognition process is still largely unknown and therefore hard to emulate on computers. The field of computer vision has grown out of the research directed towards "making a computer see." The past 15 years has seen a number of published books and articles presenting a mathematical framework for computer vision, considering problems from the perspective of algebraic projective geometry and invariant theory (Faugeras, 1993; Faugeras and Luong, 2001; Grosshans, 2005; Hartley and Zisserman, 2000; Ma et al., 2004). In this paper, we consider one of the still unsolved problems in computer vision, namely that of object recognition with unknown correspondence.

The general object recognition problem can be stated as follows: *Suppose an object is represented by a set of m_1 points and n_1 lines. Given a picture with m_2 points and n_2 lines, does the picture contain an image of the object, under a projective transformation (i.e., pinhole camera)?* Solutions to this problem have been derived when the correspondence between the object points and lines with those in the picture are known and the scenario is noise-free (Gleeson et al., 2003; Grosshans, 2005; Hartley and Zisserman, 2000). However, when the correspondence is not known, and when noise is present in the image data, this problem becomes quite difficult to solve.

Date: January 25, 2011. AT&T Labs Research Technical Report.

Key words and phrases. Object recognition, point/line projection, GRASP, continuous GRASP, global optimization, stochastic local search, nonlinear programming.

There have been some attempts to solve this object recognition problem when certain simplifying assumptions are enforced (e.g., intensity of image points are known and invariant across images (Chai and Ma, 1998; Hartley and Zisserman, 2000; Ma et al., 2004), errors are not dealt with directly in the problem formulation (Cheng et al., 1996; Scott and Longuet-Higgins, 1991; Shapiro and Brady, 1992), the correspondence is already known, and an inlier set is sought to determine the "best" resulting transformation (Chai and Ma, 1998; Hartley and Zisserman, 2000; Ma et al., 2004)). One approach in particular is the Random Sample Consensus (RANSAC) method (Fischler and Bolles, 1981). RANSAC begins by finding a minimal set of points which adequately fit to a prescribed model (i.e., an inlier set), and then enlarging the set of points according to those that also fit the resultant model. One of the main drawbacks to this approach is the potential combinatorial explosion in choosing minimal subsets of points. The authors apply the RANSAC algorithm to the location determination problem, but assume as input known correspondences between the 3-D and 2-D points.

Cass (1998) presents an approach for matching geometric features of known 3-D objects in 2-D images of a scene containing them. The geometric features consist of points, and uncertainty is represented by linear inequalities. The approach Cass took was to determine the maximal matching of point correspondences across the 3-D objects and 2-D images. In addition, a linear combination method is developed in Yi and Wang (2000) and Liu and Wang (1996) for the recognition of 3-D objects from 2-D images, and the objects can be articulated, which means that the object can morph in shape (e.g., a box with the top that opens – the image could have the top closed, partially open, or completely open). This approach stores a sample set of 2-D images of the object, from different perspectives. A linear combination of the stored 2-D images is formed and compared against the input image. If the comparison does not exceed a threshold, then the 2-D image is said to be that of the object. Of note is that the correspondence problem is not taken into account.

In this paper, we formulate the object recognition problem for points and lines as a mixed-integer nonlinear optimization problem. In particular, we derive equations for point and line correspondence directly accounting for possible noise in the image data. Using a specific type of pinhole camera, we provide a proof by which the optimization problem can be solved via a decomposition technique. This decomposition approach first solves a continuous non-convex global optimization problem to recover the camera parameters (i.e., projection matrix), and then solves a linear assignment problem to determine the correspondence between the 3-D points and lines with those in the image.

This paper is set up as follows. In Section 2, we derive equations for the point and line correspondences, explicitly accounting for noise in the image data, and provide a mixed-integer nonlinear optimization formulation for the object recognition problem. Section 3 details the decomposition approach to solve the optimization problem derived in Section 2. In Section 4, we discuss a new heuristic for continuous global optimization. This heuristic is utilized in the first step of the decomposition technique for the scenarios considered in the computational study (Section 5). Concluding remarks are provided in Section 6.

2. Object recognition problem formulation

In this section we begin by deriving equations that will be satisfied for 3-D to 2-D point and line correspondences when error is not present. We then alter these equations to explicitly account for possible uncertainty in the image points and lines, and provide a mixed-integer nonlinear optimization formulation. We begin with some notation that will be used throughout this paper.

Let $\alpha = \{\alpha_1, \ldots, \alpha_{m_1}, \alpha'_1, \ldots, \alpha'_{n_1}\}$ represent a three-dimensional object, i.e., an unordered collection of m_1 points and n_1 lines (without loss of generality, let the unprimed α's represent the points and the primed α's represent the lines). Let $\beta = \{\beta_1, \ldots, \beta_{m_2}, \beta'_1, \ldots, \beta'_{n_2}\}$ represent a two-dimensional picture, i.e., an unordered collection of m_2 points and n_2 lines (without loss of generality, let the unprimed β's represent the points and the primed β's represent the lines). For convenience, we represent the 3-D object and 2-D image in homogeneous coordinates (Gleeson et al., 2003; Grosshans, 2005; Hartley and Zisserman, 2000; Semple and Kneebone, 1952). Therefore, the 3-D point $\{x,y,z\}$ is represented as the column vector $\{x,y,z,1\}^T$ and the 2-D point $\{u,v\}$ is represented as the column vector $\{u,v,1\}^T$. An image formed by a pinhole camera (i.e., projective transformation) can then be represented as a projection from four-dimensional space to three-dimensional space, i.e., as a 3×4 matrix having 9 degrees of freedom (Hartley and Zisserman, 2000). Assume that $\Omega = [\Omega_{ij}]$ is such a matrix.

2.1. **Point equation derivation.** For an object point $\alpha_i = \{x_i, y_i, z_i, r_i\}^T$ to correspond with the point in the picture $\beta_j = \{u_j, v_j, w_j\}^T$, using Ω as the camera, it must be true that $\Omega\alpha_i = c_{ij}\beta_j$, for some $c_{ij} \neq 0$, where c_{ij} is needed to compensate for the invariance of homogeneous coordinates across scale. If we assume that $w_j \neq 0$ and $\Omega\alpha_i$ does not have its third coordinate equal to 0 (i.e., $\Omega_{31}x_i + \Omega_{32}y_i + \Omega_{33}z_i + \Omega_{34}r_i \neq 0$), then $\Omega\alpha_i = c_{ij}\beta_j$ is equivalent to the system of equations

$$\frac{\Omega_{11}x_i + \Omega_{12}y_i + \Omega_{13}z_i + \Omega_{14}r_i}{\Omega_{31}x_i + \Omega_{32}y_i + \Omega_{33}z_i + \Omega_{34}r_i} - \frac{u_j}{w_j} = 0 \tag{1}$$

$$\frac{\Omega_{21}x_i + \Omega_{22}y_i + \Omega_{23}z_i + \Omega_{24}r_i}{\Omega_{31}x_i + \Omega_{32}y_i + \Omega_{33}z_i + \Omega_{34}r_i} - \frac{v_j}{w_j} = 0. \tag{2}$$

Note that these assumptions on w_j and the third coordinate of $\Omega\alpha_i$ can be enforced in a straightforward manner (Gleeson et al., 2003; Grosshans, 2005).

If $\hat{\Omega\alpha}_i$ denotes the non-homogeneous representation of $\Omega\alpha_i$ and $\hat{\beta}_j$ denotes the non-homogeneous representation of β_j, then the system of equations (1) - (2) is equivalent to the vector equation $t_{ij} \equiv \hat{\Omega\alpha}_i - \hat{\beta}_j = \mathbf{0}$, which is equivalent to $\|t_{ij}\|_2 = 0$. Hence, when α_i and β_j correspond through the projection Ω, the Euclidean distance between $\hat{\Omega\alpha}_i$ and $\hat{\beta}_j$ will be 0.

