PL_1P - Point-line Minimal Problems under Partial Visibility in Three Views

Timothy Duff[1], Kathlén Kohn[2], Anton Leykin[3], Tomas Pajdla[4]

[1,3]Georgia Tech, Atlanta; [2]KTH Stockholm; [4]CIIRC, CTU in Prague

Abstract. We present a complete classification of minimal problems for generic arrangements of points and lines in space observed partially by three calibrated perspective cameras when each line is incident to at most one point. This is a large class of interesting minimal problems that allows missing observations in images due to occlusions and missed detections. There is an infinite number of such minimal problems; however, we show that they can be reduced to 140616 equivalence classes by removing superfluous features and relabeling the cameras. We also introduce camera-minimal problems, which are practical for designing minimal solvers, and show how to pick a simplest camera-minimal problem for each minimal problem. This simplification results in 74575 equivalence classes. Only 76 of these were known; the rest are new. To identify problems having potential for practical solving of image matching and 3D reconstruction, we present several natural subfamilies of camera-minimal problems as well as compute solution counts for all camera-minimal problems which have less than 300 solutions for generic data.

Keywords: minimal problems, calibrated cameras, 3D reconstruction

1 Introduction

Minimal problems [47,62,32,9,54,15,46,29,21,31,30,58,66,10,56,35,37,38,1,8,6,7] and [45], which we study, are 3D reconstruction problems recovering camera poses and world coordinates from given images such that random input instances have a finite positive number of solutions. They are important basic computational tasks in 3D reconstruction from images [60,61,59], image matching [55], visual odometry and localization [48,5,57,63]. Recently, a complete characterization of minimal problems for points, lines and their incidences in calibrated multi-view geometry appeared for the case of complete multi-view visibility [14]. In this paper, we extend the characterization to an important class of problems under *partial* multi-view visibility.

We provide a complete classification of minimal problems for generic arrangements of points and lines in space observed partially by three calibrated perspective cameras when each line is incident to at most one point. There is an infinite number of such minimal problems; however, we show that they can be *reduced* to 140616 equivalence classes of *reduced minimal* problems by removing superfluous features and relabeling the cameras. We compute a full description of each class in terms of the incidence structure in 3D and visibility of each 3D feature in images. All problems in every equivalence class have the same *algebraic degree*, i.e. the number of solutions over the complex numbers.

When using minimal solvers to find correct image matches by RANSAC [19,52], we often aim to recover camera parameters only. We name such reconstruction problems *camera-minimal* and reserve "minimal" for when we aim to recover 3D structure as well. Note that minimal problems are also camera-minimal but not vice versa. For instance, 50 out of the 66 problems given in [27] are non-minimal yet they all are camera-minimal. As an example, consider the problem from [27] with 3 PPP and 1 PPL correspondences. It is camera-minimal, i.e. there are 272 (in general complex) camera solutions, but it is not minimal since the line of the PPL correspondence cannot be recovered uniquely in 3D: there is a one-dimensional pencil of lines in 3D that project to the observed line in one of the images.

For each minimal problem, we delete additional superfluous features in images that can be removed without loosing camera-minimality to obtain a simplest camera-minimal problem. Thus, we introduce *terminal camera-minimal* problems. We show that, up to relabeling cameras, there are 74575 of these. They form the comprehensive list worth studying, as a solver for any camera-minimal problem can be derived from a solver for some problem on this list. Only 76 of the 74575 terminal camera-minimal problems were known — 66 problems listed in [27] plus 10 additional cases from [14] — the remaining 74499, to the best of our knowledge, are new! We find all terminal camera-minimal problems with less than 300 solutions for generic data and present other interesting cases that might be important for practical solving of image matching and 3D reconstruction.

Characterizing minimal problems under partial visibility, which allows for missing observations in images due to occlusions and missed detections, is very hard. Previous results in [14] treat the case of full visibility with no restrictions on the number of cameras and types of incidences, resulting in 30 minimal problems. By contrast, we construct a long list of interesting problems under partial visibility, even with our restrictions, i.e. having exactly three cameras and having each line incident to at most one point[1]. These restrictions make the task of enumerating problems tractable while making it still possible to account for very practical incidence cases where several existing feature detectors are applicable. For instance, SIFT [40] and LAF [42] provide quivers (points with one direction attached), which can be interpreted as lines through the points and used to compute relative camera poses [4].

2 Previous work

A large number of minimal problems appeared in the literature. See references above and [39,33,27,14] and references therein for work on general minimal prob-

[1] Under this restriction, in two cameras, the only reduced (camera-)minimal problem is the five-point problem; see Supplementary Material (SM) for an explanation.

lems. Here we review the most relevant work for minimal problems in three views related to point-line incidences and their classification.

Correspondences of non-incident points and lines in three uncalibrated views are considered in early works on the trifocal tensor [23]. Point-line incidences in the uncalibrated setup are introduced in [24] as n-quivers (points incident with n lines) and minimal problems for three 1-quivers in three affine views and three 3-quivers in three perspective views are derived. General uncalibrated multi-view constraints for points, lines and their incidences are presented in [41]. Non-incident points and lines in three uncalibrated images also appear in [50,26,51]. The cases of four points and three lines, two points and six lines, and nine lines are studied; [36] constructs a solver for nine lines. Works [17,18,16] look at lines incident to points which arise from tangent lines to curves and [4] presents a solver for that case. Results [25,1,2,3,65,64,14] introduced some of the techniques that are useful for classifying classes of minimal problems.

