

Quasiconvex Optimization for Robust Geometric Reconstruction

Qifa Ke and Takeo Kanade,
Computer Science Department, Carnegie Mellon University
{Qifa.Ke, tk}@cs.cmu.edu

Abstract

Geometric reconstruction problems in computer vision are often solved by minimizing a cost function that combines the reprojection errors in the 2D images. In this paper, we show that, for various geometric reconstruction problems, their reprojection error functions share a common and quasiconvex formulation. Based on the quasiconvexity, we present a novel quasiconvex optimization framework in which the geometric reconstruction problems are formulated as a small number of small-scale convex programs that are ready to solve. Our final reconstruction algorithm is simple and has intuitive geometric interpretation. In contrast to existing random sampling or local minimization approaches, our algorithm is deterministic and guarantees a predefined accuracy of the minimization result. We demonstrate the effectiveness of our algorithm by experiments on both synthetic and real data.

1 Introduction

Given measurements in 2D images, the goal of geometric reconstruction in computer vision is to estimate the three-dimensional information about the scene and the camera motions. Classical examples include triangulation [10], camera resectioning [4, 9], and structure from motion (see [6] for a review). The Gold standard for these estimation problems is minimizing F_s , the *average* of squared reprojection errors (model-fitting errors measured in 2D image domain). Minimizing F_s leads to maximum likelihood estimation when measurement noises follow Gaussian distribution.

Due to the camera perspective effect, the cost function F_s is highly nonlinear and often contains multiple local minima. Minimizing F_s is therefore difficult. Hartley and Schaffalitzky [5] proposed using the *pointwise maximum* of the squared reprojection errors as the cost function, which we denote as F_∞ . In contrast to F_s , it was shown that F_∞ contains only one single minimum value in its feasible domain. An approach using random line search in the parameter space was used in [5] to minimize F_∞ . The convergence behavior of random line search remains unclear. As pointed out in [5], it is difficult to perform random line search when the parameter space is high dimensional. Constrained minimization is also proposed in [5] for minimizing F_∞ . However, the constraints are nonlinear and nonconvex, making such constrained minimization a difficult problem by itself.

We can consider the model-fitting error as a function of the unknown parameters, which is termed *reprojection error function* in this paper. We show that the reprojection error functions share a *common* and *quasiconvex* formulation for the geometric reconstruction problems under our consideration. As a result, F_∞ , the pointwise maximum of a family of quasiconvex functions, is also a quasiconvex function. We then present an one-dimensional bisection algorithm to minimize the quasiconvex function F_∞ . Our algorithm consists of a small number of small-scale convex programs, specifically linear programs (LP) or second-order cone programs (SOCP). Both LP and SOCP are well-studied and existing efficient algorithms and implementations are ready to use. Compared to random line search in parameter space or local minimization approaches, our minimization approach is efficient, even when the unknowns are high dimensional. More importantly, our approach is deterministic and guarantees a predefined accuracy of the minimization result.

It has been pointed out in [5] that F_∞ is sensitive to outliers. To handle outliers, we use F_m , the pointwise m -th smallest reprojection error, as the cost function. In contrast to F_∞ or F_s , the cost function F_m is highly robust to outliers [14]. In spite of its complex formulation, in our cases F_m is still a pointwise operator of a family of quasiconvex functions. As a result, our algorithm to minimize F_∞ can be extended to efficiently minimize F_m , again by solving small-scale convex programs (LP or SOCP).

1.1 Background: geometric reconstruction problems

We present four classical examples of geometric reconstruction problems in computer vision.

1.1.1 Multi-view triangulation

We are given projection matrices of N cameras, denoted by $\{P_i, i = 1, \dots, N\}$, and the images of the unknown 3D point \mathbf{Z} in these N cameras, denoted by $\{\mathbf{x}_i, i = 1, \dots, N\}$. The task of triangulation is to estimate \mathbf{Z} from $\{P_i\}$ and $\{\mathbf{x}_i\}$. Triangulation is a necessary step in two- or multi-view 3D reconstruction, and in structure from motion.

Note that optimal triangulation algorithms [10, 8] for two-view case are not generalizable to multi-view case.

1.1.2 Camera resectioning

We are given 3D points $\{\mathbf{Z}_i, i = 1, \dots, N\}$ and their images $\{\mathbf{x}_i, i = 1, \dots, N\}$ in one camera. The task is to estimate

the camera projection matrix P from these N corresponding pairs $\{\mathbf{x}_i \leftrightarrow \mathbf{Z}_i\}$. Camera resectioning is used in camera calibration and in structure from motion.

1.1.3 Multi-view reconstruction with known rotations

In some cases the camera rotations are known, leaving only the camera positions and the 3D of the scene to be estimated [5]. For example, in vision-aided inertial navigation, accurate camera pose is available from modern gyroscopes, while the camera position information from accelerometers is still noisy [3]. Another example is that there are reconstruction methods in which the camera rotation for each frame is estimated in a first step [13].

Denote the N intrinsically calibrated cameras as $\{P_i = (R_i, -R_i C_i), i = 1, \dots, N\}$, where for each camera the rotation R_i is known, but its 3D position C_i is unknown. We are given 2D feature points $\{\mathbf{x}_{ij}\}$ over the N cameras. Here \mathbf{x}_{ij} denotes the projection of j -th 3D point \mathbf{Z}_j onto the i -th camera. The task is to estimate $\{\mathbf{Z}_j\}$ and $\{C_i\}$ from the 2D points $\{\mathbf{x}_{ij}\}$ and the camera poses $\{R_i\}$.