Suppose now that the non-homogeneous image point $\hat{\beta}_j$ has some noise, or uncertainty, associated with it. This uncertainty can be represented as a 2×2 covariance matrix C_j, which we assume is positive semi-definite. Then, even when α_i and β_j correspond through the projection Ω, the Euclidean distance between $\hat{\Omega\alpha}_i$ and $\hat{\beta}_j$ might not be 0. This is due precisely to the uncertainty involved. To account for the uncertainty in the image point, we generalize the Euclidean distance $\|t_{ij}\|_2$ to the Mahalanobis distance $t_{ij}^T C_j^{-1} t_{ij}$ (Hartley and Zisserman, 2000). When this uncertainty is Gaussian in nature (it is generally assumed to be), then we can consider the Gaussian probability density function as representing the likelihood

$$F_{ij}(\Omega) = \frac{1}{2\pi\sqrt{|C_j|}} e^{-\frac{1}{2}t_{ij}^T C_j^{-1} t_{ij}}$$

of α_i and β_j corresponding through the projection Ω.

2.2. **Line equation derivation.** Let α'_k be the object line through the two points

$$P_{k1} = \{x_{k1}, y_{k1}, z_{k1}, r_{k1}\}^T \text{ and } P_{k2} = \{x_{k2}, y_{k2}, z_{k2}, r_{k2}\}^T.$$

It is clear that α'_k is defined uniquely by any two distinct points incident to α'_k. Let β'_ℓ be the image line represented as $\beta'_\ell = \{a_\ell, b_\ell, c_\ell\}^T$, where a two-dimensional homogeneous point $q = \{u, v, w\}^T$ is incident to β'_ℓ if and only if $q^T\beta'_\ell = ua_\ell + vb_\ell + wc_\ell = 0$. Then α'_k and β'_ℓ correspond through the projection Ω precisely when both ΩP_{k1} and ΩP_{k2} are incident to β'_ℓ, i.e., when $(\Omega P_{k1})^T\beta'_\ell = 0$ and $(\Omega P_{k2})^T\beta'_\ell = 0$. Unfortunately, this approach does not generalize to the case where there is error associated with the line β'_ℓ, as these two equations are not metric oriented. Therefore, we will look at the line case from a different perspective.

For $s \in \{1,2\}$, let $\hat{\Omega P}_{ks} = \{u_{ks}, v_{ks}, 1\}^T$ denote the non-homogeneous representation of ΩP_{ks}, where

$$u_{ks} = \frac{\Omega_{11}x_{ks} + \Omega_{12}y_{ks} + \Omega_{13}z_{ks} + \Omega_{14}r_{ks}}{\Omega_{31}x_{ks} + \Omega_{32}y_{ks} + \Omega_{33}z_{ks} + \Omega_{34}r_{ks}},$$
$$v_{ks} = \frac{\Omega_{21}x_{ks} + \Omega_{22}y_{ks} + \Omega_{23}z_{ks} + \Omega_{24}r_{ks}}{\Omega_{31}x_{ks} + \Omega_{32}y_{ks} + \Omega_{33}z_{ks} + \Omega_{34}r_{ks}}.$$

It can be shown (Boyce and DiPrima, 1988) that the non-homogeneous point incident to β'_ℓ that is closest to $\hat{\Omega P}_{ks}$ is given by

$$\hat{\delta}_{ks\ell} = \{\frac{b_\ell^2 u_{ks} - a_\ell b_\ell v_{ks} - a_\ell c_\ell}{a_\ell^2 + b_\ell^2}, \frac{a_\ell^2 v_{ks} - a_\ell b_\ell u_{ks} - b_\ell c_\ell}{a_\ell^2 + b_\ell^2}\}^T.$$

Defining $t_{ks\ell} = \hat{\delta}_{ks\ell} - \hat{\Omega P}_{ks}$, then we again have, in the error-free case, α'_k and β'_ℓ correspond through the projection Ω precisely when both $t_{k1\ell} = \mathbf{0}$ and $t_{k2\ell} = \mathbf{0}$, i.e., precisely when $\|t_{k1\ell}\|_2 = \|t_{k2\ell}\|_2 = 0$. Hence, when α'_k and β'_ℓ correspond through the projection Ω, the Euclidean distances from $\hat{\Omega P}_{k1}$ to $\hat{\delta}_{k1\ell}$ and from $\hat{\Omega P}_{k2}$ to $\hat{\delta}_{k2\ell}$ will both be 0.

Suppose now that each non-homogeneous point on the image line β'_ℓ has some noise, or uncertainty, associated with it. Again, represent this uncertainty by a 2×2 covariance matrix C_ℓ. Therefore, both $\hat{\delta}_{k1\ell}$ and $\hat{\delta}_{k2\ell}$ have an associated covariance matrix C_ℓ. With this uncertainty being Gaussian in nature we can form the likelihood

$$\bar{F}_{k\ell}(\Omega) = \frac{1}{2\pi\sqrt{|C_\ell|}} e^{-\frac{1}{2}(t_{k1\ell}^T C_\ell^{-1} t_{k1\ell} + t_{k2\ell}^T C_\ell^{-1} t_{k2\ell})}$$

of α'_k and β'_ℓ corresponding through the projection Ω.

2.3. **Problem formulation.** In addition to the unknown projection matrix, Ω, we also need to determine the correspondence between the 3-D and 2-D points and lines. Hence, for each possible 3-D point α_i and 2-D point β_j, we define a binary variable ψ_{ij} to denote whether α_i and β_j should be in correspondence. We similarly define binary variables $\phi_{k\ell}$ to denote whether the 3-D line α'_k and 2-D line β'_ℓ should be in correspondence. The object

recognition problem formulation is then given by

$$\min_{\Omega,\psi,\phi} f(\Omega,\psi,\phi) = -\sum_{i=1}^{m_1}\sum_{j=1}^{m_2}\psi_{ij}F_{ij}(\Omega) - \sum_{k=1}^{n_1}\sum_{\ell=1}^{n_2}\phi_{k\ell}\bar{F}_{k\ell}(\Omega) \tag{3}$$

s.t.

$$\sum_{i=1}^{m_1}\psi_{ij} \leq 1 \;\forall\; j \in \{1,\ldots,m_2\} \tag{4}$$

$$\sum_{j=1}^{m_2}\psi_{ij} \leq 1 \;\forall\; i \in \{1,\ldots,m_1\} \tag{5}$$

$$\sum_{k=1}^{n_1}\phi_{k\ell} \leq 1 \;\forall\; \ell \in \{1,\ldots,n_2\} \tag{6}$$

$$\sum_{\ell=1}^{n_2}\phi_{k\ell} \leq 1 \;\forall\; k \in \{1,\ldots,n_1\} \tag{7}$$

$$\psi_{ij} \in \{0,1\} \;\forall\; i \in \{1,\ldots,m_1\}, \forall\; j \in \{1,\ldots,m_2\} \tag{8}$$

$$\phi_{k\ell} \in \{0,1\} \;\forall\; k \in \{1,\ldots,n_1\}, \forall\; \ell \in \{1,\ldots,n_2\} \tag{9}$$

$$\Omega \in W, \tag{10}$$

where W is the space of projective transformations. Note that constraints (4), (5), and (8) ensure that each 3-D (2-D) point will have at most one corresponding 2-D (3-D) point, and similarly constraints (6), (7), and (9) ensure that each 3-D (2-D) line will have at most one corresponding 2-D (3-D) line.

