

Deterministic Consensus Maximization with Biconvex Programming

Zhipeng Cai¹, Tat-Jun Chin¹, Huu Le², and David Suter³

¹ School of Computer Science, The University of Adelaide
{zhipeng.cai,tat-jun.chin}@adelaide.edu.au

² School of Electrical Engineering and Computer Science, Queensland University of Technology

huu.le@qut.edu.au

³ School of Computing and Security, Edith Cowan University
d.suter@ecu.edu.au

Abstract. Consensus maximization is one of the most widely used robust fitting paradigms in computer vision, and the development of algorithms for consensus maximization is an active research topic. In this paper, we propose an efficient *deterministic optimization* algorithm for consensus maximization. Given an initial solution, our method conducts a *deterministic search* that forcibly increases the consensus of the initial solution. We show how each iteration of the update can be formulated as an instance of biconvex programming, which we solve efficiently using a novel biconvex optimization algorithm. In contrast to our algorithm, previous consensus improvement techniques rely on random sampling or relaxations of the objective function, which reduce their ability to significantly improve the initial consensus. In fact, on challenging instances, the previous techniques may even return a worse off solution. Comprehensive experiments show that our algorithm can consistently and greatly improve the quality of the initial solution, without substantial cost.⁴

Keywords: Robust fitting · Consensus maximization · Biconvex programming

1 Introduction

Due to the existence of noise and outliers in real-life data, robust model fitting is necessary to enable many computer vision applications. Arguably the most prevalent robust technique is random sample consensus (RANSAC) [11], which aims to find the model that has the largest consensus set. The RANSAC algorithm approximately solves this optimization problem, by repetitively sampling minimal subsets of the data, in the hope of “hitting” an all-inlier minimal subset that gives rise to a model hypothesis with high consensus.

Many variants of RANSAC have been proposed [7]. Most variants attempt to conduct guided sampling using various heuristics, so as to speed up the retrieval

⁴ Matlab demo program is available in the supplementary material.

of all-inlier minimal subsets. Fundamentally, however, taking minimal subsets reduces the span of the data and produces biased model estimates [20, 27]. Thus, the best hypothesis found by RANSAC often has much lower consensus than the maximum achievable, especially on higher-dimensional problems. In reality, the RANSAC solution should only be taken as a rough initial estimate [9].

To “polish” a rough RANSAC solution, one can perform least squares (LS) on the consensus set of the RANSAC estimate (i.e. the Gold Standard Algorithm [12, Chap. 4]). Though justifiable from a maximum likelihood point of view, the efficacy of LS depends on having a sufficiently large consensus set to begin with.

A more useful approach is Locally Optimized RANSAC (LO-RANSAC) [9, 18], which attempts to enlarge the consensus set of an initial RANSAC estimate, by generating hypotheses from *larger-than-minimal subsets* of the consensus set.⁵ The rationale is that hypotheses fitted on a larger number of inliers typically lead to better estimates with even higher support. Ultimately, however, LO-RANSAC is also a randomized algorithm. Although it conducts a more focused sampling, the algorithm cannot guarantee improvements to the initial estimate. As we will demonstrate in Sec. 5.2, often on more challenging datasets, LO-RANSAC is unable to significantly improve upon the RANSAC result.

Due to its combinatorial nature, consensus set maximization is NP-hard [4]. While this has not deterred the development of globally optimal algorithms [21, 30, 19, 10, 6, 5, 25, 3], the fundamental intractability of the problem means that global algorithms are essentially variants of exhaustive search-and-prune procedures, whose runtime scales exponentially in the general case. While global algorithms have their place in computer vision, currently they are mostly confined to problems with low-dimensions and/or small number of measurements.

1.1 Deterministic algorithms—a new class of methods

Recently, efficient deterministic algorithms for consensus maximization are gaining attention [17, 22]. Different from random sampling, such algorithms begin with an initial solution (obtained using least squares or a random sampling method) and iteratively performs *deterministic updates* on the solution to improve its quality. While they do not strive for the global optimum, such algorithms are able to find excellent solutions due to the directed search.

To perform deterministic updating, the previous methods relax the objective function (Le et al. [17] use ℓ_1 penalization, and Purkait et al. [22] use a smooth surrogate function). Invariably this necessitates the setting of a smoothing parameter that controls the degree of relaxation, and the progressive tightening of the relaxation to ensure convergence to a good solution. As we will demonstrate in Sec. 5.4, incorrect settings of the smoothing parameter and/or its annealing rate may actually lead to a worse solution than the starting point.

⁵ This is typically invoked from within a main RANSAC routine.

1.2 Our contributions

We propose a novel deterministic optimization algorithm for consensus maximization. The overall structure of our method is a bisection search to increase the consensus of the current solution. The key to the effectiveness of our method is to formulate the feasibility test in each iteration as a *biconvex program*, which we solve efficiently via a biconvex optimization algorithm. Unlike [17, 22], our method neither relaxes the objective function, nor requires tuning of smoothing parameters. On both synthetic and real datasets, we demonstrate the superior performance of our method over previous consensus improvement techniques.

2 Problem definition

Given a set of N outlier contaminated measurements, consensus maximization aims to find the model $\mathbf{x} \in D$ that is consistent with the largest data subset

$$\underset{\mathbf{x} \in D}{\text{maximize}} \quad \mathcal{I}(\mathbf{x}), \quad (1)$$

where D is the domain of model parameters (more details later), and

$$\mathcal{I}(\mathbf{x}) = \sum_{i=1}^N \mathbb{I}(r_i(\mathbf{x}) \leq \epsilon) \quad (2)$$

counts the number of inliers (consensus) of \mathbf{x} . Function $r_i(\mathbf{x})$ gives the *residual* of the i -th measurement w.r.t. \mathbf{x} , ϵ is the inlier threshold and \mathbb{I} is the indicator function which returns 1 if the input statement is true and 0 otherwise.