1.1.4 Planar homography estimation

Two images of points on a 3D scene plane are related by a planar homography H , a 3×3 non-singular matrix. Given N correspondences $\{\mathbf{x}_i \leftrightarrow \mathbf{x}'_i, i = 1, \dots, N\}$, the task is to estimate H such that $\mathbf{x}'_i = H\mathbf{x}_i$.

2 The cost function

In this section, we define the reconstruction error metric at each individual 2D measurement, and the cost functions that combine reconstruction errors from individual 2D measurements.

2.1 Error metric for one 2D measurement

We use Triangulation as an example to illustrate three often-used error metrics for an individual 2D measurement.

2.1.1 Algebraic distance

Denote $\tilde{\mathbf{x}}_i = (\mathbf{x}_i; 1)$ the homogeneous coordinates of the 2D measurement \mathbf{x}_i , we have the following linear equation:

$$k_i \tilde{\mathbf{x}}_i = P_i \mathbf{Z} \quad (1)$$

Here \mathbf{Z} is also expressed in homogeneous coordinates. The algebraic distance for \mathbf{x}_i is then defined by:

$$f_i(\mathbf{Z}) = \|k_i \tilde{\mathbf{x}}_i - P_i \mathbf{Z}\|_2 \quad (2)$$

Linear least-squares can be applied to estimate \mathbf{Z} by minimizing the sum of squared algebraic distances. Since the algebraic distance is not geometrically or statistically meaningful, the algebraic reconstruction is not reliable (see [14, 5]).

2.1.2 Distance in 3D space

In the case of calibrated cameras, $f_i(\mathbf{Z})$ can be defined as the distance from the 3D point \mathbf{Z} to the ray back-projected from \mathbf{x}_i . In the case of two views, this distance function




Figure 1. Distance between $\mathbf{x} = (u, v)$ and $\hat{\mathbf{x}} = (\hat{u}, \hat{v})$. The solid square shows the contour on which the L_1 norm error $e_1 = |e_u| + |e_v| = \alpha$, while the dash line shows the contour on which the L_2 norm error $e_2 = \sqrt{e_u^2 + e_v^2} = \alpha$. Here $e_u = (u - \hat{u})$, and $e_v = (v - \hat{v})$.

leads to the midpoint estimation method where \mathbf{Z} is given by midpoint of the perpendicular between the two rays.

When a camera is further away from the 3D point \mathbf{Z} , the camera has larger uncertainty on \mathbf{Z} . Distance metric in 3D space cannot take such uncertainty into account. As a result, the reconstruction result is unstable when the 3D point is far away from cameras.

2.1.3 Reprojection error in the image

The reprojection error is defined as the distance in the 2D image domain between \mathbf{x}_i and its reprojection $\hat{\mathbf{x}}_i = \pi_i(\mathbf{Z})$:

$$f_i(\mathbf{Z}) = \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_l = \|\mathbf{x}_i - \pi_i(\mathbf{Z})\|_l \quad (3)$$

where $\hat{\mathbf{x}}_i = \pi_i(\mathbf{Z})$ is the reprojection of \mathbf{Z} in the image of camera P_i , and $\|\cdot\|_l$ denotes some vector norm. Both \mathbf{x}_i and $\hat{\mathbf{x}}_i$ are in 2D Cartesian coordinates.

We choose reprojection error metric since it has a well-defined geometric meaning and it leads to maximum likelihood estimation. For example, when L_2 norm is used in Eq. (3), the reprojection error f_i is the *Euclidean distance* between \mathbf{x}_i and $\hat{\mathbf{x}}_i$. We can also use L_1 norm. Its geometric meaning is shown in Fig. 1.

2.2 Generalized reprojection error function

Definition 1. The general formulation of reprojection error function:

$$f(\mathbf{X}) = \frac{p(\mathbf{X})}{q(\mathbf{X})} \quad (4)$$

where

- $\mathbf{X} \in \mathbb{R}^n$ is the unknown vector to be estimated;
- $p(\mathbf{X})$ is a convex function, and $p(\mathbf{X}) \geq 0$.
- $q(\mathbf{X})$ is a linear function, and $q(\mathbf{X}) > 0$;

In the following we show that most reprojection error functions are special cases of the above general formulation. For a geometric reconstruction problem, if its reprojection error function conforms to the general formulation in Eq. (4), the algorithms we present in this paper can be applied to solve such reconstruction problem.

2.2.1 Reprojection error function in the image plane

Result 1. For the reconstruction problems in Section 1.1, the reprojection error function defined in the image domain conforms to the general formulation in Definition 1.

Proof. For the problems in Section 1.1, the reprojection of $\mathbf{x} = (u, v)$ in the image can be written as:

$$\hat{\mathbf{x}} = \begin{pmatrix} \mathbf{a}^\top \mathbf{X} & \mathbf{b}^\top \mathbf{X} \\ \mathbf{c}^\top \mathbf{X} & \mathbf{c}^\top \mathbf{X} \end{pmatrix}^\top \quad (5)$$

Here \mathbf{X} is the vector to be estimated. \mathbf{a} , \mathbf{b} , and \mathbf{c} are known vectors. For example, in the triangulation problem, they are the three rows of the camera matrix \mathbf{P} , respectively.

The reprojection error function is:

$$f(\mathbf{X}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_l = \left\| \frac{1}{q(\mathbf{X})} (p_u(\mathbf{X}), p_v(\mathbf{X})) \right\|_l, \quad (6)$$

where $\|\cdot\|_l$ is the vector norm, and

$$\begin{aligned} p_u(\mathbf{X}) &= (\mathbf{u}\mathbf{c}^\top - \mathbf{a}^\top)\mathbf{X}, \\ p_v(\mathbf{X}) &= (\mathbf{v}\mathbf{c}^\top - \mathbf{b}^\top)\mathbf{X}, \\ q(\mathbf{X}) &= \mathbf{c}^\top \mathbf{X}. \end{aligned} \quad (7)$$

It is obvious that $q(\mathbf{X})$ is a linear function of \mathbf{X} .