3. PROBLEM DECOMPOSITION

A general pinhole camera, Ω, can be represented by $\Omega = K[R|t]$, where K is an internal camera calibration matrix (of size 3×3, with 3 degrees of freedom ($dof(K) = 3$)), R is a 3-D rotation matrix (of size 3×3, with $dof(R) = 3$), and t is a translation vector (of size 3×1, with $dof(t) = 3$) (Hartley and Zisserman, 2000). When K is set to the identity matrix, then the resulting projection matrix Ω preserves distances between points. In what follows, we make use of the following theorem.

Theorem 3.1. *Let $\tilde{f}(\Omega) = -\sum_{i=1}^{m_1}\sum_{j=1}^{m_2}F_{ij}(\Omega) - \sum_{k=1}^{n_1}\sum_{\ell=1}^{n_2}\bar{F}_{k\ell}(\Omega)$ and N_c^p (N_c^q) represent the truthful number of point (line) correspondences between the object and image. Suppose the following:* ***i)*** *Ω preserves distances between points, i.e., $\|x-y\|_2 = \|\Omega(x)-\Omega(y)\|_2$;* ***ii)*** *$N_c = N_c^p + N_c^q \geq dof(\Omega)$; and* ***iii)*** *The object point and line geometry is sufficiently random.*

Then, in the case where there is no noise in the image (i.e., covariance matrices are **0***), the Ω that minimizes $\tilde{f}(\Omega)$ is precisely the Ω mapping the points and lines of the 3-D object onto their truthful correspondences in the image.*

Proof. Let $\tilde{f}'(\Omega)$ equal $\tilde{f}(\Omega)$ when no random noise is present in the image. Then

$$\begin{aligned}
\tilde{f}'(\Omega) &= \lim_{\substack{C_j \to 0 \,\forall\, j \\ C_\ell \to 0 \,\forall\, \ell}} \tilde{f}(\Omega) \\
&= \lim_{\substack{C_j \to 0 \,\forall\, j \\ C_\ell \to 0 \,\forall\, \ell}} \Big[-\sum_{i=1}^{m_1}\sum_{j=1}^{m_2} F_{ij}(\Omega) - \sum_{k=1}^{n_1}\sum_{\ell=1}^{n_2} \bar{F}_{k\ell}(\Omega) \Big] \\
&= -\sum_{i=1}^{m_1}\sum_{j=1}^{m_2} \lim_{C_j \to 0 \,\forall\, j} F_{ij}(\Omega) - \sum_{k=1}^{n_1}\sum_{\ell=1}^{n_2} \lim_{C_\ell \to 0 \,\forall\, \ell} \bar{F}_{k\ell}(\Omega) \\
&= -\sum_{i=1}^{m_1}\sum_{j=1}^{m_2} \delta(\|t_{ij}\|^2) - \sum_{k=1}^{n_1}\sum_{\ell=1}^{n_2} \delta(\|t_{k\ell}\|^2)
\end{aligned}$$

where $\delta(x)$ is the Dirac-delta, or unit-impulse, function, i.e., $\delta(x-a) = 0$ if $x \neq a$ and $\int_{-\infty}^{\infty} \delta(x)\mathrm{d}x = 1$ (Dirac, 1926; Greenberg, 1988).

Now, $N_c = N_c^p + N_c^q \leq \min[m_1, m_2] + \min[n_1, n_2]$. Without loss of generality, assume that the first N_c^p object points correspond with the first N_c^p image points and the first N_c^q object lines correspond with the first N_c^q image lines.

For each $i \in \{1,\ldots,m_1\},\ j \in \{1,\ldots,m_2\}$, let $\Omega_{ij} = \{\Omega \in W \mid \hat{\beta}_j = \hat{\Omega}\alpha_i\}$. Similarly, for each $k \in \{1,\ldots,n_1\},\ \ell \in \{1,\ldots,n_2\}$, let $\check{\Omega}_{k\ell} = \{\Omega \in W \mid (\Omega P_{k1})^T \beta'_\ell = 0,\ (\Omega P_{k2})^T \beta'_\ell = 0,$ P_{k1} and P_{k2} distinct points incident to $\alpha'_k\}$. Since there is no random noise and $N_c \geq dof(\Omega)$, $\bar{\Omega} = \Big[\underset{i\in\{1,\ldots,N_c^p\}}{\cap} \Omega_{ii}\Big] \bigcap \Big[\underset{k\in\{1,\ldots,N_c^q\}}{\cap} \check{\Omega}_{kk}\Big]$ is well-defined and unique.

Choose any sets of pairs $\{\{i_1,j_1\},\ldots,\{i_r,j_r\}\}$ and $\{\{k_1,\ell_1\},\ldots,\{k_s,\ell_s\}\}$ with the following properties:
a) $i_t \in \{1,\ldots,m_1\}\ \forall\, t \in \{1,\ldots,r\}$ and $i_1,\ldots,i_r$ are distinct;
b) $j_t \in \{1,\ldots,m_2\}\ \forall\, t \in \{1,\ldots,r\}$ and $j_1,\ldots,j_r$ are distinct;
c) $k_t \in \{1,\ldots,n_1\}\ \forall\, t \in \{1,\ldots,s\}$ and $k_1,\ldots,k_s$ are distinct;
d) $\ell_t \in \{1,\ldots,n_2\}\ \forall\, t \in \{1,\ldots,s\}$ and $\ell_1,\ldots,\ell_s$ are distinct;
e) $r+s \geq dof(\Omega)$;
f) $\Big(\exists\, t \in \{1,\ldots,r\} \ni i_t \neq j_t\Big) \vee \Big(\exists\, t \in \{1,\ldots,s\} \ni k_t \neq \ell_t\Big)$;

Since the 3-D point and line geometry is sufficiently random,

$$\mathbf{P}\Big(\Big[\underset{t\in\{1,\ldots,r\}}{\cap} \Omega_{i_t j_t}\Big] \bigcap \Big[\underset{t\in\{1,\ldots,s\}}{\cap} \check{\Omega}_{k_t \ell_t}\Big] = \emptyset\Big) = 1.$$

Therefore, for any sets of pairs $\{\{i_1,j_1\},\ldots,\{i_r,j_r\}\}$ and $\{\{k_1,\ell_1\},\ldots,\{k_s,\ell_s\}\}$ satisfying the above conditions a) through f), the probability that there exists an Ω projecting each of the 3-D points α_{i_t} onto the 2-D points β_{j_t} $(t = 1,\ldots,r)$ and each of the 3-D lines α'_{k_t} onto the 2-D lines β'_{ℓ_t} $(t = 1,\ldots,s)$ is zero. Hence, the global minimum of $\tilde{f}'(\Omega)$ occurs at $\Omega = \bar{\Omega}$, i.e., at the Ω mapping the N_c^p 3-D points and N_c^q 3-D lines onto their truthful correspondences. □

Theorem 3.1 provides the basis of our decomposition approach. When there is no noise present in the image we can solve the optimization Problem (3) - (10) by first finding the Ω that minimizes $\tilde{f}(\Omega)$. Projecting all the 3-D points and lines onto the image via the recovered Ω, the second step is to apply a linear assignment algorithm to determine the 3-D to 2-D point and line correspondences. This is different from the RANSAC approach (Fischler and Bolles, 1981) in that we do not make use of known (or assumed) 3-D to

2-D point/line correspondences in order to determine a potential projection matrix. The global minimum of the objective function (3), as shown in Theorem 3.1, in the absence of noise and degenerate point/line configurations, occurs at the mode of the potential projection matrices, which is precisely the truthful projection matrix. We make note of the following. First, $\tilde{f}(\Omega)$ is non-convex in Ω, and while there exist algorithms for finding a *local minimum* in a finite number of iterations, it is well known that finding the global minimum is inherently unsolvable in a finite number of steps (Boender and Romeijn, 1995; Zhan, 2005). Second, in practice, there is normally some noise in the image data, thus making the global minimum of $\tilde{f}(\Omega)$, only an approximation to the true projection matrix. To solve the first step in the decomposition approach, we make use of a new global optimization heuristic, Continuous GRASP, or simply C-GRASP, as detailed in Section 4. The linear assignment algorithm used for the second step of the decomposition is an implementation of the JVC algorithm (Jonker and Volgenant, 1987).