Work [27] classifies camera-minimal problems in 3 calibrated views that can be formulated with linear constraints on the trifocal tensor [22]. It presents 66 camera-minimal problems, all covered in our classification as terminal camera-minimal problems. Among them are 16 reduced minimal problems, out of which 2 are with full visibility and 14 with partial visibility. The remaining 50 problems are not minimal.

A complete characterization of minimal problems for points, lines and their incidences in calibrated multi-view geometry for the case of complete multi-view visibility is presented in [14]. It gives 30 minimal problems. Among them, 17 problems include exactly three cameras but only 12 of them (3002_1, 3002_2, 3010_0, 2005_3, 2005_4, 2005_5, 2013_2, 2013_3, 2021_1, 1024_4, 1032_2, 1040_0 in Tab. 1 of [14]) meet our restrictions on incidences. These 12 cases are all terminal camera-minimal as well as reduced minimal. Notice that the remaining 5 problems (3100_0, 2103_1, 2103_2, 2103_3, 2111_1 in Tab. 1 of [14]) are not considered in this paper because collinearity of more than two points cannot be modeled in the setting of this paper.

This paper can be seen as an extension of [27] and [14] to a much larger class of problems in three calibrated views under partial multi-view visibility.

3 Problem Specification

Our results apply to problems in which points, lines, and point-line incidences are partially observed. We model intersecting lines by requiring that each intersection point of two lines has to be one of the points in the point-line problem.

Definition 1. A *point-line problem* is a tuple $(p, l, \mathcal{I}, \mathcal{O})$ specifying that p points and l lines in space satisfy a given incidence relation

$$\mathcal{I} \subset \{1, \ldots, p\} \times \{1, \ldots, l\},$$

where $(i,j) \in \mathcal{I}$ means that the i-th point is on the j-th line, and are projected to $m = |\mathcal{O}|$ views with

$$\mathcal{O} = ((\mathcal{P}_1, \mathcal{L}_1), \ldots, (\mathcal{P}_m, \mathcal{L}_m))$$

describing which points and lines are observed by each camera — view v contains exactly the points in $\mathcal{P}_v \subset \{1, \ldots, p\}$ and the lines in $\mathcal{L}_v \subset \{1, \ldots, l\}$.

For $\mathcal{I}$ we assume *realizability* (the incidence relations are realizable by some point-line arrangement in $\mathbb{R}^3$) and *completeness* (every incidence which is automatically implied by the incidences in $\mathcal{I}$ must also be contained in $\mathcal{I}$).

For $\mathcal{O}$ we assume that if a camera observes two lines that meet according to $\mathcal{I}$ then it observes their point of intersection.

Note that, for instance, the realizability assumption implies that two distinct lines cannot have more than one point in common. Our assumption on $\mathcal{O}$ is natural — the set $\mathcal{I}$ of incidences describes all the knowledge about which lines intersect in space, as well as in the images. An *instance* of a point-line problem is specified by the following data:

(1) A point-line arrangement in space consisting of p points $X_1, \ldots, X_p$ and l lines $L_1, \ldots, L_l$ in $\mathbb{P}^3$ which are incident exactly as specified by $\mathcal{I}$. Hence, the point X_i is on the line L_j if and only if $(i,j) \in \mathcal{I}$. We write

$$\mathcal{X}_{p,l,\mathcal{I}} = \left\{ (X, L) \in \left(\mathbb{P}^3\right)^p \times \left(\mathbb{G}_{1,3}\right)^l \mid \forall (i,j) \in \mathcal{I} : X_i \in L_j \right\}$$

for the associated *variety of point-line arrangements*. Note that this variety also contains degenerate arrangements, where not all points and lines have to be pairwise distinct or where there are more incidences between points and lines than those specified by $\mathcal{I}$.

(2) A list of m calibrated cameras which are represented by matrices

$$P_1 = [R_1 \mid t_1], \ldots, P_m = [R_m \mid t_m]$$

with $R_1, \ldots, R_m \in \mathrm{SO}(3)$ and $t_1, \ldots, t_m \in \mathbb{R}^3$.

(3) The *joint image* consisting of the projections $\{x_{v,i} \mid i \in \mathcal{P}_v\} \subset \mathbb{P}^2$ of the points $X_1, \ldots, X_p$ and the projections $\{\ell_{v,j} \mid j \in \mathcal{L}_v\} \subset \mathbb{G}_{1,2}$ of the lines $L_1, \ldots, L_l$ by the cameras $P_1, \ldots, P_m$ to the views $v = 1, \ldots, m$. We denote by $\rho = \sum_{v=1}^m |\mathcal{P}_v|$ and $\lambda = \sum_{v=1}^m |\mathcal{L}_v|$ the total numbers of observed points and lines, and write

$$\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}} = \left\{ (x, \ell) \in \left(\mathbb{P}^2\right)^\rho \times \left(\mathbb{G}_{1,2}\right)^\lambda \;\middle|\; \begin{array}{l} \forall v = 1, \ldots, m \; \forall i \in \mathcal{P}_v \; \forall j \in \mathcal{L}_v : \\ (i,j) \in \mathcal{I} \Rightarrow x_{v,i} \in \ell_{v,j} \end{array} \right\}$$

for the *image variety* which consists of all m-tuples of 2D-arrangements of the points and lines specified by $\mathcal{O}$ which satisfy the incidences specified by $\mathcal{I}$. We note that an m-tuple in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ is not necessarily a joint image of a common point-line arrangement in $\mathbb{P}^3$.