In this paper, we consider affine or Euclidean reconstruction¹. The cheirality constraint (see [6]), which states that the 3D points visible in the image must be in front of the camera, can then be expressed as $\mathbf{c}^\top \mathbf{X} > 0$ ². Therefore, we have $q(\mathbf{X}) > 0$. The reprojection error function in Eq. (6) can then be rewritten as:

$$f(\mathbf{X}) = \frac{1}{q(\mathbf{X})} \|(p_u(\mathbf{X}), p_v(\mathbf{X}))\|_l = \frac{p(\mathbf{X})}{q(\mathbf{X})} \quad (8)$$

Any norm function $g(\mathbf{y}) = \|\mathbf{y}\|_l$ is a convex function of \mathbf{y} . The function $h(\mathbf{X}) = (p_u(\mathbf{X}), p_v(\mathbf{X}))$ is an affine function of \mathbf{X} . The composition of a convex function g and an affine function h , denoted by $g \circ h$, is a convex function. Therefore, $p(\mathbf{X}) = (g \circ h)(\mathbf{X})$ is a convex function of \mathbf{X} . It is obvious that $p(\mathbf{X}) \geq 0$. \square

When uncertainty on the location of each 2D feature point is available, it can be shown that the uncertainty-weighted reprojection error function still conforms to the general formulation in Definition 1.

2.2.2 Angular reprojection error function

When the camera is calibrated, the angle θ between the observed ray \mathbf{x} and the reprojection ray $\mathbf{r} = (\mathbf{a}, \mathbf{b}, \mathbf{c})^\top \mathbf{X}$ can be used to define the reprojection error [8, 5]:

$$f(\mathbf{X}) = |\tan(\theta)| = \left| \frac{\mathbf{x} \times \mathbf{r}}{\mathbf{x}^\top \mathbf{r}} \right| \quad (9)$$

where \times denotes cross-product. We choose $\tan(\theta)$ since it is a monotonically-increasing function of θ when $\theta \in [0, \pi/2)$. The cheirality constraint can be enforced by $|\theta| < \pi/2$, which leads to $q(\mathbf{X}) = \mathbf{x}^\top \mathbf{r} > 0$. It is easy to verify

¹In a way similar to the method briefed in [5], our algorithm in this paper can be extended to projective reconstruction.

²In planar homography estimation, the cheirality constraint $\mathbf{p}_3'^\top \mathbf{X} > 0$ can be rewritten as $\mathbf{h}_3^\top \mathbf{x} > 0$, by using the following facts: 1) \mathbf{X} is on a 3D plane; 2) homography $\mathbf{H} = \mathbf{A} - \mathbf{b}\mathbf{v}^\top$, where $\mathbf{P}' = [\mathbf{A}|\mathbf{b}]$ is the second camera. Here $\mathbf{p}_3'^\top$ is the third row of \mathbf{P}' , and \mathbf{h}_3^\top the third row of \mathbf{H} .

that $q(\mathbf{X})$ is a linear function of \mathbf{X} , and $p(\mathbf{X}) = |\mathbf{x} \times \mathbf{r}|$ is convex in \mathbf{X} . Therefore, the angular reprojection error function $f(\mathbf{X}) = \frac{p(\mathbf{X})}{q(\mathbf{X})}$ conforms to the general form in Definition 1.

2.3 Combining reprojection errors into cost function

The often used cost function F_s in geometric reconstruction is defined as the average of the squared L_2 -norm reprojection errors:

$$F_s = \frac{1}{M} \sum_i f_i^2(\mathbf{X}) \quad (10)$$

where M is the total number of 2D measurements (points). F_s is difficult to minimize as it is highly nonlinear and contains multiple local minima [5].

Hartley and Schaffalitzky [5] proposed using the pointwise maximum of the reprojection errors as the cost function:

$$F_\infty(\mathbf{X}) = \max_i f_i(\mathbf{X}) \quad (11)$$

It was shown in [5] that $F_\infty(\mathbf{X})$ contains only one single minimum value in its domain, and is therefore easier to minimize than $F_s(\mathbf{X})$. But as is also pointed out in [5], $F_\infty(\mathbf{X})$ is sensitive to outliers.

To deal with the outliers, we propose using the *pointwise* m -th smallest reprojection errors as the cost function:

$$F_m(\mathbf{X}) = \text{mth}_i f_i(\mathbf{X}) \quad (12)$$

It is obvious that F_∞ is a special case of F_m when $m = N$. F_m is a highly robust function. For example, when $m = \lfloor N/2 \rfloor$, it is the median operator. Minimizing F_m leads to least-median optimization [14], which can handle noisy measurements with up to 50% of outliers.

3 Minimizing the cost function

Both F_∞ and F_m are constructed from pointwise operations on a family of functions. They are not differentiable at many points. As a result, classical gradient-based approaches are not applicable to minimizing them. Random line search in the parameter space was proposed in [5] to minimize F_∞ , and random sampling [14, 11] is often used to detect outliers and to minimize F_m . These randomized approaches are not scalable when the unknowns are high-dimensional. They do not guarantee convergence either.

In this section, we show that the general reprojection error function (Definition 1) is quasiconvex. Such quasiconvexity enables us to design a deterministic and efficient algorithm to minimizing F_∞ and F_m .