Fig. 1 illustrates the objective function $\mathcal{I}(\mathbf{x})$. As can be appreciated from the inlier counting operations, $\mathcal{I}(\mathbf{x})$ is a step function with uninformative gradients.

2.1 The update problem

Let $\tilde{\mathbf{x}}$ be an initial solution to (1); we wish to improve $\tilde{\mathbf{x}}$ to yield a better solution. We define this task formally as

$$\text{find } \mathbf{x} \in D, \text{ such that } \mathcal{I}(\mathbf{x}) \geq \delta, \quad (3)$$




Fig. 1: Illustrating the update problem. Given the current solution $\tilde{\mathbf{x}}$ and a target consensus δ , where $\delta > \mathcal{I}(\tilde{\mathbf{x}})$, the update problem (3) aims to find another solution $\hat{\mathbf{x}}$ with $\mathcal{I}(\hat{\mathbf{x}}) \geq \delta$. Later in Sec. 4, problem (3) will be embedded in a broader algorithm that searches over δ to realize deterministic consensus maximization.

where $\delta > \mathcal{I}(\tilde{\mathbf{x}})$ is a target consensus value. See Fig. 1 for an illustration. For now, assume that δ is given; later in Sec. 4 we will embed (3) in a broader algorithm to search over δ .

Also, although (3) does not demand that the revised solution be “close” to $\tilde{\mathbf{x}}$, it is strategic to employ $\tilde{\mathbf{x}}$ as a starting point to perform the update. In Sec. 3, we will propose such an algorithm that is able to efficiently solve (3).

2.2 Residual functions and solvable models

Before embarking on a solution for (3), it is vital to first elaborate on the form of $r_i(\mathbf{x})$ and the type of models that can be fitted by the proposed algorithm. Following previous works [14, 15, 5], we focus on residual functions of the form

$$r_i(\mathbf{x}) = \frac{q_i(\mathbf{x})}{p_i(\mathbf{x})}, \quad (4)$$

where $q_i(\mathbf{x})$ is convex quadratic and $p_i(\mathbf{x})$ is linear. We also insist that $p_i(\mathbf{x})$ positive. We call $r_i(\mathbf{x})$ the *quasiconvex geometric residual* since it is quasiconvex [2, Sec. 3.4.1] in the domain

$$D = \{\mathbf{x} \in \mathbb{R}^d \mid p_i(\mathbf{x}) > 0, i = 1, \dots, N\}, \quad (5)$$

Note that D in the above form specifies a convex domain in \mathbb{R}^d .

Many model fitting problems in computer vision have residuals of the type (4). For example, in multiple view triangulation where we aim to estimate the 3D point $\mathbf{x} \in \mathbb{R}^3$ from multiple (possibly incorrect) 2D observations $\{\mathbf{u}_i\}_{i=1}^N$,

$$r_i(\mathbf{x}) = \frac{\|(\mathbf{P}_i^{(1:2)} - \mathbf{u}_i \mathbf{P}_i^{(3)})\tilde{\mathbf{x}}\|_2}{\mathbf{P}_i^{(3)}\tilde{\mathbf{x}}} \quad (6)$$

is the reprojection error in the i -th camera, where $\tilde{\mathbf{x}} = [\mathbf{x}^T \ 1]^T$,

$$\mathbf{P}_i = \begin{bmatrix} \mathbf{P}_i^{(1:2)} \\ \mathbf{P}_i^{(3)} \end{bmatrix} \in \mathbb{R}^{3 \times 4} \quad (7)$$

is the i -th camera matrix with $\mathbf{P}_i^{(1:2)}$ and $\mathbf{P}_i^{(3)}$ respectively being the first-two rows and third row of \mathbf{P} . Insisting that \mathbf{x} lies in the convex domain $D = \{\mathbf{x} \in \mathbb{R}^3 \mid \mathbf{P}_i^{(3)}\tilde{\mathbf{x}} > 0, \forall i\}$ ensures that the estimated \mathbf{x} lies in front of all the cameras.

Other model fitting problems with quasiconvex geometric residuals include homography fitting, camera resectioning, and the known rotation problem; see [14] for details and other examples. However, note that fundamental matrix estimation is not a quasiconvex problem [14]; in Sec. 5, we will show how the proposed technique can be adapted to robustly estimate the fundamental matrix.

3 Solving the update problem

As the decision version of (1), the update problem (3) is NP-complete [4] and thus can only be approximately solved. In this section, we propose an algorithm that works well in practice, i.e., able to significantly improve $\tilde{\mathbf{x}}$.

3.1 Reformulation as continuous optimization

With quasicovex geometric residuals (4), the inequality $r_i(\mathbf{x}) \leq \epsilon$ becomes

$$q_i(\mathbf{x}) - \epsilon p_i(\mathbf{x}) \leq 0. \quad (8)$$

Since $q_i(\mathbf{x})$ is convex and $p_i(\mathbf{x})$ is linear, the constraint (8) specifies a convex region in D . Defining

$$r'_i(\mathbf{x}) := q_i(\mathbf{x}) - \epsilon p_i(\mathbf{x}) \quad (9)$$

and introducing for each $r'_i(\mathbf{x})$ an indicator variable $y_i \in [0, 1]$ and a slack variable $s_i \geq 0$, we can write (3) using complementarity constraints [13] as

$$\text{find} \quad \mathbf{x} \in D \quad (10a)$$

$$\text{subject to} \quad \sum_i y_i \geq \delta, \quad (10b)$$

$$y_i \in [0, 1], \quad \forall i, \quad (10c)$$

$$y_i s_i = 0, \quad \forall i, \quad (10d)$$

$$s_i - r'_i(\mathbf{x}) \geq 0, \quad \forall i, \quad (10e)$$

$$s_i \geq 0, \quad \forall i. \quad (10f)$$

Intuitively, y_i reflects whether the i -th datum is an inlier w.r.t. \mathbf{x} . In the following, we establish the integrality of y_i and the equivalence between (10) and (3).