Given a joint image, we want to recover an arrangement in space and cameras yielding the given joint image. We refer to a pair of such an arrangement and such a list of m cameras as a *solution* of the point-line problem for the given joint image.

To fix the arbitrary space coordinate system [22], we set $P_1 = [I\,|\,0]$ and the first coordinate of t_2 to 1. So our *camera configurations* are parameterized by

$$\mathcal{C}_m = \left\{ (P_1, \dots, P_m) \in \left(\mathbb{R}^{3\times 4}\right)^m \;\middle|\; \begin{array}{l} P_i = [R_i \mid t_i],\, R_i \in \mathrm{SO}(3),\, t_i \in \mathbb{R}^3, \\ R_1 = I,\, t_1 = 0,\, t_{2,1} = 1 \end{array} \right\}.$$

We will always assume that the camera positions in an instance of a point-line problem are sufficiently generic such that the points and lines in the views are in generic positions with respect to the specified incidences $\mathcal{I}$.

We say that a point-line problem is *minimal* if a generic image tuple in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ has a nonzero finite number of solutions. We may phrase this formally:

Definition 2. Let $\Phi_{p,l,\mathcal{I},\mathcal{O}} : \mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m \dashrightarrow \mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ denote the *joint camera map*, which sends a point-line arrangement in space and m cameras to the resulting joint image. We say that the point-line problem $(p, l, \mathcal{I}, \mathcal{O})$ is *minimal* if
- $\Phi_{p,l,\mathcal{I},\mathcal{O}}$ is a *dominant map*[2], i.e. a generic element (x, ℓ) in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ has a solution, so $\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell) \neq \emptyset$, and
- the preimage $\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell)$ of a generic element (x, ℓ) in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ is finite.

Remark 1. We require the joint camera map in Definition 2 to be dominant because we want solutions to minimal problems to be stable under perturbation of the image data that preserves the incidences $\mathcal{I}$. A classical example of a problem which is not stable under perturbation in images is the problem of four points in three calibrated views [49].

Over the complex numbers, the cardinality of the preimage $\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell)$ is the same for every *generic* joint image (x, ℓ) of a minimal point-line problem $(p, l, \mathcal{I}, \mathcal{O})$. We refer to this cardinality as the *degree* of the minimal problem.

In many applications, one is only interested in recovering the camera poses, and not the points and lines in 3D. Hence, we say that a point-line problem is *camera-minimal* if, given a generic image tuple in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$, it has a nonzero finite number of possible camera poses. Formally, this means:

Definition 3. Let $\gamma : \mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m \to \mathcal{C}_m$ denote the projection onto the second factor. We say that the point-line problem $(p, l, \mathcal{I}, \mathcal{O})$ is *camera-minimal* if
- its joint camera map $\Phi_{p,l,\mathcal{I},\mathcal{O}}$ is dominant, and
- $\gamma(\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell))$ is finite for a generic element (x, ℓ) in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$.

The cardinality over $\mathbb{C}$ of a generic $\gamma(\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell))$ is the *camera-degree* of $(p, l, \mathcal{I}, \mathcal{O})$.

Remark 2. Every minimal point-line problem is camera-minimal, but not necessarily the other way around. In the setting of complete visibility (i.e. where every camera observes all points and all lines), both notions coincide [14, Cor. 2].

[2] In birational geometry, dominant maps are analogs of surjective maps.

In [14], all minimal point-line problems with complete visibility are described, including their degrees. It is a natural question if one can extend the classification in [14] to all point-line problems with *partial visibility*. A first obstruction is that there are minimal point-line problems for arbitrarily many cameras[3], whereas the result in [14] shows that minimal point-line problems with complete visibility exist only for at most six views. Moreover, as we see in the following sections, deriving a classification for partial visibility seems more difficult and involves more elaborate tools. Hence, in this article, we only aim for classifying point-line problems *in three views*[4]. We also restrict our attention to point-line problems satisfying the following assumption:

Definition 4. We say that a point-line problem is a *PL_1P* if each line in 3D is incident to at most one point.

This assumption makes our analysis easier, since the point-line arrangement in space of a PL_1P is a collection of the following independent *local features*:
- *free line* (i.e. a line which is not incident to any point), and
- *point with k pins* where $k = 0, 1, 2, \dots$ (i.e. a point with k incident lines).

In the following, we shortly write *pin* for a line passing through a point. We stress that a pin refers only to the line itself, rather than the incident point. A first consequence of restricting our attention to PL_1Ps is the following fact, which fails for general point-line problems[5].