3.1 Minimization by feasibility

Instead of random search or sampling, let us look at a minimization approach that uses the classic bisection search in the range domain of F_∞ and F_m .

For the vision problems in which we are interested, the image size is bounded. Therefore, it is realistic to assume

Algorithm: $\min_{\mathbf{X}} F(\mathbf{X})$.

- 1: Given $l \leq F^*$, $h \geq F^*$, and the tolerance $\varepsilon > 0$.
- 2: **while** $(h - l) > \varepsilon$ **do**
- 3: $\alpha = (h + l)/2$.
- 4: Solve the feasibility problem (14).
- 5: **if** (14) is feasible, **then** $h = \alpha$;
- 6: **else** $l = \alpha$.
- 7: **end while**

Figure 2. Classical bisection algorithm to pin down the optimal value by searching in the one-dimensional range domain.

that $l \leq F(\mathbf{X}) \leq h$, where $F(\mathbf{X})$ is the cost function. For $\alpha \in [l, h]$, denote S_α the α -sublevel set of $F(\mathbf{X})$:

$$S_\alpha = \{\mathbf{X} \mid F(\mathbf{X}) \leq \alpha\} \quad (13)$$

If S_α is non-empty, then we know that F^* , the minimum value of $F(\mathbf{X})$, satisfies $F^* \leq \alpha$. Otherwise, we have $F^* > \alpha$. Determining whether S_α is empty or not can be achieved by solving the following feasibility problem:

$$\begin{aligned} &\text{find } \mathbf{X} \\ &s.t. \mathbf{X} \in S_\alpha \end{aligned} \quad (14)$$

Based on the above observation, we can use the bisection algorithm (see [2]) to pin down the optimal value of $F(\mathbf{X})$ by solving a sequence of feasibility problems. Fig. 2 shows the basic procedure of the algorithm. It starts with a range $[l, h]$ that is known to contain F^* . Then we solve the feasibility problem at its mid-point $\alpha = (l + h)/2$. If it is feasible, then the optimal value F^* is in the lower half of the interval and we can shrink $[l, h]$ to $[l, \alpha]$. Otherwise, F^* must be in the upper half of the interval and we shrink $[l, h]$ to $[\alpha, h]$. The algorithm then continues on the identified half of the interval.

As we can see, at each iteration the range is shrunk by half, and the bisection algorithm is guaranteed to converge in $\lceil \log_2((h - l)/\varepsilon) \rceil$ iterations. For example, $[0, 100]$ allows the re-projection error to be as many as 100 pixels, which is guaranteed to contain the optimal value F^* . If we choose $\varepsilon = 0.5$ pixel, the algorithm will converge in only $\lceil \log_2 200 \rceil = 8$ iterations. Note that the number of iterations is independent of the dimension of the unknown \mathbf{X} , indicating that the algorithm is suitable for solving high dimensional problems. More importantly, the optimal value we derive is guaranteed to be less than $\varepsilon = 0.5$ pixel away from the true minimum value.

3.2 Quasiconvex functions

The bisection algorithm in Fig. 2 is simple, deterministic, and it converges in a small number of iterations. It can even be applied to minimizing cost functions with multiple minima. The *critical step* in the algorithm is solving the feasibility problem in Eq. (14), which could be a hard problem by itself if the α -sublevel set of $F(\mathbf{X})$ is complicated. How-




Figure 3. A quasiconvex function. All of its α -sublevel sets $\{S_\alpha\}$ are convex. But this quasiconvex function is not convex, as can be seen from the line segment \overline{AB} that lies below the function.

ever, if S_α is convex, then Eq. (14) is a convex feasibility problem [2] that can be solved efficiently. A function with such convex α -sublevel set is called a *quasiconvex* function:

Definition 2. (see [2]) A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is **quasiconvex** if its domain $\text{dom}(f)$ and all its sublevel sets

$$S_\alpha = \{x \in \text{dom}(f) \mid f(x) \leq \alpha\},$$

for $\alpha \in \mathbb{R}$, are convex.

A convex function has convex sublevel sets, and therefore, is quasiconvex. The reverse is not true in general. Fig. 3 shows an example of quasiconvex function that is not convex. The dash-line segment that lies below the function indicates the non-convexity of the function.

The reprojection error functions are not convex due to camera perspective effect, but they are quasiconvex:

Result 2. A reprojection error function that conforms to the general form defined in Eq. (4) is a quasiconvex function.

Proof. For any $\alpha > 0$, the α -sublevel set of $f(\mathbf{X})$ is:

$$\begin{aligned} S_\alpha &= \{\mathbf{X} \mid f(\mathbf{X}) \leq \alpha\} \\ &= \{\mathbf{X} \mid p(\mathbf{X}) - \alpha q(\mathbf{X}) \leq 0, q(\mathbf{X}) > 0\} \end{aligned}$$

From the definition of the general reprojection error function (Definition 1), we know that $p(\mathbf{X})$ is a convex function, and $-\alpha q(\mathbf{X})$ is a linear function and, therefore, a convex function. The sum of these two convex functions $\phi(\mathbf{X}) = p(\mathbf{X}) - \alpha q(\mathbf{X})$ is still a convex function. A sublevel set of a convex function is a convex set. As a result, S_α is a convex set since it is the intersection of two convex sets: the zero sublevel set of $\phi(\mathbf{X})$, and the half space defined by $q(\mathbf{X}) > 0$. Since $\text{dom}(f) = \mathbb{R}^n$ and S_α are all convex, we conclude that $f(\mathbf{X})$ is quasiconvex. \square

3.3 Minimizing cost function F_∞

Result 3. $F_\infty(\mathbf{X})$, the pointwise maximum of quasiconvex reprojection error functions $f_i(\mathbf{X})$, is also quasiconvex.