Lemma 1. *Problems (10) and (3) are equivalent.*

Proof. Observe that for any \mathbf{x} ,

a1: If $r'_i(\mathbf{x}) > 0$, the i -th datum is outlying to \mathbf{x} , and (10d) and (10e) will force $s_i \geq r'_i(\mathbf{x}) > 0$ and $y_i = 0$.

a2: If $r'_i(\mathbf{x}) \leq 0$, the i -th datum is inlying to \mathbf{x} , and (10f) and (10d) allow s_i and y_i to have only one of the following settings: **a2.1:** $s_i > 0$ and $y_i = 0$; or **a2.2:** $s_i = 0$ and y_i being indeterminate.

If \mathbf{x} is infeasible for (3), i.e., $\mathcal{I}(\mathbf{x}) < \delta$, condition **a1** ensures that (10b) is violated, hence \mathbf{x} is also infeasible for (10). Conversely, if \mathbf{x} is infeasible for (10), i.e., $\sum_i y_i < \delta$, then $\mathcal{I}(\mathbf{x}) < \delta$, hence \mathbf{x} is also infeasible for (3).

If \mathbf{x} is feasible for (3), we can always set $y_i = 1$ and $s_i = 0$ for all inliers to satisfy (10b), ensuring the feasibility of \mathbf{x} to (10). Conversely, if \mathbf{x} is feasible for (10), by **a1** there are at least δ inliers, thus \mathbf{x} is also feasible to (3). \square

From the computational standpoint, (10) is no easier to solve than (3). However, by constructing a cost function from the bilinear constraints (10d), we

arrive at the following continuous optimization problem

$$\begin{aligned} & \underset{\mathbf{x} \in D, \mathbf{s} \in \mathbb{R}^N, \mathbf{y} \in \mathbb{R}^N}{\text{minimize}} && \sum_i y_i s_i && (11a) \end{aligned}$$

$$\begin{aligned} & \text{subject to} && \sum_i y_i \geq \delta, && (11b) \end{aligned}$$

$$y_i \in [0, 1], \quad \forall i, \quad (11c)$$

$$s_i - r'_i(\mathbf{x}) \geq 0, \quad \forall i, \quad (11d)$$

$$s_i \geq 0, \quad \forall i, \quad (11e)$$

where $\mathbf{s} = [s_1, \dots, s_N]^T$ and $\mathbf{y} = [y_1, \dots, y_N]^T$. The following lemma establishes the equivalence between (11) and (3).

Lemma 2. *If the globally optimal value of (11) is zero, then there exists \mathbf{x} that satisfies the update problem (3).*

Proof. Due to (11c) and (11e), the objective value of (11) is lower bounded by zero. Let $(\mathbf{x}^*, \mathbf{s}^*, \mathbf{y}^*)$ be a global minimizer of (11). If $\sum_i y_i^* s_i^* = 0$, then \mathbf{x}^* satisfies all the constraints in (10), thus \mathbf{x}^* is feasible to (3). \square

3.2 Biconvex optimization algorithm

Although all the constraints in (11) are convex (including $\mathbf{x} \in D$), the objective function is not convex. Nonetheless, the primary value of formulation (11) is to enable the usage of convex solvers to approximately solve the update problem. Note also that (11) does not require any smoothing parameters.

To this end, observe that (11) is in fact an instance of *biconvex programming* [1]. If we fix \mathbf{x} and \mathbf{s} , (11) reduces to the linear program (LP)

$$\begin{aligned} & \underset{\mathbf{y} \in \mathbb{R}^N}{\text{minimize}} && \sum_i y_i s_i && (12a) \end{aligned}$$

$$\begin{aligned} & \text{subject to} && \sum_i y_i \geq \delta, && (12b) \end{aligned}$$

$$y_i \in [0, 1], \quad \forall i, \quad (12c)$$

which can be solved in close form.⁶ On the other hand, if we fix \mathbf{y} , (11) reduces to the second order cone program (SOCP)

$$\begin{aligned} & \underset{\mathbf{x} \in D, \mathbf{s} \in \mathbb{R}^N}{\text{minimize}} && \sum_i y_i s_i && (13a) \end{aligned}$$

$$\begin{aligned} & \text{subject to} && s_i - r'_i(\mathbf{x}) \geq 0, \quad \forall i, && (13b) \end{aligned}$$

$$s_i \geq 0, \quad \forall i. \quad (13c)$$

⁶ Set $y_i = 1$ if s_i is one of the δ -smallest slacks, and $y_i = 0$ otherwise.

Algorithm 1 Biconvex optimization (BCO) for the continuous problem (11).

Require: Initial solution $\tilde{\mathbf{x}}$, target consensus δ .

- 1: Initialize $\hat{\mathbf{x}} \leftarrow \tilde{\mathbf{x}}$, set $\hat{\mathbf{s}}$ using (14).
 - 2: **while** not converged **do**
 - 3: $\hat{\mathbf{y}} \leftarrow$ solve LP (12).
 - 4: $(\hat{\mathbf{x}}, \hat{\mathbf{s}}) \leftarrow$ solve SOCP (13).
 - 5: **end while**
 - 6: **return** $\hat{\mathbf{x}}$, $\hat{\mathbf{s}}$ and $\hat{\mathbf{y}}$.
-

Note that s_i does not have influence if the corresponding $y_i = 0$; these slack variables can be removed from the problem to speed up optimization.⁷

The proposed algorithm (called *Biconvex Optimization* or *BCO*) is simple: we initialize \mathbf{x} as the starting $\tilde{\mathbf{x}}$ from (3), and set the slacks as

$$s_i = \max\{0, r'_i(\tilde{\mathbf{x}})\}, \quad \forall i. \quad (14)$$

Then, we alternate between solving the LP and SOCP until convergence. Since (11) is lower-bounded by zero, and each invocation of the LP and SOCP are guaranteed to reduce the cost, BCO will always converge to a *local optimum* $(\hat{\mathbf{x}}, \hat{\mathbf{s}}, \hat{\mathbf{y}})$.