In summary, the decomposition approach we propose to solve Problem (3) - (10) is given by:

(1) Use the C-GRASP heuristic to find the Ω that minimizes $\tilde{f}(\Omega)$.
(2) Apply the JVC linear assignment algorithm to determine the association between the 3-D points and lines with those in the image.

4. CONTINUOUS GRASP

Feo and Resende (1989; 1995) describe the metaheuristic GRASP (greedy random adaptive search procedures) as a multi-start local search procedure, where each GRASP iteration consists of two phases, a construction phase and a local search phase. In the construction phase, interactions between greediness and randomization generate a diverse set of good-quality solutions. The local search phase attempts to improve the solutions found by construction. The best solution over all of the multi-start iterations is retained as the final solution.

Hirsch et al. (2007b) describe C-GRASP, an adaptation of GRASP to solve continuous global optimization problems (see also Hirsch (2006)). C-GRASP works by discretizing the domain into a uniform grid. Both the construction and local improvement phases move along points on the grid. As the algorithm progresses, the grid adaptively becomes more dense. C-GRASP resembles GRASP in that it is a multi-start stochastic search metaheuristic that uses a randomized greedy construction procedure to generate starting solutions for a local improvement algorithm. The main difference is that an iteration of C-GRASP does not consist of a single greedy randomized construction followed by local improvement, but rather a series of construction-local improvement cycles with the output of construction serving as the input of the local improvement, as in GRASP. Unlike GRASP, however, the output of the C-GRASP local improvement procedure serves as the input of the C-GRASP construction procedure.

Hirsch et al. (2010) proposed modifications to the original C-GRASP algorithm, resulting in a significant decrease in the number of objective function evaluations required to converge to the global optimum. The modified C-GRASP heuristic was applied to a sensor registration problem in Hirsch et al. (2006), which is similar to the object recognition problem considered presently. It was also used to solve systems of non-linear equations (Hirsch et al., 2009) and to determine the relationship between drugs and adverse reactions (Hirsch et al., 2007a). In the remainder of this section, we detail this version of the C-GRASP heuristic.

```
procedure C-GRASP(n, ℓ, u, f(·), MaxRS, h_s, h_e, ρ_lo)
1     f* ← ∞;
2     for j = 1, ..., MaxRS do
3           x ← UnifRand(ℓ, u);
4           h ← h_s;
5           while h ≥ h_e do
6                 Impr_C ← false;
7                 Impr_L ← false;
8                 [x, Impr_C] ← ConstructGreedyRandomized(x, f(·), n, h, ℓ, u, Impr_C);
9                 [x, Impr_L] ← LocalImprovement(x, f(·), n, h, ℓ, u, ρ_lo, Impr_L);
10                if f(x) < f* then
11                      x* ← x;
12                      f* ← f(x);
13                end if
14                if Impr_C = false and Impr_L = false then
15                      h ← h/2;      /* make grid more dense */
16                end if
17          end while
18    end for
19    return(x*);
end C-GRASP;
```

FIGURE 1. Pseudo-code for C-GRASP.

4.1. **The heuristic.** Pseudo-code for the C-GRASP heuristic is shown in Figure 1. The procedure takes as input the problem dimension n, lower and upper bound vectors ℓ and u, the objective function $f(\cdot)$, as well as the parameters MaxRS, h_s, h_e, and ρ_{lo}. MaxRS defines the number of multi-starts to perform in the C-GRASP algorithm, h_s and h_e define the starting and ending discretization levels for each multi-start, and ρ_{lo} defines the portion of the neighborhood of the current solution that is searched during the local improvement phase of the algorithm.

Line 1 of the pseudo-code initializes the objective function value f^* of the best solution found to infinity. We perform MaxRS independent multi-starts. For each multi-start, Line 3 initializes the solution x to a random point distributed uniformly over the box in $\mathbb{R}^n$ defined by ℓ and u. The parameter h, that controls the discretization density of the search space, is initialized to h_s in line 4. The code sequentially executes lines 6 to 16 as long as $h \geq h_e$. The construction and local improvement phases are then called in lines 8 and 9, respectively. The solution returned from the local improvement procedure is compared against the current best solution in line 10. If the returned solution has a smaller objective value than the current best solution, then, in lines 11– 12, the current best solution is updated with the returned solution. In line 14, if the variables $\texttt{Impr}_\texttt{C}$ and $\texttt{Impr}_\texttt{L}$ are still set to false, then the grid density is increased by halving h, in line 15. The variable $\texttt{Impr}_\texttt{C}$ is false upon return from the construction procedure if and only if no improvement is made in the construction phase. Section 4.3 shows that the $\texttt{Impr}_\texttt{L}$ variable is false on return from the local improvement procedure if and only if the input solution to local improvement is determined to be an *h-local minimum*. We increase the grid density at this stage because

repeating the construction procedure with the same grid density will not improve the solution. This approach allows C-GRASP to start with a coarse discretization and adaptively increase the density as needed, thereby intensifying the search in a more dense discretization when a good solution has been found. The best solution found, at the end of the `MaxRS` multi-starts, is returned in line 19.

4.2. **Construction procedure.** In this section, we describe in detail the construction procedure. The construction algorithm combines greediness and randomization to produce a diverse set of good-quality solutions from which to start the local improvement phase. The construction algorithm is shown in Figure 2. The input is a solution vector x. To start, line 1 of the algorithm allows all coordinates of x to change in the current construction call (i.e. these coordinates are unfixed). In turn, in line 10 of the pseudo-code, if `ReUse` is `false`, a line search is performed in each unfixed coordinate direction i of x with the other $n-1$ coordinates of x held at their current values. In lines 10 and 11 of the pseudo-code, the value z_i for the i-th coordinate that minimizes the objective function, together with the objective function value g_i, are saved. In line 11, $\check{x}^i$ denotes x with the i-th coordinate set to z_i.

After looping through all unfixed coordinates (lines 7–16), in lines 17–23 a restricted candidate list (RCL) is formed containing the unfixed coordinates i whose g_i values are less than or equal to $\underline{g} + \alpha \cdot (\bar{g} - \underline{g})$, where $\bar{g}$ and $\underline{g}$ are, respectively, the maximum and minimum g_i values over all currently unfixed coordinates of x, and $\alpha \in [0,1]$ is randomly determined in line 2. In line 24, a coordinate is chosen at random from the `RCL`, say coordinate $j \in$ `RCL`. Line 25 checks whether x_j and z_j are equal. If so, line 26 sets `ReUse` to the value `true`. Otherwise, in lines 28–30, `ReUse` is set to `false`, $\texttt{Impr}_C$ is set to `true`, and x_j is set to equal z_j. Finally, in line 30, we fix coordinate j of x, by removing j from the set `UnFixed`. Choosing a coordinate by selecting at random from the `RCL` ensures both greediness and randomness in the construction phase. The above procedure is continued until all of the n coordinates of x have been fixed. At that stage, x and $\texttt{Impr}_C$ are returned from the construction procedure.