Proof. The α -sublevel set S_α of $F_\infty(\mathbf{X})$ is:

$$\begin{aligned} S_\alpha &= \{\mathbf{X} \mid \max_i f_i(\mathbf{X}) \leq \alpha\} \\ &= \{\mathbf{X} \mid f_i(\mathbf{X}) \leq \alpha, i = 1, 2, \dots, N\} \\ &= \bigcap_{i=1}^N S_\alpha^i \end{aligned}$$

Here S_α^i is the α -sublevel set of the reprojection error function $f_i(\mathbf{X})$. From Result 2, we know that $\{S_\alpha^i\}$ are all convex sets. As a result, their intersection S_α is also a convex set. Therefore, $F_\infty(\mathbf{X})$ is a quasiconvex function. \square

Due to its quasiconvexity, F_∞ can be efficiently minimized by the bisection algorithm in Fig. 2. The convex set $S_\alpha = \bigcap_{i=1}^N S_\alpha^i$ can be expressed as:
 $S_\alpha = \{\mathbf{X} \mid q_i(\mathbf{X}) > 0; p_i(\mathbf{X}) - \alpha q_i(\mathbf{X}) \leq 0; i = 1, \dots, N\}$
The feasibility problem of the bisection algorithm in Eq. (14) can now be solved by the following *convex* program:

$$\begin{aligned} \min_{\mathbf{X}, \gamma} \quad & (15) \\ \text{s.t.} \quad & -q_i(\mathbf{X}) + \epsilon \leq \gamma, \\ & p_i(\mathbf{X}) - \alpha q_i(\mathbf{X}) \leq \gamma, \\ & i = 1, \dots, N. \end{aligned}$$

Here ϵ is a small positive number. Denote γ^* the optimal value of (15). If $\gamma^* \leq 0$, then S_α of $F_\infty(\mathbf{X})$ is nonempty, and the problem in (14) is feasible; otherwise (14) is infeasible. Note that we do not need to solve (15) with high accuracy. The algorithm terminates whenever $\gamma \leq 0$ is satisfied, or whenever a dual feasible point is found with positive dual objective (which means $\gamma^* > 0$).

3.4 Minimizing robust cost function F_m

$F_\infty(\mathbf{X})$ is sensitive to outliers [5]. To deal with outliers, we use the robust cost function $F_m(\mathbf{X})$, which is defined as the m -th smallest reprojection error (see Eq. (12)). $F_m(\mathbf{X})$ is not a quasiconvex function, except for $m = N$, in which case F_m becomes F_∞ .

However, since $F_m(\mathbf{X})$ is a pointwise function of a family of quasiconvex functions $\{f_i(\mathbf{X})\}$, its α -sublevel set can still be represented by the convex sublevel sets of these quasiconvex functions. As a result, we are able to extend the bisection algorithm to efficiently minimize F_m .

3.4.1 The α -sublevel set of F_m

A point \mathbf{X}_0 belongs to the α -sublevel set of $F_m(\mathbf{X})$ if and only if there exists a group of m α -sublevel sets whose intersection contains the point \mathbf{X}_0 .

Result 4. Denote S_α the α -sublevel set of $F_m(\mathbf{X})$. For any \mathbf{X}_0 , $\mathbf{X}_0 \in S_\alpha$ if and only if $\mathbf{X}_0 \in_m \{S_\alpha^1, S_\alpha^2, \dots, S_\alpha^N\}$. Here S_α^i is the α -sublevel set of $f_i(\mathbf{X})$. The symbol \in_m means that there exist m sublevel sets in $\{S_\alpha^1, S_\alpha^2, \dots, S_\alpha^N\}$ such that \mathbf{X}_0 is inside the intersection of these m sublevel sets.

Proof. For any \mathbf{X}_0 , we sort the N reprojection errors

$$f_1(\mathbf{X}_0), f_2(\mathbf{X}_0), \dots, f_N(\mathbf{X}_0)$$

into the nondecreasing order

$$f_{(1)}(\mathbf{X}_0) \leq \dots \leq f_{(m)}(\mathbf{X}_0) \leq \dots \leq f_{(N)}(\mathbf{X}_0) \quad (16)$$

For the necessary condition, if $\mathbf{X}_0 \in S_\alpha$, then we have $F_m(\mathbf{X}_0) = f_{(m)}(\mathbf{X}_0) \leq \alpha$. The first m smallest reprojection errors $\{f_{(i)}(\mathbf{X}_0), i = 1, \dots, m\}$ in Eq. (16) must therefore satisfy $f_{(i)}(\mathbf{X}_0) \leq \alpha$. As a result, \mathbf{X}_0 belongs to the intersection of the m α -sublevel sets of the first m functions in Eq. (16).