In respect to solving the update problem (3), if the local optimum $(\hat{\mathbf{x}}, \hat{\mathbf{s}}, \hat{\mathbf{y}})$ turns out to be the global optimum (i.e., $\sum_i \hat{y}_i \hat{s}_i = 0$), then $\hat{\mathbf{x}}$ is a solution to (3), i.e., $\mathcal{I}(\hat{\mathbf{x}}) \geq \delta$. Else, $\hat{\mathbf{x}}$ might still represent an improved solution over $\tilde{\mathbf{x}}$. Compared to randomized search, our method is by design more capable of improving $\tilde{\mathbf{x}}$. This is because optimizing (11) naturally reduces the residual of outliers that “should be” an inlier, i.e., with $y_i = 1$, which may still lead to a local refinement, i.e., $\mathcal{I}(\hat{\mathbf{x}}) > \delta_l = \mathcal{I}(\tilde{\mathbf{x}})$, regardless of whether problem (3) is feasible or not. In the next section, we will construct an effective deterministic consensus maximization technique based on Algorithm 1.

4 Main algorithm—deterministic consensus maximization

Given an initial solution $\mathbf{x}^{(0)}$ to (1), e.g., obtained using least squares or a random sampling heuristic, we wish to update $\mathbf{x}^{(0)}$ to a better solution. The main structure of our proposed algorithm is simple: we conduct bisection over the consensus value to search for a better solution; see Algorithm 2.

A lower and upper bound δ_l and δ_h for the consensus, which are initialized respectively to $\mathcal{I}(\mathbf{x}^{(0)})$ and N , are maintained and progressively tightened. Let $\tilde{\mathbf{x}}$ be the current best solution (initialized to $\mathbf{x}^{(0)}$); then, the midpoint $\delta = \lfloor 0.5(\delta_l + \delta_h) \rfloor$ is obtained and the update problem via the continuous biconvex formulation (11) is solved using Algorithm 1. If the solution $\hat{\mathbf{x}}$ for (11) has a higher quality than the incumbent, $\tilde{\mathbf{x}}$ is revised to become $\hat{\mathbf{x}}$ and δ_l is increased to $\mathcal{I}(\hat{\mathbf{x}})$. And if $\mathcal{I}(\hat{\mathbf{x}}) < \delta$, δ_h is decreased to δ . Algorithm 2 ends when $\delta_h = \delta_l + 1$.

⁷ Given the optimal $\hat{\mathbf{x}}$ for (13), the values of the slack variables that did not participate in the problem can be obtained as $s_i = \max\{0, r'_i(\hat{\mathbf{x}})\}$.

Algorithm 2 Bisection (non-global) for deterministic consensus maximization.

Require: Initial solution $\mathbf{x}^{(0)}$ for (1) obtained using least squares or random sampling.

```

1:  $\tilde{\mathbf{x}} \leftarrow \mathbf{x}^{(0)}$ ,  $\delta_h \leftarrow N$ ,  $\delta_l \leftarrow \mathcal{I}(\mathbf{x}^{(0)})$ .
2: while  $\delta_h > \delta_l + 1$  do
3:    $\delta \leftarrow \lfloor 0.5(\delta_l + \delta_h) \rfloor$ .
4:    $(\hat{\mathbf{x}}, \hat{\mathbf{s}}, \hat{\mathbf{y}}) \leftarrow \text{BCO}(\tilde{\mathbf{x}}, \delta)$  (see Algorithm 1).
5:   if  $\mathcal{I}(\hat{\mathbf{x}}) > \mathcal{I}(\tilde{\mathbf{x}})$  then
6:      $\tilde{\mathbf{x}} \leftarrow \hat{\mathbf{x}}$ ,  $\delta_l \leftarrow \mathcal{I}(\hat{\mathbf{x}})$ .
7:   end if
8:   if  $\mathcal{I}(\hat{\mathbf{x}}) < \delta$  then
9:      $\delta_h \leftarrow \delta$ .
10:  end if
11: end while
12: return  $\tilde{\mathbf{x}}, \delta_l$ .

```

Since the “feasibility test” in Algorithm 2 (Step 4) is solved via a non-convex subroutine, the bisection technique does not guarantee finding the global solution, i.e., the quality of the final solution may underestimate the maximum achievable quality. However, our technique is fundamentally advantageous compared to previous methods [9, 17, 22] since it is not subject to the vagaries of randomization or require tuning of hyperparameters. Empirical results in the next section will demonstrate the effectiveness of the proposed algorithm.

5 Results

We call the proposed algorithm *IBCO* (for *iterative biconvex optimization*). We compared IBCO against the following random sampling methods:

- RANSAC (RS) [11] (baseline): the confidence ρ was set to 0.99 for computing the termination threshold.
- PROSAC (PS) [8] and Guided MLESAC (GMS) [26] (RS variants with guided sampling): only tested for fundamental matrix and homography estimation since inlier priors like matching scores for correspondences were needed.
- LO-RANSAC (LRS) [9]: subset size in inner sampling was set to half of the current consensus size, and the max number of inner iterations was set to 10.
- Fixing LO-RANSAC (FLRS) [18]: subset size in inner sampling was set to $7 \times$ minimal subset size, and the max number of inner iterations was set to 50.
- USAC [23]: a modern technique that combines ideas from PS and LRS.⁸ USAC was evaluated only on fundamental matrix and homography estimation since the available code only implements these models.