Lemma 1. *The degree and camera-degree of a minimal PL_1P coincide.*

Proof. Proofs of all lemmas, theorems and justification of results are in SM.

We will see that there are *infinitely many* (camera-)minimal PL_1Ps in three views. However, we can partition them into finitely many classes such that all PL_1Ps in the same class are closely related; in particular, they have the same (camera-)degree. For this classification, we pursue the following strategy:
Step 1: We introduce *reduced* PL_1Ps as the canonical representatives of the finitely many classes of minimal PL_1Ps we aim to find (see Section 4).
Step 2: Basic principles from algebraic geometry brought up in [14] imply

Lemma 2. *A point-line problem $(p, l, \mathcal{I}, \mathcal{O})$ is minimal if and only if*
- *it is* balanced, *i.e.* $\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m) = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}})$, *and*
- *its joint camera map $\Phi_{p,l,\mathcal{I},\mathcal{O}}$ is dominant.*

We identify a *finite* list of reduced balanced PL_1Ps in three views that contains all reduced minimal PL_1Ps (see Section 5).
Step 3: We explicitly describe the relation of reduced camera-minimal problems to reduced minimal ones, which implies that there are only finitely many reduced camera-minimal PL_1Ps in three views (see Section 6).

[3] See SM discussion of camera registration.
[4] See SM for a discussion on two views.
[5] See SM for an example.

Step 4: For each of the finitely many balanced PL_1Ps identified in Step 2, we check if its joint camera map is dominant. This provides us with a complete catalog of all reduced (camera-)minimal PL_1Ps in three views (see Section 7).

In addition to the classification, we compute the camera-degrees of all reduced camera-minimal PL_1Ps in three views whose camera-degree is less than 300 (see Section 8 for this and related results on natural subfamilies of PL_1Ps.)

4 Reduced PL_1Ps

From a given PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ we can obtain a new PL_1P by forgetting some points and lines, both in space and in the views. Formally, if $\mathcal{P}' \subset \{1, \ldots, p\}$ and $\mathcal{L}' \subset \{1, \ldots, l\}$ are the sets of points and lines which are *not* forgotten, the new PL_1P is $(p', l', \mathcal{I}', \mathcal{O}')$ with $p' = |\mathcal{P}'|$, $l' = |\mathcal{L}'|$, $\mathcal{I}' = \{(i, j) \in \mathcal{I} \mid i \in \mathcal{P}', j \in \mathcal{L}'\}$, and $\mathcal{O}' = ((\mathcal{P}'_1, \mathcal{L}'_1), \ldots, (\mathcal{P}'_m, \mathcal{L}'_m))$, where $\mathcal{P}'_v = \mathcal{P}_v \cap \mathcal{P}'$ and $\mathcal{L}'_v = \mathcal{L}_v \cap \mathcal{L}'$. This induces natural projections Π and π between the domains and codomains of the joint camera maps which forget the points and lines *not* in $\mathcal{P}'$ and $\mathcal{L}'$.

$$\begin{array}{ccc} \mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m & \xrightarrow{\Phi = \Phi_{p,l,\mathcal{I},\mathcal{O}}} & \mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}} \\ \Pi \downarrow & & \downarrow \pi \\ \mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_m & \xrightarrow[\Phi' = \Phi_{p',l',\mathcal{I}',\mathcal{O}'}]{} & \mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'} \end{array}$$

In the following, we shortly write $\Phi = \Phi_{p,l,\mathcal{I},\mathcal{O}}$ and $\Phi' = \Phi_{p',l',\mathcal{I}',\mathcal{O}'}$.

Definition 5. We say that $(p, l, \mathcal{I}, \mathcal{O})$ is *reducible* to $(p', l', \mathcal{I}', \mathcal{O}')$ if
- for each forgotten point, at most one of its pins is kept, and
- a generic solution $S' = ((X', L'), P) \in \mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_m$ of $(p', l', \mathcal{I}', \mathcal{O}')$ can be lifted to a solution of $(p, l, \mathcal{I}, \mathcal{O})$ for generic input images in $\pi^{-1}(\Phi'(S'))$.
  In other words, for a generic $S' = ((X', L'), P) \in \mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_m$ and a generic $(x, \ell) \in \pi^{-1}(\Phi'(S'))$, there is a point-line arrangement $(X, L) \in \mathcal{X}_{p,l,\mathcal{I}}$ such that $\Phi((X, L), P) = (x, \ell)$ and $\Pi((X, L), P) = S'$.

Theorem 1. *If a PL_1P is minimal and reducible to another PL_1P, then both are minimal and have the same degree.*

We can partition *all* (infinitely many) minimal PL_1Ps in three views into finitely many classes using this reduction process. Each class is represented by a unique PL_1P that is *reduced*, i.e. not reducible to another PL_1P.