Note that the `ReUse` variable is utilized to speed up computations by avoiding unnecessary line searches. More details can be found in Hirsch et al. (2010). The parameter α controls the size of the `RCL` and therefore determines the mix of greediness and randomness in the construction procedure. Different values of α throughout the run allow some construction phases to be more greedy while others to be more random.

4.3. **Local improvement procedure.** C-GRASP makes no use of derivatives. Though derivatives can be easily computed for many functions, they are not always available or efficiently computable for all functions. The local improvement phase (with pseudo-code shown in Figure 3) can be seen as *approximating* the role of the gradient of the objective function $f(\cdot)$. From a given input point $x \in \mathbb{R}^n$, the local improvement algorithm generates a neighborhood, and determines at which points in the neighborhood, if any, the objective function improves. If an improving point is found, it is made the current point and the local search continues from the new solution.

Let $\bar{x} \in \mathbb{R}^n$ be the current solution and h be the current grid discretization parameter. Define

$$S_h(\bar{x}) = \{x \in S \mid \ell \leq x \leq u,\ x = \bar{x} + \tau \cdot h,\ \tau \in \mathbb{Z}^n\}$$

to be the set of points in S that are integer steps (of size h) away from $\bar{x}$. Let

$$B_h(\bar{x}) = \{x \in S \mid x = \bar{x} + h \cdot (x' - \bar{x})/\|x' - \bar{x}\|,\ x' \in S_h(\bar{x}) \setminus \{\bar{x}\}\}$$

```
procedure ConstructGreedyRandomized(x, f(·), n, h, ℓ, u, Impr_C)
1     UnFixed ← {1,...,n};
2     α ← UnifRand(0,1);
3     ReUse ← false;
4     while UnFixed ≠ ∅ do
5           g ← +∞;
6           ḡ ← −∞;
7           for i = 1,...,n do
8                 if i ∈ UnFixed then
9                       if ReUse = false then
10                            z_i ← LineSearch(x, h, i, n, f(·), ℓ, u);
11                            g_i ← f(x^i);
12                      end if
13                      if g > g_i then g ← g_i;
14                      if ḡ < g_i then ḡ ← g_i;
15                end if
16          end for
17          RCL ← ∅;
18          Threshold ← g + α·(ḡ − g);
19          for i = 1,...,n do
20                if i ∈ UnFixed and g_i ≤ Threshold then
21                      RCL ← RCL ∪ {i};
22                end if
23          end for
24          j ← RandomlySelectElement(RCL);
25          if x_j = z_j then
26                ReUse ← true;
27          else
28                x_j ← z_j;
29                ReUse ← false;
30                Impr_C ← true;
31          end if
32          UnFixed ← UnFixed \ {j};      /* Fix coordinate j. */
33    end while
34    return(x, Impr_C);
end ConstructGreedyRandomized;
```

FIGURE 2. Pseudo-code for C-GRASP construction phase.

be the projection of the points in $S_h(\bar{x}) \setminus \{\bar{x}\}$ onto the hyper-sphere centered at $\bar{x}$ of radius h. The *h-neighborhood* of the point $\bar{x}$ is defined as the set of points in $B_h(\bar{x})$.

The local improvement procedure is given a starting solution $x \in S \subseteq \mathbb{R}^n$. The current best local improvement solution x^* is initialized to x in line 1. Lines 3 and 4 determine the number of grid points, based on the current value of the discretization parameter h, and the maximum number of points in $B_h(x^*)$ that are to be examined. This number of grid points is defined by the parameter ρ_{lo} which is the portion of the neighborhood which is to be

```
procedure LocalImprovement(x, f(·), n, h, ℓ, u, ρ_lo, Impr_L)
1     x* ← x;
2     f* ← f(x);
3     NumGridPoints ← ∏_{i=1}^{n} ⌈(u_i − ℓ_i)/h⌉;
4     MaxPointsToExamine ← ⌈ρ_lo · NumGridPoints⌉;
5     NumPointsExamined ← 0;
6     while NumPointsExamined ≤ MaxPointsToExamine do
7           NumPointsExamined ← NumPointsExamined + 1;
8           x ← RandomlySelectElement(B_h(x*));
9           if ℓ ≤ x ≤ u and f(x) < f* then
10                x* ← x;
11                f* ← f(x);
12                Impr_L ← true;
13                NumPointsExamined ← 0;
14          end if
15    end while
16    return(x*, Impr_L);
end LocalImprovement;
```

FIGURE 3. Pseudo-code for C-GRASP local improvement phase.

examined. If all of these points are examined and no improving point is found, the current solution x^* is considered an *h-local minimum*.

Starting at the point x^*, in the loop in lines 6–15, the algorithm randomly selects `MaxPointsToExamine` points in $B_h(x^*)$, one at a time. In line 9, if the current point x selected from $B_h(x^*)$ is feasible and is better than x^*, then x^* is set to x, $\texttt{Impr}_\texttt{L}$ is set to `true`, and the process restarts with x^* as the starting solution. $\texttt{Impr}_\texttt{L}$ is used to determine whether the local improvement procedure improved the best solution. Local improvement is terminated when an *h-local minimum* solution x^* is found. At that point, x^* and $\texttt{Impr}_\texttt{L}$ are returned from the local improvement procedure.

5. COMPUTATIONAL STUDY

To illustrate our approach for solving the object recognition optimization problem (3) - (10), we consider three classes of scenarios. Each class is designated by the number of 3-D and 2-D points and lines, as well as the number of truthful 3-D to 2-D point and line correspondences. The details for each scenario class are given in Table 1. This table lists the number of 3-D points and lines, the number of 2-D points and lines, and the number of true 3-D and 2-D point and line correspondences. For each class, we created 10 test instances, in the following way. Each instance started with a randomly determined set of 3-D points and lines. A projection matrix Ω_{truth} was also randomly chosen. Based on the number of truthful correspondences between 3-D and 2-D points and lines, a subset of the 3-D points and lines were projected using Ω_{truth}, onto the 2-D image plane. Additional 2-D points and lines were added randomly to the image plane. All 2-D points and lines were then subjected to Gaussian noise, with zero mean and standard deviation values from the vector $\{0, 0.1, 0.2, \ldots, 0.9, 1.0\}$.

Each randomly generated scenario instance was input to the C-GRASP algorithm and run for a fixed number of multi-starts. The C-GRASP parameters are listed in Table 2.

TABLE 1. Object recognition scenario classes.

Scenario	Points			Lines		
Class	3-D	2-D	Num True Corr.	3-D	2-D	Num True Corr.
1	15	12	10	-	-	-
2	-	-	-	16	13	10
3	6	9	5	5	10	4

TABLE 2. Object recognition C-GRASP parameters.

Number Multi-Starts	100
h_s	0.10
h_e	0.01
ρ_{lo}	0.70

When the C-GRASP algorithm finished executing for a scenario instance, using the best projection matrix found, $\Omega_{computed}$, the 3-D points and lines from the scenario instance were projected onto the 2-D image plane, and the JVC linear assignment algorithm (Jonker and Volgenant, 1987) was applied to determine the optimal assignment of 3-D points and lines to those on the 2-D image plane.