Except USAC which was implemented in C++, the other sampling methods were based on MATLAB [16]. Also, least squares was executed on the final consensus set to refine the results of all the random sampling methods.

⁸ Code from <https://http://www.cs.unc.edu/~rraguram/usac/USAC-1.0.zip>.

Convex subproblem	LP	SOCP
Solvers used	Gurobi	Sedumi
Methods using the solver	EP, SS	IBCO

Table 1: Convex solvers used in deterministic methods.

In addition to the random sampling methods, we also compared IBCO against the following deterministic consensus maximization algorithms:

- Exact Penalty (EP) method [17]: The method⁹ was retuned for best performance on our data: we set the penalty parameter α to 1.5 for fundamental matrix estimation and 0.5 for all other problems. The annealing rate κ for the penalty parameter was set to 5 for linear regression and 2D homography estimation and 1.5 for triangulation and fundamental matrix estimation.
- Smooth Surrogate (SS) method [22]: Using our own implementation. The smoothing parameter γ was set to 0.01 as suggested in [22].

For the deterministic methods, Table 1 lists the convex solvers used for their respective subproblems. Further, results for these methods with both FLRS and random initialization ($\mathbf{x}^{(0)}$ was generated randomly) were provided, in order to show separately the performance with good (FLRS) and bad (random) initialization. We also tested least squares initialization, but under high outlier rates, its effectiveness was no better than random initialization. All experiments were executed on a laptop with Intel Core 2.60GHz i7 CPU and 16GB RAM.

5.1 Robust linear regression on synthetic data

Data of size $N = 1000$ for 8-dimensional linear regression (i.e., $\mathbf{x} \in \mathbb{R}^8$) were synthetically generated. In linear regression, the residual takes the form

$$r_i(\mathbf{x}) = \|\mathbf{a}_i^T \mathbf{x} - b_i\|_2, \quad (15)$$

which is a special case of (4) (set $p_i(\mathbf{x}) = 1$), and each datum is represented by $\{\mathbf{a}_i \in \mathbb{R}^8, b_i \in \mathbb{R}\}$. First, the independent measurements $\{\mathbf{a}_i\}_{i=1}^N$ and parameter vector \mathbf{x} were randomly sampled. The dependent measurements were computed as $b_i = \mathbf{a}_i^T \mathbf{x}$ and added with noise uniformly distributed between $[-0.3, 0.3]$. A subset of $\eta\%$ of the dependent measurements were then randomly selected and added with Gaussian noise of $\sigma = 1.5$ to create outliers. To guarantee the outlier rate, each outlier is regenerated until the noise is not within $[-0.3, 0.3]$. The inlier threshold ϵ for (1) was set to 0.3.

Fig. 2 shows the optimized consensus, runtime and model accuracy of the methods for $\eta \in \{0, 5, \dots, 70, 75\}$, averaged over 10 runs for each data instance. Note that the actual outlier rate was sometimes slightly lower than expected since the largest consensus set included some outliers with low noise value. For $\eta = 75$ the actual outlier rate was around 72% (see Fig. 2(a)). To prevent inaccurate analysis caused by this phenomenon, results for $\eta > 75$ were not provided.

⁹ Code from <https://cs.adelaide.edu.au/~hhu/>.

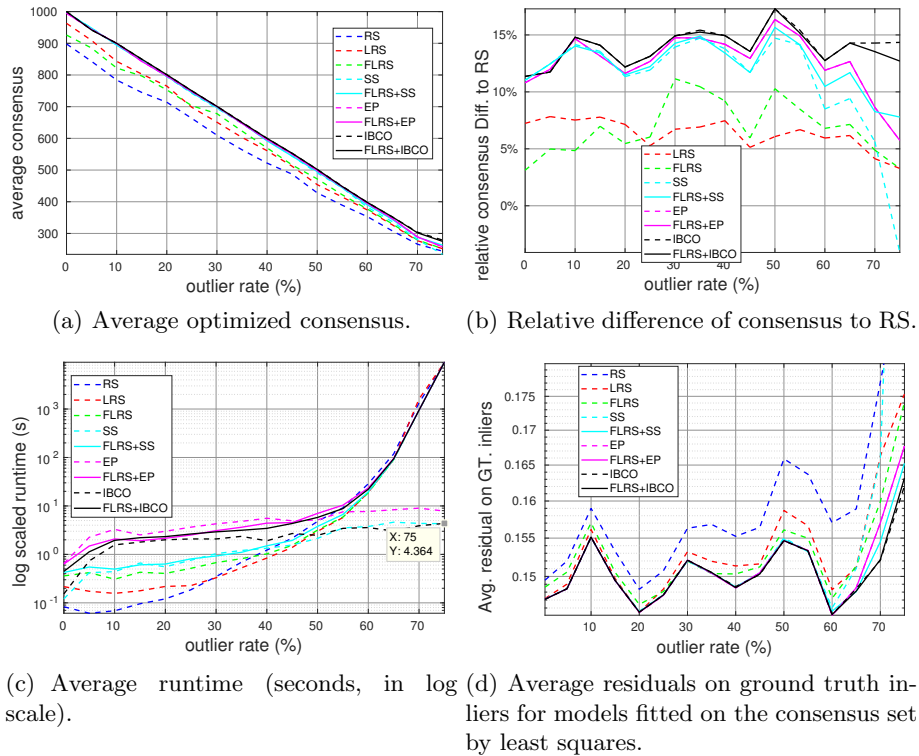
Fig. 2: Robust linear regression results with varied η (approx. outlier rate).