Theorem 2. *A minimal PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ in three views is reducible to a unique reduced PL_1P $(p', l', \mathcal{I}', \mathcal{O}')$. The corresponding projection Π forgets:*
- *every pin that is observed in exactly two views such that both views also observe the point of the pin (it does not matter if the third view observes the point or not, but it must not see the line), e.g.* 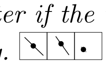 *is reduced to* 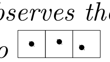

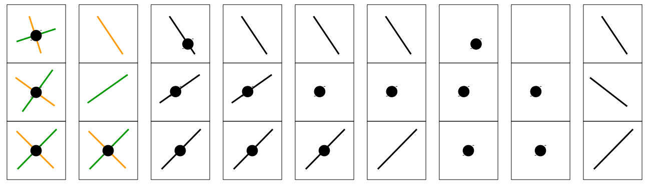

#	$c_{2,0}$	$c^a_{2,1}$ $c^b_{2,1}$ $c^c_{2,1}$	$c_{1,0}$	$c^a_{1,1}$ $c^b_{1,1}$ $c^c_{1,1}$	$c^a_{1,2}$ $\vdots$ $c^f_{1,2}$	$c^a_{1,3}$ $c^b_{1,3}$ $c^c_{1,3}$	$c_{0,0}$	$c^a_{0,1}$ $c^b_{0,1}$ $c^c_{0,1}$	f
3D	7	7	5	5	5	5	3	3	4
2D	12	8	9	8	7	6	6	4	6

Table 1. How points with two / one / zero pins and free lines can be observed in the three views of a reduced minimal PL_1P (up to permuting the views). The rows "3D" and "2D" show the degrees of freedom of each local feature in 3-space and in the three views. The row "#" fixes notation for a signature introduced in Section 5.

- *every free line that is observed in exactly two views to reduce* 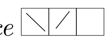 *to* 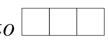
- *every point that has exactly one pin and is viewed like* 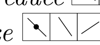 *to get* 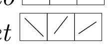
- *every point together with its single pin if it is viewed like* 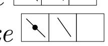 *to get* 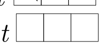

In addition, applying inverses of these reductions to a minimal PL_1P in three views results in a minimal PL_1P.

Hence, it is enough to classify all reduced minimal PL_1Ps. We will see that there are finitely many reduced minimal PL_1Ps in three views. To count them, we need to understand how they look.

Theorem 3. *A reduced minimal PL_1P in three views has at most one point with three or more pins. If such a point exists,*
- *it has at most seven pins,*
- *and the point and all its pins are observed in all three views.*

All other local features are viewed as in Table 1.

5 Balanced PL1Ps

A reduced minimal PL_1P in three views is uniquely determined by a *signature*, a vector consisting of 27 numbers $(c_7, \ldots, c_3, c_{2,0}, c^a_{2,1}, \ldots, f)$, that specifies how often each local feature occurs in space and how often it is observed in a certain way by the cameras. By Theorem 3, the local features in such a PL_1P are free lines or points with at most seven pins. We denote by f the number of free lines and write $c_3, c_4, \ldots, c_7$ for the numbers of points with three, four, …, seven pins. By Theorem 3, these local features are completely observed by the cameras. The row "#" in Table 1 shows our notation for the numbers of points with zero, one or two pins that are viewed in a certain way. For instance, $c_{2,0}$ counts how many points with two pins are completely observed by the cameras. Moreover,

$c^a_{2,1}, c^b_{2,1}, c^c_{2,1}$ are the numbers of points with two pins that are partially observed like 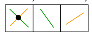 or 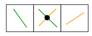 or 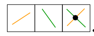. Here the upper index a, b, c distinguishes the three different permutations of this local feature in the three views (note: as the two pins can be relabeled, there are only three and not six permutations). Similarly, upper indices distinguish different permutations of partially viewed points with at most one pin; see Table 1. We also note that assigning arbitrary 27 non-negative integers to $c_7, \ldots, f$ describes a unique $\mathrm{PL}_1\mathrm{P}$ in three views, which is reduced by construction (see Thm. 2 and 3) but not necessarily minimal.

Due to Lemma 2, every minimal $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ is balanced, i.e. it satisfies $\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m) = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}})$. To compute the dimension of $\mathcal{X}_{p,l,\mathcal{I}}$, we need to know the degrees of freedom of each local feature in 3-space. For free lines and points with at most two pins, this is given in the row "3D" in Table 1. More generally, a point in space with k pins has $3 + 2k$ degrees of freedom. Hence, a reduced minimal $\mathrm{PL}_1\mathrm{P}$ in three views satisfies

$$\begin{aligned}\dim(\mathcal{X}_{p,l,\mathcal{I}}) = {} & 17c_7 + 15c_6 + 13c_5 + 11c_4 + 9c_3 + 7(c_{2,0} + c^a_{2,1} + c^b_{2,1} + c^c_{2,1}) \\ & + 5(c_{1,0} + c^a_{1,1} + c^b_{1,1} + c^c_{1,1} + c^a_{1,2} + \ldots + c^f_{1,2} + c^a_{1,3} + c^b_{1,3} + c^c_{1,3}) \\ & + 3(c_{0,0} + c^a_{0,1} + c^b_{0,1} + c^c_{0,1}) + 4f.\end{aligned}$$