We begin our discussion of the results by showing an example from Class 1. Tables 3 and 4 list the set of 15 3-D and 12 2-D points, respectively, for a particular instance from the Class 1 data, when the noise level is set to 0.6. The first two columns of Table 5 present the truthful correspondences between the 3-D object points and the 2-D image points. The truthful projection matrix to project the 3-D object points onto their corresponding 2-D image points is given by

$$\Omega_{truth} = \begin{bmatrix} 0.391 & 0.583 & -0.713 & 1.000 \\ -0.882 & 0.015 & -0.471 & -3.558 \\ -0.263 & 0.813 & 0.520 & -0.767 \end{bmatrix}.$$

After running our algorithm (C-GRASP, along with JVC), we determined a projection matrix of

$$\Omega_{computed} = \begin{bmatrix} 0.3975 & 0.5755 & -0.7147 & 2.5375 \\ -0.8689 & -0.0144 & -0.4949 & -1.0250 \\ -0.2951 & 0.8177 & 0.4943 & -2.3750 \end{bmatrix}$$

and a 3-D to 2-D point correspondence as given in the first and third columns of Table 5. In addition, Figure 4 displays the 3-D points, projected using $\Omega_{computed}$ onto the 2-D image plane, the 2-D points, and the computed correspondences (via a green dashed line). As can be seen from the table and figure, for this test instance our approach determines the correct correspondence of 9 of the 10 truthful.

Figures 5 – 7 display the average number of correspondences correctly determined for the three classes. As can be seen from these figures, the overall trend is for our decomposition algorithm to perform better when there is not a lot of noise in the image, and degrade gracefully as the noise increases. Further evidence for this claim is shown in Figures 8 – 10. These figures display the average distance between the 3-D points and lines, projected

TABLE 3. Class 1 Example, 3-D object points.

Point	x	y	z
1	−20	−49	9
2	−35	5	28
3	−20	43	−47
4	−50	39	35
5	−7	11	18
6	17	19	30
7	32	43	41
8	25	−5	−12
9	−34	−33	−3
10	5	11	−42
11	2	1	42
12	41	−14	17
13	−6	27	14
14	−46	32	49
15	46	−19	−1

TABLE 4. Class 1 Example, 2-D image points.

Point	x	y
A	1.36391260	−0.29771817
B	−1.09763397	0.52575762
C	3.44958041	2.45901026
D	−0.33330668	0.39606649
E	−0.42118354	−0.29559591
F	−0.10371727	−1.25394914
G	0.19888931	−1.07164390
H	−0.92978162	1.13549158
I	1.45428633	−1.35456172
J	−2.62180995	−0.79750469
K	−0.01296734	2.41371005
L	−1.91425568	1.01449136

TABLE 5. Class 1 Example - truthful and computed correspondences between 3-D object points and 2-D image points. This can also be seen in Figure 4.

Projected	$2-D$ Correspondences	
$3-D$ Points	Truthful	Computed
1	–	A
2	B	B
3	–	–
4	C	C
5	E	D
6	F	F
7	–	E
8	–	–
9	H	H
10	G	G
11	I	I
12	J	J
13	–	–
14	K	K
15	L	L

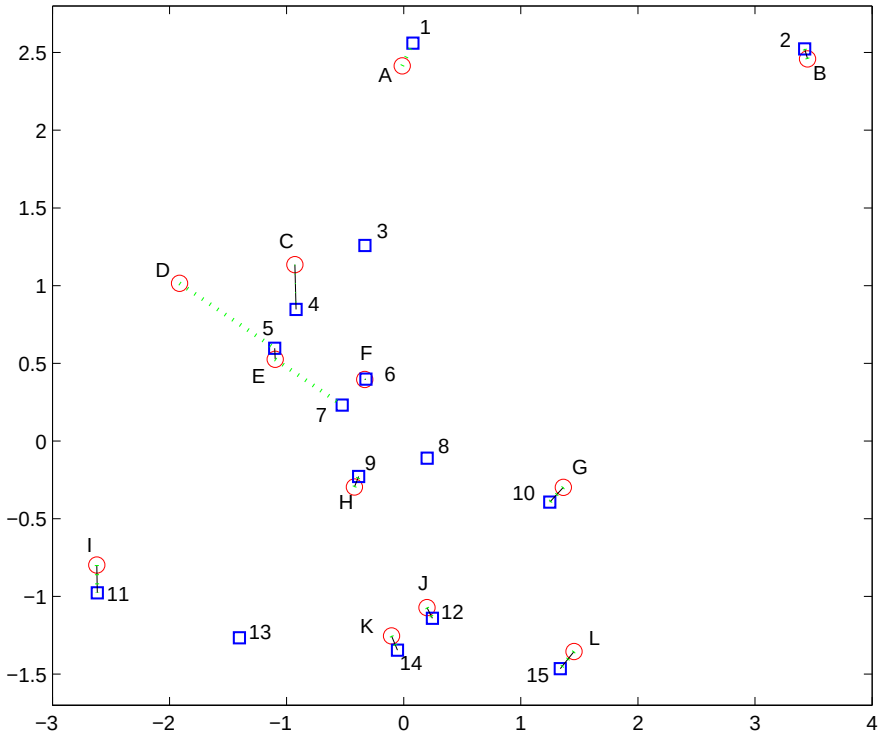

FIGURE 4. Example Class 1 scenario. The squares represent the 3-D points, projected onto the 2-D image plane using $\Omega_{computed}$. The circles represent the 2-D image points. The solid lines represent the truthful correspondences, and the dashed lines represent the computed correspondences.

using $\Omega_{computed}$ onto the 2-D image plane, and their truthful 2-D point and line correspondences. Again we note that aside from some slight deviations, we see the trend of our decomposition approach working quite well when the noise level is small, and degrading as the noise level increases. As is clear from the figures, for all three scenario classes, the decomposition approach performs quite well. Of note is that adding lines to the scenario seems to add significant complexity to the surface of the objective function, thus taking longer for the C-GRASP algorithm to run to completion and longer on average to locate approximate the correct solution, from a time and multi-start perspective.

6. CONCLUSIONS

In this paper, we have examined a problem from computer vision: the recognition of a 3-D object, represented by points and lines in an image, when the correspondence of points and lines is not known *a priori*. We have formulated this problem as a mixed-integer nonlinear optimization problem, explicitly accounting for possible noise in the 2-D data. For a

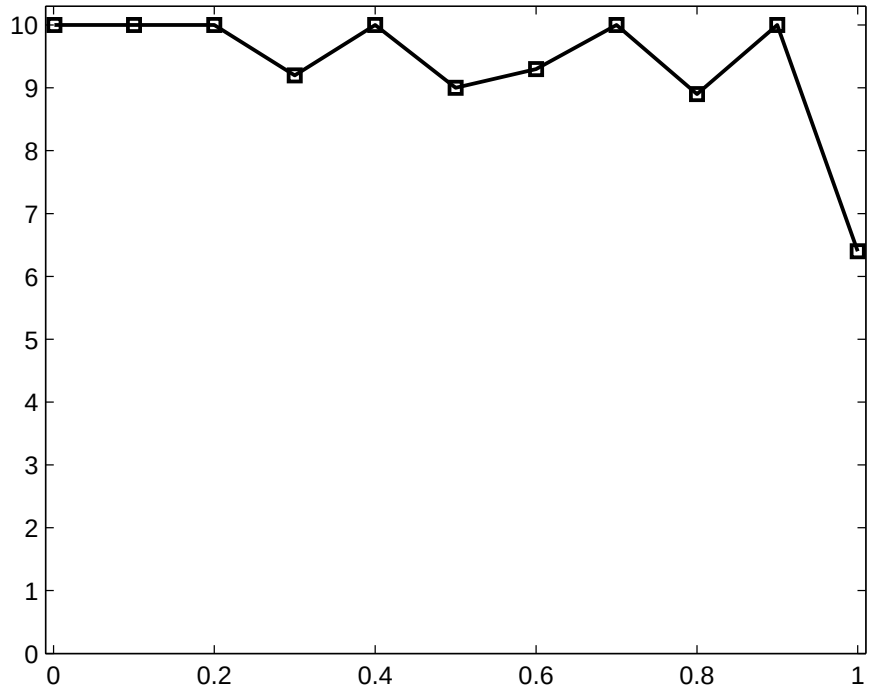

FIGURE 5. Class 1 average number of correct correspondences computed, as a function of noise level. Maximum possible correspondences for this class is 10.

pinhole camera (with specific internal calibration parameters), we decomposed the mixed-integer nonlinear optimization problem into a two step process of first determining the best projection matrix transforming the 3-D points and lines onto the 2-D picture and then using a linear assignment algorithm to determine the correspondences between the points and lines of the object and those of the image. Computational studies have shown that this approach does a very good job of determining the correct projection matrix and correspondences. Future research will be geared towards reducing the time of the C-GRASP heuristic, in an effort to make this approach suitable for a real-time system.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the editors and anonymous referees for their insightful comments and suggestions. This research was partially supported by grants from the U.S. Air Force and The Defense Threat Reduction Agency (DTRA) of the United States Department of Defense.