For the sufficient condition, suppose \mathbf{X}_0 is in the intersection of the following m sublevel sets: $\{S_\alpha^{(i)}, i = 1, \dots, m\}$, where $S_\alpha^{(i)}$ is the α -sublevel set of $f^{(i)}$. We must have:

$$f^{(i)}(\mathbf{X}_0) \leq \alpha, \quad i = 1, \dots, m \quad (17)$$

Now if $F_m(\mathbf{X}_0) = f_{(m)}(\mathbf{X}_0) > \alpha$, then from the sorted sequence in Eq. (16) we know that the number of less-than- α reprojection errors is less than m . This contradicts Eq. (17) where there are m less-than- α reprojection errors. Therefore we have $F_m(\mathbf{X}_0) \leq \alpha$, i.e., $\mathbf{X}_0 \in S_\alpha$. \square

3.4.2 Feasibility by convex program

From Result 4, the feasibility problem in the bisection algorithm to minimizing F_m can be rewritten as:

$$\begin{aligned} \text{find } \mathbf{X} \quad & (18) \\ \text{s.t. } \mathbf{X} \in_m \{S_\alpha^1, S_\alpha^2, \dots, S_\alpha^N\} \end{aligned}$$

In other words, we need to determine if there exist m α -sublevel sets whose common intersection is non-empty. A straightforward approach is to check the feasibility of every possible group of m sublevel sets, where for each group its feasibility can be exactly determined by the convex program of Eq. (15). In worst case, this requires $\binom{N}{m}$ convex programs to solve Eq. (18), which is good for small N . When N is large, we use the following single convex program to determine the feasibility problem in Eq. (18):

Result 5. Denote $\gamma^* = (\gamma_1^*, \gamma_2^*, \dots, \gamma_N^*)$ the optimal value of the following convex program achieving at \mathbf{X}^* :

$$\begin{aligned} \min_{\mathbf{X}, \gamma} \quad & \gamma_1 + \gamma_2 + \dots + \gamma_N \quad (19) \\ \text{s.t.} \quad & -q_i(\mathbf{X}) + \epsilon \leq \gamma_i, \\ & p_i(\mathbf{X}) - \alpha q_i(\mathbf{X}) \leq \gamma_i, \\ & \gamma_i \geq 0, \\ & i = 1, \dots, N. \end{aligned}$$

Here ϵ is a small positive number. Denote g the number of zero elements in γ^* . If $g \geq m$, then the problem defined by Eq. (18) must be feasible; otherwise we consider Eq. (18) infeasible.

γ_i^* is called the infeasibility of $f_i(\mathbf{X}^*)$. For any sublevel set S_α^i , if its corresponding infeasibility $\gamma_i^* = 0$, then \mathbf{X}^* is inside S_α^i . As a result, the condition $g \geq m$ is sufficient for Eq. (18) to be feasible, since these g sublevel sets contain the common point \mathbf{X}^* .

While $g \geq m$ is a sufficient condition, it is an approximated necessary condition for Eq. (18) to be feasible. The




Figure 4. When the camera matrix is normalized appropriately, the infeasibility γ_i^* is the distance from \mathbf{X}^* to the α -convex cone. For comparison purpose, d_i is the distance to the ray back-projected from 2D measurement \mathbf{x}_i .

exact conclusion about the infeasibility of Eq. (18) requires checking the feasibility of $\binom{N}{m}$ groups of m sublevel sets, or using integer programming to find the optimal point \mathbf{X}^* that minimizes the number of infeasibilities (the number of nonzero components in γ^*). Result 5 finds the minimum sum-of-infeasibilities $\sum_i \gamma_i^*$, and uses it to approximate the minimum number of infeasibilities. With such approximation, the bisection algorithm gives an upper bound on the true minimum value of F_m . The sum of infeasibility $\|\gamma\|_1 = \sum_i \gamma_i$ is by itself a robust metric (L_1 norm is a robust metric), especially in our cases where the magnitude of outliers in the 2D measurements is bound by the image size. As a result, the bisection algorithm using Result 5 can usually achieve a tight upper bound on the true minimum value of F_m . We can further improve the result by using weighted sum of infeasibility $\mathbf{w}^\top \gamma$ in Eq. (19), where $\mathbf{w} = (w_1, \dots, w_N)$ is the weight for each measurement, and $w_i \in [0, 1]$ can be set according to its corresponding reprojection error to down-weight outliers.

Fig. 4 illustrates the meaning of γ_i^* . When the camera matrix is *normalized* appropriately, the infeasibility γ_i^* is the distance from \mathbf{X}^* to the convex cone if \mathbf{X}^* is outside the cone. If \mathbf{X}^* is inside the cone S_α^i , then $\gamma_i^* = 0$. The sum-of-infeasibilities $\sum_i \gamma_i^*$ is therefore the sum of distances from \mathbf{X}^* to the convex cones that do not contain \mathbf{X}^* . When \mathbf{X}^* goes further away from the camera C_i , the camera C_i has larger uncertainty on \mathbf{X}^* . Such varying uncertainty is taken into account by γ_i^* as it is the distance to the cone, and the cone becomes larger as \mathbf{X}^* goes further away from the camera C_i . This is in contrast to the distance to back-projected ray in 3D space (see Fig. 4).

3.5 Feasibility by LP or SOCP

When L_1 - or L_2 -norm error metric is used in defining the reprojection error function, the convex program for feasibility becomes small-scale linear programs (LP) or second-order convex programs (SOCP), respectively.

3.5.1 L_1 -norm error metric leads to LP

When L_1 -norm error metric is used, the convex program in Eq. (19) becomes the following linear program:




Figure 5. Geometric illustration of 3D reconstruction using convex feasibility. The algorithm seeks the minimum cone size with which at least m cones have non-empty intersection.

$$\begin{aligned} \min_{\mathbf{X}, \gamma} \quad & \gamma_1 + \gamma_1 + \dots + \gamma_N \\ \text{s.t.} \quad & -q_i(\mathbf{X}) + \epsilon \leq \gamma_i, \\ & -\alpha q_i(\mathbf{X}) + p_{ui}(\mathbf{X}) - p_{vi}(\mathbf{X}) \leq \gamma_i, \\ & -\alpha q_i(\mathbf{X}) + p_{ui}(\mathbf{X}) + p_{vi}(\mathbf{X}) \leq \gamma_i, \\ & -\alpha q_i(\mathbf{X}) - p_{ui}(\mathbf{X}) - p_{vi}(\mathbf{X}) \leq \gamma_i, \\ & -\alpha q_i(\mathbf{X}) - p_{ui}(\mathbf{X}) + p_{vi}(\mathbf{X}) \leq \gamma_i, \\ & \gamma_i \geq 0, \quad i = 1, \dots, N. \end{aligned} \quad (20)$$

Here p_{ui} , p_{vi} , and q_i are all linear functions of \mathbf{X} (see Eq. (7) for the definition).