Fig. 2(b) demonstrates for each method the relative consensus difference to RS. It is evident that both IBCO variants outperformed other methods in general. Unlike other methods, whose improvement to RS was low at high outlier rates, both IBCO variants were consistently better than RS by more than 11%. Though IBCO was only marginally better than EP for outlier rates lower than 65%, Fig. 2(a) shows that for most of the data instances, both IBCO variants found consensus very close or exactly equal to the maximum achievable. The cost of IBCO was fairly practical (less than 5 seconds for all data instances, see the data tip in Fig. 2(c)). Also the runtime of the random sampling methods (RS, LRS, FLRS) increased exponentially with η . Hence, at high η , the major cost of FLRS+EP, FLRS+SS and FLRS+IBCO came from FLRS.

To demonstrate the significance of having higher consensus, we further performed least squares fitting on the consensus set of each method. Given a least squares fitted model \mathbf{x}_{LS} , define the average residual on ground truth inliers (the data assigned with less than 0.3 noise level) as:

$$e(\mathbf{x}_{LS}) = \frac{\sum_{i^* \in \mathcal{I}^*} r_{i^*}(\mathbf{x}_{LS})}{|\mathcal{I}^*|}, \quad (16)$$

where \mathcal{I}^* was the set of all ground truth inliers. Fig. 2(d) shows $e(\mathbf{x}_{LS})$ for all methods on all data instances. Generally, higher consensus led to a lower average residual, suggesting a more accurate model.

5.2 Homography estimation

Five image pairs from the NYC Library dataset [29] were used for 2D homography estimation. On each image pair, SIFT correspondences were produced by the VLFeat toolbox [28] and used as inputs. Fig. 3 depicts examples of inputs, as well as consensus sets from FLRS and FLRS+IBCO. The transfer error in one image [12, Sec. 4.2.2] was used as the distance measurement. The inlier threshold ϵ was set to 4 pixels. The 4-Point algorithm [12, Sec. 4.7.1] was used in all random sampling approaches for model fitting on minimal samples.

Fig. 4, shows the quantitative results, averaged over 50 runs. Though marginally costlier than SS and random approaches, both IBCO variants found considerably larger consensus sets than other methods for all data. Meanwhile, different from the linear regression case, EP no longer had similar result quality to IBCO. Also note that for challenging problems, e.g., *Ceiling1* and *Sign*, the two IBCO variants were the only methods that returned much higher consensus than RS.

5.3 Triangulation

Five feature tracks from the NotreDame dataset [24] were selected for triangulation, i.e., estimating the 3D coordinates. The input from each feature track contained a set of camera matrices and the corresponding 2D feature coordinates. The re-projection error was used as the distance measurement [15] and the inlier threshold ϵ was set to 1 pixel. The size of minimal samples was 2 (views) for all RANSAC variants. The results are demonstrated in Fig. 5. For triangulation, the quality of the initial solution largely affected the performance of EP, SS and IBCO. Initialized with FLRS, IBCO managed to find much larger consensus sets than all other methods.

5.4 Effectiveness of refinement

Though all deterministic methods were provided with reliable initial FLRS solutions, IBCO was the only one that effectively refined all FLRS results. EP and SS sometimes even converged to worse than initial solutions. To illustrate these effects, Fig. 6 shows the solution quality during the iterations of the three deterministic methods (initialized by FLRS) on *Ceiling1* for homography estimation and *Point 16* for triangulation. In contrast to EP and SS which progressively made the initial solution worse, IBCO steadily improved the initial solution.

It may be possible to rectify the behaviour of EP and SS by choosing more appropriate smoothing parameters and/or their annealing rates. However, the need for data-dependent tuning makes EP and SS less attractive than IBCO.

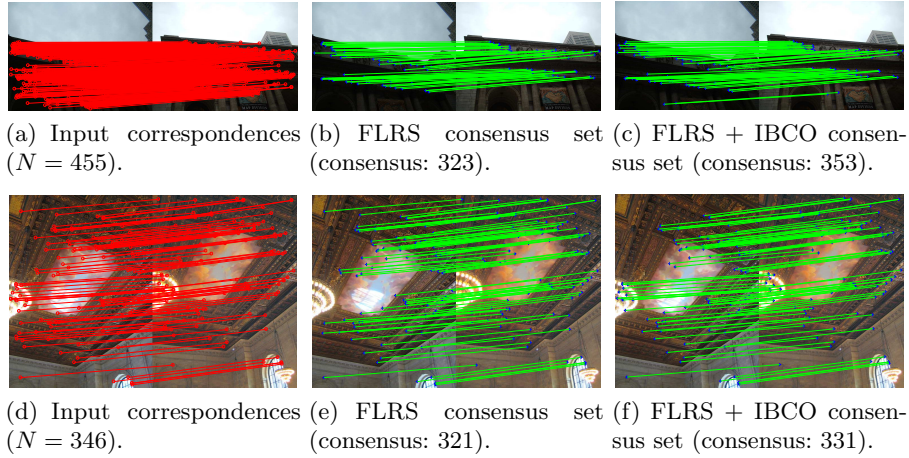


Fig. 3: Data and results of robust homography estimation for *Building1* (top) and *Ceiling1* (bottom). Consensus sets were downsampled for visual clearance.

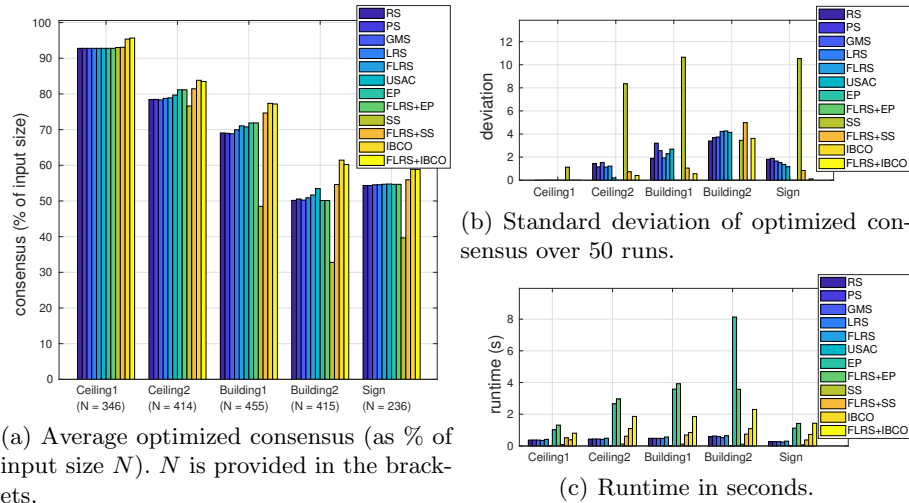


Fig. 4: Robust homography estimation results.