Similarly, the degrees of freedom of each local feature in the three views are shown in row "2D" in Table 1. For instance, if a point with two pins is viewed like 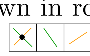, then it has eight degrees of freedom in the three views: 2+1+1 in the first view, 2 in the second view, and 2 in the third view. Since a point with k pins for $k = 3, \ldots, 7$ is completely observed by the cameras, it has $3(2 + k)$ degrees of freedom in the three views. Therefore, we have

$$\begin{aligned}\dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}) = {} & 27c_7 + 24c_6 + 21c_5 + 18c_4 + 15c_3 + 12c_{2,0} + 8(c^a_{2,1} + c^b_{2,1} + c^c_{2,1}) \\ & + 9c_{1,0} + 8(c^a_{1,1} + c^b_{1,1} + c^c_{1,1}) + 7(c^a_{1,2} + \ldots + c^f_{1,2}) + 6(c^a_{1,3} + c^b_{1,3} + c^c_{1,3}) \\ & + 6c_{0,0} + 4(c^a_{0,1} + c^b_{0,1} + c^c_{0,1}) + 6f.\end{aligned} \quad (1)$$

As $\dim(\mathcal{C}_3) = 11$, the balanced equality $\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m) = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}})$ for a reduced minimal $\mathrm{PL}_1\mathrm{P}$ in three views is $11 = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}) - \dim(\mathcal{X}_{p,l,\mathcal{I}})$, i.e.

$$\begin{aligned}11 = {} & 10c_7 + 9c_6 + 8c_5 + 7c_4 + 6c_3 + 5c_{2,0} + (c^a_{2,1} + c^b_{2,1} + c^c_{2,1}) \\ & + 4c_{1,0} + 3(c^a_{1,1} + c^b_{1,1} + c^c_{1,1}) + 2(c^a_{1,2} + \ldots + c^f_{1,2}) + (c^a_{1,3} + c^b_{1,3} + c^c_{1,3}) \\ & + 3c_{0,0} + (c^a_{0,1} + c^b_{0,1} + c^c_{0,1}) + 2f.\end{aligned} \quad (2)$$

The linear equation (2) has 845161 non-negative integer solutions[6]. Each of these is a signature that represents a $\mathrm{PL}_1\mathrm{P}$ in three views which is reduced and balanced. Thus, it remains to check which of the 845161 signatures represent *minimal* $\mathrm{PL}_1\mathrm{Ps}$.

Some of the 845161 solutions of (2) yield *label-equivalent* $\mathrm{PL}_1\mathrm{Ps}$, i.e. $\mathrm{PL}_1\mathrm{Ps}$ which are the same up to relabeling the three views. It turns out that there

[6] See SM for details on how to solve it.

143494 such label-equivalence classes of PL_1Ps given by solutions to (2)[7]. So all in all, we have to check 143494 PL_1Ps for minimality, namely one representative for each label-equivalence class.

6 Camera-Minimal PL1Ps

As in the case of minimal problems, we can understand all camera-minimal PL_1Ps from the reduced ones (see also Theorem 7).

Theorem 4. *If a PL_1P is reducible to another PL_1P, then either none of them is camera-minimal or both are camera-minimal. In the latter case, their camera-degrees are equal.*

In order to understand how reduced camera-minimal PL_1Ps look, in comparison to reduced minimal PL_1Ps as described in Theorem 3, we define a pin to be *dangling* if it is viewed by exactly one camera. Dangling pins are not determined uniquely by the camera observations, and hence they appear in PL_1Ps that are camera-minimal but not minimal.

Theorem 5. *The local features of a reduced camera-minimal PL_1P in three views are viewed as described in Theorem 3 plus as in the following three additional cases:*

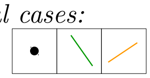

point with two pins, both are dangling

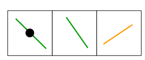

point with two pins, one of which is dangling

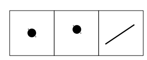

point with one pin, which is dangling

Remark 3. For a dangling pin L of a reduced camera-minimal PL_1P, the point X incident to the pin L is uniquely reconstructible. Since L is viewed by exactly one camera, it belongs to the planar pencil of lines which are incident to X and have the same image as L. Thus we see that L is not uniquely reconstructible from its image.

The next theorem relates minimal and camera-minimal PL_1Ps. By adding more constraints to images, we make configurations in space uniquely reconstructible.

Theorem 6. *The following replacements in images lift a reduced camera-minimal PL_1P in three views to a reduced minimal PL_1P (cf. Thm. 5 and Table 1):*

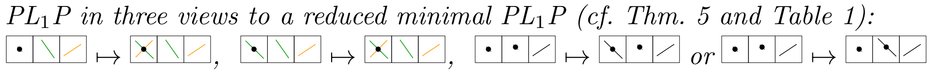

Moreover, the camera-degrees of both PL_1Ps are the same.

This has two important implications for classifying (camera-)minimal PL_1Ps. First, reversing the replacements in Theorem 6 transforms each reduced camera-minimal PL_1P in three views into a *terminal* PL_1P of the same camera-degree.

[7] See SM for details on how to compute this.

(a)

camera-degree	64	80	144	160	216	224	240	256	264	272	288
# problems	13	9	3	547	7	2	159	2	2	11	4

(b)

camera-degree	80	160	216	240	256	264	272	288	304	312	320	352	360	368	376
# problems	9	173	4	80	2	2	2	1	5	2	213	3	9	3	1
camera-degree	384	392	400	408	416	424	432	448	456	464	472	480	488	496	
# problems	2	9	14	2	6	10	2	7	11	4	1	96	12	9	

Table 2. Distribution of camera-degrees of terminal camera-minimal PL_1Ps in three calibrated views with: (a) camera-degree less than 300, (b) at most one pin per point and camera-degree less than 500.