3.5.2 L_2 -norm error metric leads to SOCP

When L_2 -norm error metric is used, Eq. (19) becomes:

$$\begin{aligned} \min_{\mathbf{X}, \gamma} \quad & \gamma_1 + \gamma_1 + \dots + \gamma_N \\ \text{s.t.} \quad & -q_i(\mathbf{X}) + \epsilon \leq \gamma_i, \\ & \|\mathbf{A}_i \mathbf{X}\|_2 \leq \alpha q_i(\mathbf{X}) + \gamma_i, \\ & \gamma_i \geq 0, \quad i = 1, \dots, N. \end{aligned} \quad (21)$$

Here

$$\mathbf{A}_i = \begin{pmatrix} u_i \mathbf{c}_i^\top - \mathbf{a}_i^\top \\ v_i \mathbf{c}_i^\top - \mathbf{b}_i^\top \end{pmatrix}$$

is a 2×3 matrix, and \mathbf{a} , \mathbf{b} , and \mathbf{c} are known vectors (see Eq. (5) for the notation). $\alpha q_i(\mathbf{X}) + \gamma_i$ is a linear function of \mathbf{X} . Therefore, the inequality

$$\|\mathbf{A}_i \mathbf{X}\|_2 \leq \alpha q_i(\mathbf{X}) + \gamma_i$$

defines a second order convex cone [2]. As a result, Eq. (21) is a second-order cone programming (SOCP).

3.6 Geometric interpretation

The minimization algorithm we presented in this section has intuitive geometric interpretation. We use multi-view triangulation as an example to illustrate. For each 2D measurement, the camera optical center and the six linear inequalities in Eq.(20) form a convex cone S_α^i in front of the camera in the 3D space, as shown in Fig. 5. The cone size is determined by α . For any point inside the convex cone S_α^i , its reprojection error must be less than α . If the common intersection of at least m convex cones is not empty, then we conclude that there exists at least one point \mathbf{X}_0 in the 3D space such that the cost function $F_m(\mathbf{X}_0) \leq \alpha$. Minimizing $F_m(\mathbf{X})$ is therefore equivalent to adjusting α , the size of

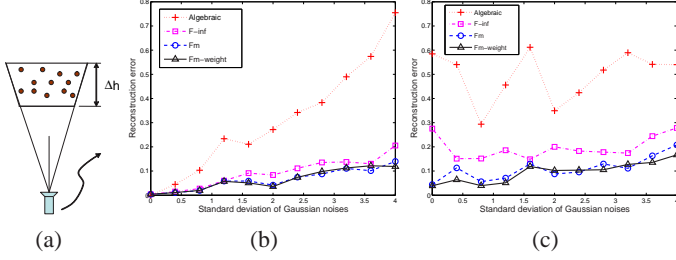


Figure 6. Multi-view triangulation: synthetic data. (a): The camera is rotating and translating, imaging a 3D scene consists of 40 points; (b): Reconstruction errors (normalized by Eq. (22)) with zero-mean Gaussian noises added to 2D point coordinates; (c): Reconstruction errors under both Gaussian noises and outliers.

the convex cone, until we find the minimum α with which the intersection of at least m convex cones is non-empty.

Note that as a camera is further away from the 3D point \mathbf{X}_0 , it has weaker constraint, since the convex cone size at \mathbf{X}_0 becomes larger. This is a nice property since the further away from the camera, the larger uncertainty about the 3D position the camera has.

4 Experiments

We apply our quasiconvex optimization algorithm to multi-view triangulation and sequential structure from motion (SFM)(see [1]), and evaluate the performance using both synthetic and real data.

4.1 Synthetic data

The synthetic scene contains forty 3D points, distributed at different depth, that are imaged by a moving synthetic camera, as shown in Fig. 6(a). We use 10 consecutive views in the triangulation. Controlled zero-mean Gaussian noises and outliers are added to the 2D points. We apply our algorithm to minimize three cost functions F_{∞} , F_m , and F_m^w . Here F_m^w denotes F_m with weighted sum-of-infeasibilities used in Eq. (19). The reconstruction results from the algebraic approach (see Section 2.1.1) are included for comparison purpose.

Fig. 6 shows the average reconstruction errors, where (b) shows results when Gaussian noises are added to the 2D positions at increasing variances, and (c) shows the results with both Gaussian noises and 50% of outliers. The reconstruction error is normalized by

$$err = \frac{\|\mathbf{Z} - \mathbf{Z}_T\|_2}{\|\mathbf{Z}_T\|_2} \quad (22)$$

where \mathbf{Z}_T is the known ground truth of 3D position, and \mathbf{Z} is the triangulation result. As we can see, the algebraic approach has poor performance when there are noises or outliers, while our quasiconvex optimization successfully minimizes F_{∞} , F_m , and F_m^w . Without outliers, F_{∞} , F_m , and F_m^w have similar performance, with F_m and F_m^w better than F_{∞} when the noises become larger. When there are




Figure 7. Multi-view triangulation: corridor sequence. (a): The first image of this 11-frame sequence; (b): Reconstruction errors (normalized by Eq. (22)) with added zero-mean Gaussian noises; (c): Reconstruction errors with both Gaussian noises and outliers. For each feature track, the outliers are added to 1 to 3 views, depending on the number of views in which the corresponding 3D point is visible. (d): Reconstruction errors with increased strength of outliers (ranged from 5 to 40 pixels).

outliers, the performance of F_{∞} degrades quickly.