5.5 Fundamental matrix estimation

Image pairs from the two-view geometry corpus of CMP¹⁰ were used for fundamental matrix estimation. As in homography estimation, SIFT correspondences were used as the input data. Since the Sampson error [12, Sec. 11.4.3] and the reprojection error [12, Sec. 11.4.1] for fundamental matrix estimation are not linear or quasiconvex, the deterministic algorithms (EP, SS, IBCO) cannot be

¹⁰ <http://cmp.felk.cvut.cz/data/geometry2view/>

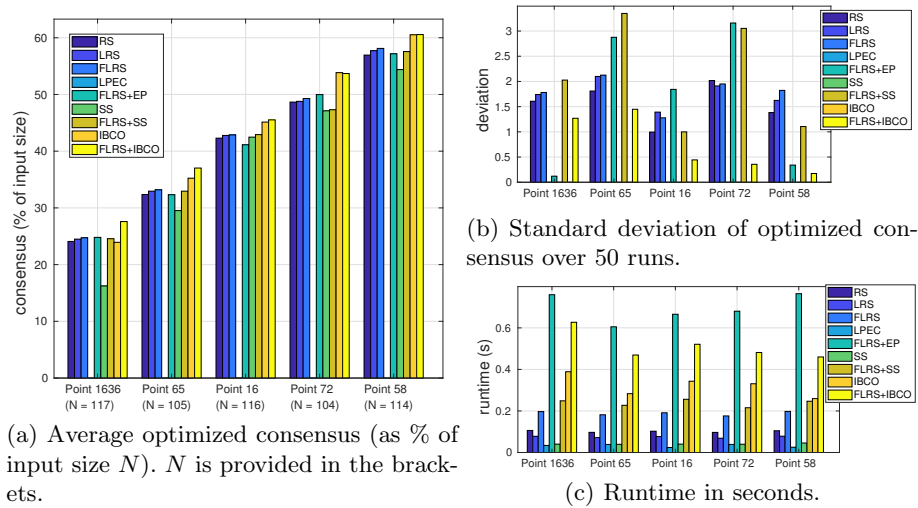


Fig. 5: Robust triangulation results.




Fig. 6: Consensus size in each iteration, given FLRS results as the initialization. Observe that EP and SS converged to worse off solutions.

directly applied. Thus, we linearize the epipolar constraint and use the algebraic error [12, Sec. 11.3] as the residual. The inlier threshold ϵ was set to 0.006 for all data.

Further, a valid fundamental matrix satisfies the rank-2 constraint [12, Sec. 11.1.1], which is non-convex. For EP, SS, IBCO, we impose the rank-2 constraint using SVD after each parameter vector updates (for IBCO, after each BCO run).

Fig. 7 depicts sample image pairs and generated SIFT correspondences, as well as consensus sets from FLRS and FLRS+IBCO. The seven-point method [12, Sec. 11.1.2] was used in USAC and the normalized 8-point algorithm [12, Sec. 11.2] was used in all other RANSAC variants.

As shown in Fig. 8(a), unlike EP and SS who failed to refine the initial FLRS results for all the tested data, IBCO was still effective even though the problem contains non-convex constraints.

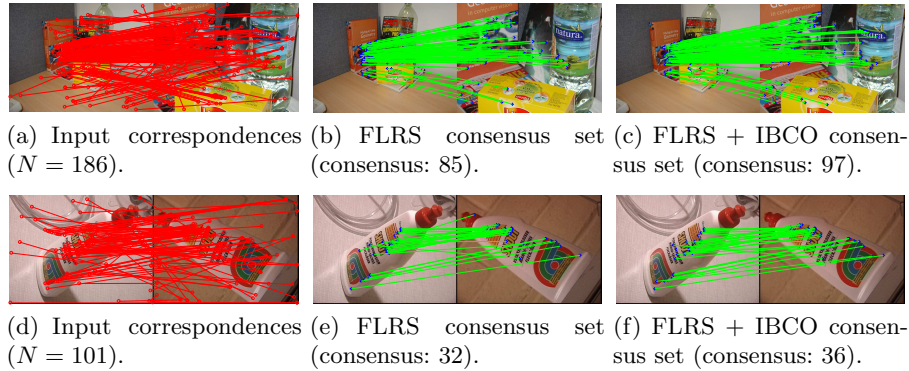


Fig. 7: Data and results of fundamental matrix estimation for *zoom* (top) and *shout* (bottom).

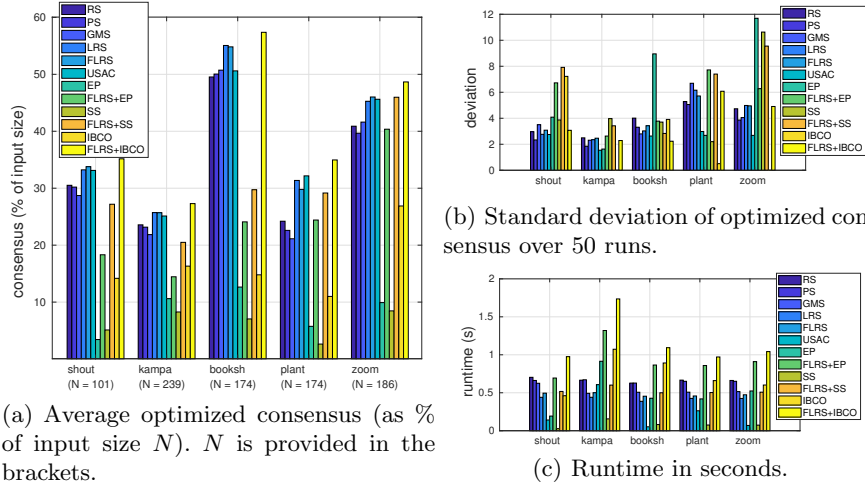


Fig. 8: Robust fundamental matrix estimation results.