Definition 6. We say that a camera-minimal PL_1P in three views is *terminal* if it is reduced and does not view local features like 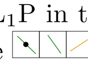 or 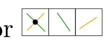 or 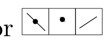.

Hence, to classify *all* camera-minimal PL_1Ps in three views, it is enough to find the terminal ones. Secondly, Theorem 6 implies for minimal PL_1Ps the following.

Corollary 1. *Consider a minimal PL_1P in three views. After replacing a single occurrence of* 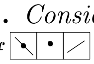 *with* 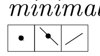 *(or the other way around), the resulting PL_1P is minimal and has the same degree.*

At the end of Section 5, we defined two PL_1Ps to be label-equivalent if they are the same up to relabeling the views. We note that the swap described in Corollary 1 does *not* preserve the label-equivalence class of a PL_1P. Instead, we say that two PL_1Ps in three views are *swap&label-equivalent* if one can be transformed into the other by relabeling the views and applying (any number of times) the swap in Corollary 1. We conclude that either all PL_1Ps in the same swap&label-equivalence class are minimal and have the same degree, or none of them is minimal. Moreover, the lift in Theorem 6 yields the following.

Corollary 2. *The swap&label-equivalence classes of reduced minimal PL_1Ps in three views are in a camera-degree preserving one-to-one correspondence with the label-equivalence classes of terminal camera-minimal PL_1Ps in three views.*

Hence, we do not have to check minimality for all 143494 label-equivalence classes of PL_1Ps given by solutions to (2), that we found at the end of Section 5. Instead it is enough to consider the swap&label-equivalence classes of the solutions to (2). It turns out that there are **76446** such classes[8]. So to find *all* (camera-)minimal PL_1Ps in three views, we only have to check 76446 PL_1Ps for minimality, namely one representative for each of the swap&label-equivalence classes.
Finally, we present the analog to Theorem 2 and describe how all camera-minimal PL_1Ps are obtained from the reduced camera-minimal ones.

[8] See SM for details on how to compute this.

Theorem 7. *A camera-minimal PL_1P in three views is reducible to a unique reduced PL_1P. The corresponding projection forgets:*
- *everything that is forgotten in Theorem 2*
- *every line (free or pin) that is not observed in any view*
- *every free line that is observed in exactly one view to reduce* 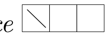 *to* 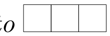
- *every pin that is observed in exactly one view such that the view also observes the point of the pin (it does not matter if the other two views observe the point or not, but they must not see the line), e.g.* 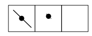 *is reduced to* 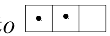
- *every point without pins that is observed in at most one view, e.g.* 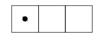 *is reduced to* 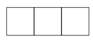
- *every point that has exactly one pin if the point is not observed in any view, e.g. a pin viewed like* 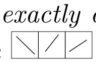 *becomes a free line viewed like* 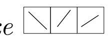
- *every point together with its single pin if it is viewed like* 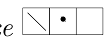 *to get* 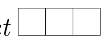
- *every point together with all its pins if the point has at least two pins and the point is not observed in any view, e.g.* 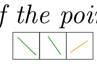 *is reduced to* 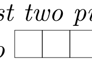

7 Checking minimality

To show that a balanced point-line problem $(p, l, \mathcal{I}, \mathcal{O})$ is minimal, it is equivalent to show that the Jacobian of the joint camera map $\Phi_{p,l,\mathcal{I},\mathcal{O}}$ at some point $(X, P) \in \mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m$ has full rank, i.e. rank given by the formula in equation (1). This follows from Lemma 2, as explained in [14]. On the implementation level, this minimality criterion requires writing down local coordinates for the various projective spaces and Grassmannians. To take advantage of fast exact arithmetic and linear algebra, we ran each test with random inputs (X, P) over a finite field $\mathbb{F}_q$ for some large prime q. We observe that *false positives*[9] for these tests are impossible. To guard against *false negatives*, we re-run the test on remaining non-minimal candidates for different choices of q. Moreover, as a byproduct of our degree computations, we obtain yet another test of minimality, following the same procedure as [14, Algorithm 1].
The computation described above detects non-minimality for 2878 of the 143494 label-equivalence classes of $\mathrm{PL}_1\mathrm{Ps}$ given by solutions to (2). Among the 76446 swap&label-equivalence classes, 1871 are not minimal.

Result 8 *In three calibrated views, up to relabeling cameras, there are*
- **140616** $= 143494 - 2878$ *reduced minimal PL_1Ps and*
- **74575** $= 76446 - 1871$ *terminal camera-minimal PL_1Ps.*

8 Computing degrees

From the perspective of solving minimal problems, it is highly desirable to compute all degrees of our minimal $\mathrm{PL}_1\mathrm{Ps}$. In particular, we wish to identify prob-

[9] Since we are testing minimality, being minimal is the positive outcome. See SM for detailed explanation why false positives cannot occur.