4.2 Real data with “ground truth”

We use the *corridor* sequence³ in which the camera is moving forward along the corridor. Fig. 7(a) shows the first frame of this 11-frame sequence. Along with the sequence, the 2D feature tracks, camera projection matrices, and 3D points are also provided. We use 2D feature tracks and camera matrices for triangulation, and compare the recovered 3D against the provided “ground truth”.

Controlled zero-mean Gaussian and/or outliers are added to the 2D feature coordinates. Fig. 7(b) and (c) show the reconstruction errors. The results are consistent with those from the synthetic data experiment. Again, our quasiconvex optimization successfully minimizes F_{∞} , F_m , and F_m^w .

We observed that F_{∞} is determined by outliers. Its performance depends on the “strength” of the outliers. Fig 7(d) shows the results where the strength of one outlier is increased. As we can see, the performance from F_{∞} degrades quickly when outlier strength is increased. F_m^w performs better than F_m when outlier strength is large. When the 2D feature tracking error is less than 25 pixels, F_m performs as well as F_m^w , indicating that in real scenarios F_m is usually good enough.

³<http://www.robots.ox.ac.uk/~vgg/data1.html>




Figure 8. Multi-view triangulation in sequential SFM. The camera is moved (largely forward motion) around inside the office. (a): The first, middle, and last frame of the 450-frame sequence (image size 360×240), with tracked points superimposed. (b): Top-down view of the reconstruction results of camera trajectory and 3D points. The yellow lines show the optical axis of the recovered cameras. The red circle indicates the 3D points corresponding to the chair.

4.3 Application: sequential structure from motion

Our target application is vision-aided small and micro aerial vehicle navigation, in which sequential SFM is applied to estimate both the camera motions and the 3D. We apply our multi-view triangulation using F_m minimization to the sequential SFM.

A 450-frame image sequence is taken by a mini camera that was moved around by hand in an office. Fig. 8(a) shows the first, middle, and last frames in this sequence. The camera is mostly moving forward, which is typical for a micro aerial vehicle. The forward motion makes the 3D estimation very challenging. Moreover, the images captured by the mini camera have low quality, resulting in noisy 2D feature tracking. We therefore seek to use as many frames as possible in triangulating a 3D point.

Fig. 8(b) shows the final reconstruction result (without global bundle adjustment). The red circle indicates the points from the chair visible both in the first and the last image. In the 3D view, the reconstruction of those points at the end of the sequence aligns very well with their reconstruction at the beginning of the sequence, indicating a good estimation of both the 3D and the camera motions.

5 Conclusion

We have presented a novel quasiconvex optimization framework to geometric reconstruction problems. Our algorithm is an efficient bisection search in the *one-dimensional* range domain, with each search step accomplished by a small-scale convex program that can be efficiently solved. We derived the algorithm based on sound mathematical grounds, and the algorithm is essentially free of parameter tuning. The final algorithm is simple, deterministic, and

has very intuitive geometric interpretation. We have demonstrated the effectiveness of our approach, using both synthetic and real data.

We identified the general quasiconvex formulation of the reprojection error functions, therefore our quasiconvex optimization framework can be potentially applied to many other estimation problems. We are investigating the applications of our approach to space carving [7], multi-baseline stereo reconstruction, and efficient bundle adjustment [12] in structure from motion.

Acknowledgements

This research was supported by USAF under grant No. F49620-03-1-0381 managed by Dr. Robert Sierakowski, Chief Scientist AFRL/MN, and Dr. Neal Glassman and Dr. Sharon Heise, Program Directors, AFOSR.

References

- [1] P. A. Beardsley, A. Zisserman, and D. W. Murray. Sequential updating of projective and affine structure from motion. *IJCV*, 23:235–259, June 1997.
- [2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [3] D. Diel, P. DeBitetto, and S. Teller. Epipolar constraints for vision-aided inertial navigation. In *IEEE Workshop on Motion and Video Computing*, 2005.
- [4] R. Haralick, C. Lee, K. Ottenberg, and M. Nolle. Review and analysis of solutions of the 3-point perspective pose estimation problem. *IJCV*, 13(3):331–356, 1994.
- [5] R. I. Hartley and F. Schaffalitzky. L_∞ minimization in geometric reconstruction problems. In *CVPR 2004*.
- [6] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [7] K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. *IJCV*, 38(3):199–218, 2000.
- [8] J. Oliensis. Exact two-image structure from motion. *PAMI*, 24(12):1618–1633, 2002.
- [9] L. Quan and Z. Lan. Linear $n \geq 4$ -point pose determination. In *ICCV 1998*.
- [10] R. Hartley and P. Sturm. Triangulation. *Computer Vision and Image Understanding*, 68(2):146–157, 1997.
- [11] P. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78(1):138–156, 2000.
- [12] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle adjustment – a modern synthesis. *Springer Lecture Notes in Computer Science*, 1883:298–372, 2000.
- [13] M. Uyttendaele, A. Criminisi, S. Kang, S. Winder, R. Hartley, and R. Szeliski. High-quality image-based interactive exploration of real-world environments. In *IEEE Computer Graphics and Applications*, 2004.
- [14] Z. Zhang. Determining the epipolar geometry and its uncertainty: A review. *IJCV*, 27(2):161–195, 1998.