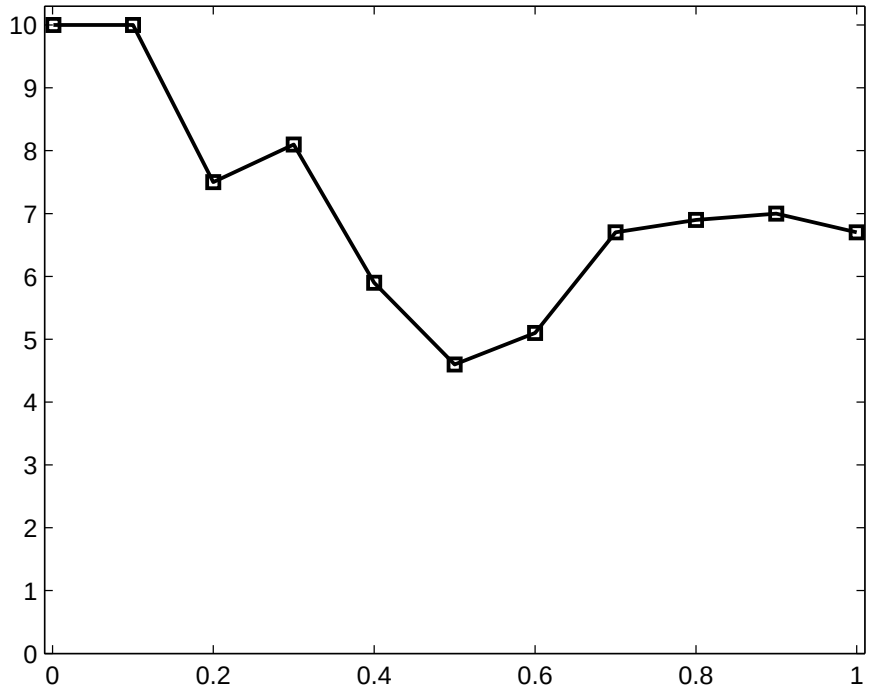

FIGURE 6. Class 2 average number of correct correspondences computed, as a function of noise level. Maximum possible correspondences for this class is 10.

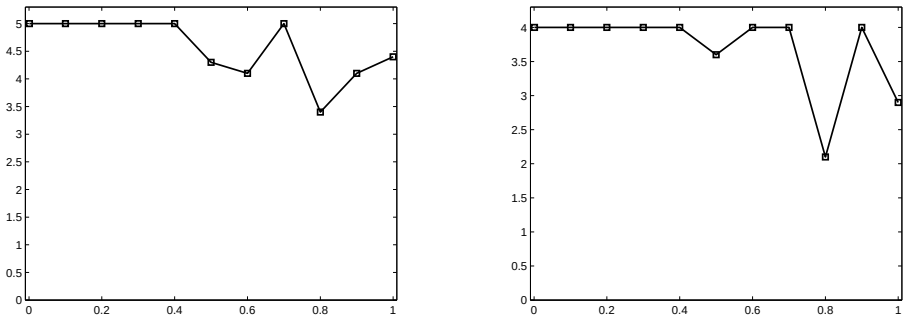

FIGURE 7. Class 3 average number of correct correspondences computed, as a function of noise level. The plot on the left is for the point correspondences, and the plot on the right is for the line correspondences. Maximum possible point and line correspondences for this class is 5 and 4, respectively.

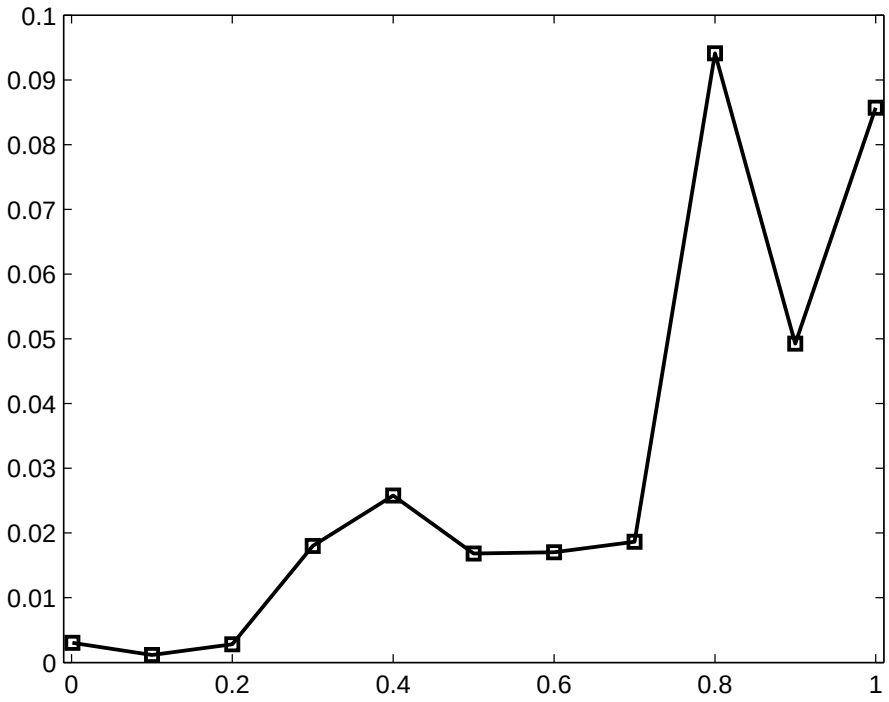

FIGURE 8. Class 1 average distance between 3-D points and lines, projected using $\Omega_{computed}$, and their truthful corresponding 2-D points and lines.

REFERENCES

C. G. E. Boender and H. E. Romeijn. Stochastic methods. In R. Horst and P. M. Pardalos, editors, *Handbook of Global Optimization*, pages 829–869. Kluwer Academic Publishers, 1995.

W.E. Boyce and R.C. DiPrima. *Calculus.* John Wiley & Sons, 1988.

T.A. Cass. Robust affine structure matching for 3D object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20:1265–1274, 1998.

J. Chai and S. D. Ma. Robust epipolar geometry estimation using genetic algorithm. *Pattern Recognition Letters*, 19:829–838, 1998.

Y. Cheng, V. Wu, R. T. Collins, A. R. Hanson, and E. M. Riseman. Maximum weight bipartite matching technique and its application to solve image feature matching. In *Proc. of the SPIE Conference on Visual Communication and Image Processing*, 1996.

P.A.M. Dirac. The physical interpretation of quantum mechanics. *Proceedings of the Royal Society of London*, A113:621–641, 1926.

O. Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint*. The MIT Press, 1993.

O. Faugeras and Q. T. Luong. *The Geometry of Multiple Images: The laws that govern the formation of multiple images of a scene and some of their applications*. The MIT Press,

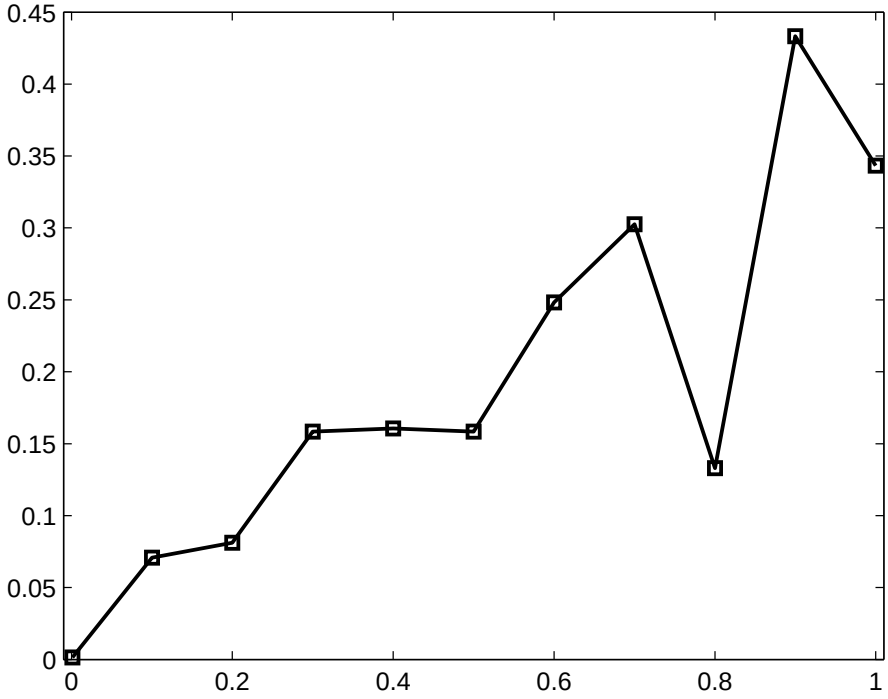

FIGURE 9. Class 2 average distance between 3-D points and lines, projected using $\Omega_{computed}$, and their truthful corresponding 2-D points and lines.

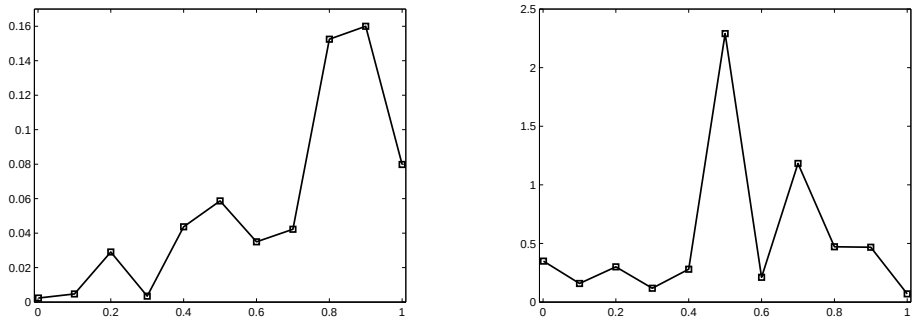

FIGURE 10. Class 3 average distance between 3-D points and lines, projected using $\Omega_{computed}$, and their truthful corresponding 2-D points and lines. The plot on the left is for the point correspondences, and the plot on the right is for the line correspondences.

2001.

T. A. Feo and M. G. C. Resende. A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters*, 8:67–71, 1989.

T. A. Feo and M. G. C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133, 1995.

M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24:381–395, 1981.

R. F. Gleeson, F. D. Grosshans, M. J. Hirsch, and R. M. Williams. Algorithms for the recognition of 2D images of *m* points and *n* lines in 3D. *Image and Vision Computing*, 21:497–504, 2003.

M.D. Greenberg. *Advance engineering mathematics*. Prentice Hall, 1988.

F. D. Grosshans. On the equations relating a three-dimensional object and its two-dimensional images. *Advances in Applied Mathematics*, 34:366–392, 2005.

R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

M. J. Hirsch. *GRASP-based heuristics for continuous global optimization problems*. PhD thesis, University of Florida, December 2006.

M.J. Hirsch, P.M. Pardalos, and M.G.C. Resende. Sensor registration in a sensor network by continuous GRASP. In *Proc. of the IEEE Military Communications Conference (MILCOM)*, 2006.

M.J. Hirsch, C.N. Meneses, P.M. Pardalos, M.A. Ragle, and M.G C.Resende. A continuous GRASP to determine the relationship between drugs and adverse reactions. In O. Seref, O. E. Kundakcioglu, and P.M. Pardalos, editors, *Data Mining, Systems Analysis, and Optimization in Biomedicine*, volume 953, pages 106–121. American Institute of Physics, 2007a.

M.J. Hirsch, C.N. Meneses, P.M. Pardalos, and M.G.C. Resende. Global optimization by continuous GRASP. *Optimization Letters*, 1:201–212, 2007b.

M.J. Hirsch, P.M. Pardalos, and M.G.C. Resende. Solving systems of nonlinear equations using continuous GRASP. *Nonlinear Analysis: Real World Applications*, 10:2000–2006, 2009.

M.J. Hirsch, P.M. Pardalos, and M.G.C. Resende. Speeding up continuous GRASP. *European Journal of Operational Research*, 205:507–521, 2010.

R. Jonker and A. Volgenant. A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing*, 39:325–340, 1987.

B. Liu and P. Wang. Recognition of 3D objects from 2D images – Some issues. *Lecture Notes in Computer Science*, 1121:995–998, 1996.

Y. Ma, S. Soatto, J. Kosecka, and S. S. Sastry. *An Invitation to 3-D Vision: From Images to Geometric Models*. Springer-Verlag, 2004.

G. L. Scott and H. C. Longuet-Higgins. An algorithm for associating the features of two images. *Proc. of the Royal Society of London* ***B***, 244:21–26, 1991.

J. G. Semple and G. T. Kneebone. *Algebraic Projective Geometry*. Oxford University Press, 1952.

L. S. Shapiro and J. M. Brady. Feature based correspondence: an eigenvector approach. *Image and Vision Computing*, 10(5):283–288, 1992.

M. Yi and P.S. Wang. 3D artificial objects recognition under virtual environment. In *Proc. of the International Conference on Pattern Recognition*, pages 1995–2004, 2000.

L. Zhan. Fast stochastic global optimization methods and their applications to cluster crystallization and protein docking. *PhD. thesis, University of Waterloo*, 2005.

(Michael J. Hirsch) RAYTHEON COMPANY, INTELLIGENCE AND INFORMATION SYSTEMS, 300 SENTINEL DRIVE, ANNAPOLIS JUNCTION, MD, 20701, USA.
E-mail address: mjh8787@ufl.edu

(Panos M. Pardalos) DEPARTMENT OF INDUSTRIAL AND SYSTEMS ENGINEERING, UNIVERSITY OF FLORIDA, 303 WEIL HALL, GAINESVILLE, FL, 32611, USA.
E-mail address: pardalos@ufl.edu

(Mauricio G. C. Resende) ALGORITHMS AND OPTIMIZATION RESEARCH DEPARTMENT, AT&T LABS RESEARCH, 180 PARK AVENUE, ROOM C241, FLORHAM PARK, NJ 07932 USA.
E-mail address: mgcr@research.att.com