6 Conclusions

We proposed a novel deterministic algorithm for consensus maximization with non-linear residuals. The basis of our method lies in reformulating the decision version of consensus maximization into an instance of biconvex programming, which enables the use of bisection for efficient guided search. Compared to other deterministic methods, our method does not relax the objective of consensus maximization problem and is free from the tuning of smoothing parameters, which makes it much more effective at refining the initial solution. Experiments show that our method is able to greatly improve upon initial results from widely used random sampling heuristics.

Acknowledgements This work was supported by the ARC grant DP160103490.

References

1. Biconvex optimization. https://en.wikipedia.org/wiki/Biconvex_optimization
2. Boyd, S., Vandenberghe, L.: Convex optimization. Cambridge university press (2004)
3. Campbell, D., Petersson, L., Kneip, L., Li, H.: Globally-optimal inlier set maximization for simultaneous camera pose and feature correspondence. arXiv preprint arXiv:1709.09384 (2017)
4. Chin, T.J., Cai, Z., Neumann, F.: Robust fitting in computer vision: Easy or hard? arXiv preprint arXiv:1802.06464 (2018)
5. Chin, T.J., Heng Kee, Y., Eriksson, A., Neumann, F.: Guaranteed outlier removal with mixed integer linear programs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5858–5866 (2016)
6. Chin, T.J., Purkait, P., Eriksson, A., Suter, D.: Efficient globally optimal consensus maximisation with tree search. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2413–2421 (2015)
7. Choi, S., Kim, T., Yu, W.: Performance evaluation of RANSAC family. In: British Machine Vision Conference (BMVC) (2009)
8. Chum, O., Matas, J.: Matching with prosac-progressive sample consensus. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 1, pp. 220–226. IEEE (2005)
9. Chum, O., Matas, J., Kittler, J.: Locally optimized ransac. In: DAGM. Springer (2003)
10. Enqvist, O., Ask, E., Kahl, F., Åström, K.: Tractable algorithms for robust model estimation. *International Journal of Computer Vision* **112**(1), 115–129 (2015)
11. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6), 381–395 (1981)
12. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge university press (2003)
13. Hu, J., Mitchell, J.E., Pang, J.S., Yu, B.: On linear programs with linear complementarity constraints. *Journal of Global Optimization* **53**(1), 29–51 (2012)
14. Kahl, F., Hartley, R.: Multiple-view geometry under the l_∞ -norm. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**(9), 1603–1617 (2008)
15. Ke, Q., Kanade, T.: Quasiconvex optimization for robust geometric reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(10) (2007)
16. Kovesi, P.D.: MATLAB and Octave functions for computer vision and image processing, available from: <<http://www.peterkovesi.com/matlabfns/>>
17. Le, H., Chin, T.J., Suter, D.: An exact penalty method for locally convergent maximum consensus. In: Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on. IEEE (2017)
18. Lebeda, K., Matas, J., Chum, O.: Fixing the locally optimized ransac–full experimental evaluation. In: British Machine Vision Conference. pp. 1–11. Citeseer (2012)
19. Li, H.: Consensus set maximization with guaranteed global optimality for robust geometry estimation. In: Computer Vision, 2009 IEEE 12th International Conference on. pp. 1074–1080. IEEE (2009)
20. Meer, P.: Robust techniques for computer vision. *Emerging topics in computer vision* pp. 107–190 (2004)

21. Olsson, C., Enqvist, O., Kahl, F.: A polynomial-time bound for matching and registration with outliers. In: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on. pp. 1–8. IEEE (2008)
22. Purkait, P., Zach, C., Eriksson, A.: Maximum consensus parameter estimation by reweighted ℓ_1 methods. In: International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition. pp. 312–327. Springer (2017)
23. Raguram, R., Chum, O., Pollefeys, M., Matas, J., Frahm, J.M.: Usac: a universal framework for random sample consensus. *IEEE transactions on pattern analysis and machine intelligence* **35**(8), 2022–2038 (2013)
24. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3d. In: *ACM transactions on graphics (TOG)*. vol. 25, pp. 835–846. ACM (2006)
25. Speciale, P., Paudel, D.P., Oswald, M.R., Kroeger, T., Gool, L.V., Pollefeys, M.: Consensus maximization with linear matrix inequality constraints. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5048–5056. IEEE (2017)
26. Tordoff, B.J., Murray, D.W.: Guided-mlesac: Faster image transform estimation by using matching priors. *IEEE transactions on pattern analysis and machine intelligence* **27**(10), 1523–1535 (2005)
27. Tran, Q.H., Chin, T.J., Chojnacki, W., Suter, D.: Sampling minimal subsets with large spans for robust estimation. *International journal of computer vision* **106**(1), 93–112 (2014)
28. Vedaldi, A., Fulkerson, B.: VLFeat: An open and portable library of computer vision algorithms. <http://www.vlfeat.org/> (2008)
29. Wilson, K., Snavely, N.: Robust global translations with ldsfm. In: Proceedings of the European Conference on Computer Vision (ECCV) (2014)
30. Zheng, Y., Sugimoto, S., Okutomi, M.: Deterministically maximizing feasible subsystem for robust model fitting with unit norm constraint. In: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on. pp. 1825–1832. IEEE (2011)