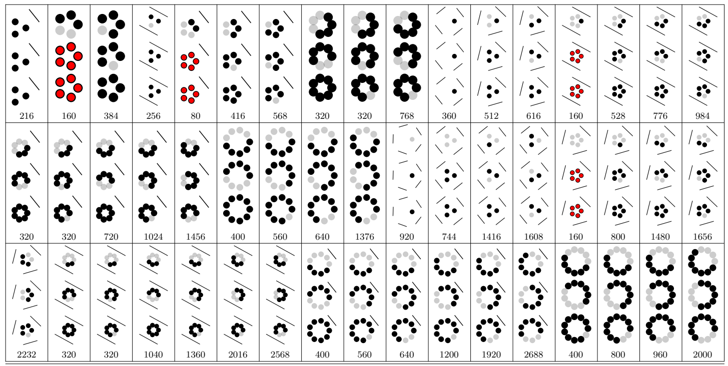

Table 3. Reduced minimal PL_0Ps and their degrees. Points not visible in a given view are indicated in grey. Five-point subproblems are indicated in red.

lems with small degrees that may be of practical interest. Since some problems in this list are known from prior work [27] to have large degrees (> 1000), our main technique is a monodromy approach based on numerical homotopy continuation. Our implementation in Macaulay2 [20,13] is similar to that used in prior work [14]. The next result shows that there are many interesting problems with small degrees that are solvable with existing solving technology [34,4].

Result 9 *There are* 759 *(up to relabeling cameras) terminal camera-minimal* PL_1Ps *in three calibrated views with camera-degree less than* 300*. Their camera-degree distribution is shown in Table 2(a).*

It is also interesting to look at problems with simple incidence structure, since they are easier to detect in images. This motivates the following.

Definition 7. We say that a point-line problem is a $PL_\kappa P$ if each line in 3D is incident to at most κ points.

Definition 7 generalizes Definition 4 of PL_1Ps to get a hierarchy of subfamilies of point-line problems: $PL_0P \subset PL_1P \subset PL_2P \subset \dots$
For instance, PL_0Ps consist of free points and free lines only[10]. The family of PL_1Ps contains many problems involving at most one pin per point—see Result 10 and Table 2(b). Such features are readily provided by the SIFT [40] detector. PL_1Ps can also include features with two pins per point, which are readily provided by the LAF [43] detector, which can also be used to get PL_2Ps if all 3 LAF points are used. More complex incidences (e.g. PL_3Ps) can be obtained from line t-junction detectors [67].

Result 10 *There are* 9533 *(up to relabeling cameras) terminal camera-minimal* PL_1Ps *in three calibrated views which have at most one pin per point.* ***694*** *of them have camera-degree less than* ***500****. Their camera-degree distribution is shown in Table 2(b).*

[10] We note that reduced minimal PL_0Ps are terminal.

Result 11 *There are* 51 *(up to relabeling cameras) reduced minimal PL_0Ps in three calibrated views. They are depicted together with their degrees in Table 3.*

Note that there are four problems in Table 3 that are *extensions* of the classical minimal problem of five points in two views. This implies that the relative pose of the two cameras can be determined from the five point correspondences (highlighted in red). As to the remaining camera, each of these four problems can be interpreted as a *camera registration* problem (the first one is known as P3P [28]): given a set of points and lines in the world and their images for a camera, find that camera pose. Note that the solution counts indicate that there are 8, 4, 8, and 8 solutions to the corresponding four camera registration problems. Similar degrees were previously reported for camera registration from 3D points and lines for perspective and generalized cameras [12,11,53,44,45].

Result 12 *We determined all PL_1Ps in three calibrated views that are extensions of the five-points minimal problem. Of them, up to relabeling cameras,*
- 6300 *are reduced minimal,*
- 61 *of the* 6300 *correspond to camera registration problems (see SM)*
- 3648 *are terminal camera-minimal.*

9 Conclusion

We have explicitly classified all reduced minimal and camera-minimal problems in three calibrated views for configurations of points and lines when lines contain at most one point.

The number of (camera-)minimal problems in our classification is large. Apart from constructing a database of all these problems, we identify interesting subfamilies where the number of the problems is relatively small (see Tables Table 3, Table 2 in this article and Table 4 in SM.)

Another part of our computational effort focused on determining algebraic degrees of the (camera-)minimal problems. The degree of a problem provides a measure of complexity of a solver one may want to construct. The smaller the degree, the more plausible it is that a problem could be used in practice: Table 2 shows the degree distributions for problems of degree less than 300.

Our code is available at `https://github.com/timduff35/PL1P`.

Acknowledgements

We thank ICERM (NSF DMS-1439786 and the Simons Foundation grant 507536). We acknowledge: T. Duff and A. Leykin - NSF DMS-1719968 and the Algorithms and Randomness Center at Georgia Tech and the Max Planck Institute for Mathematics in the Sciences in Leipzig; K. Kohn - the Knut and Alice Wallenberg Foundation: WASP (Wallenberg AI, Autonomous Systems and Software Program) AI/Math initiative; T. Pajdla EU Reg. Dev. Fund IMPACT No. CZ.02.1.01/0.0/0.0/15 003/0000468, EU H2020 ARtwin No. 856994, and EU H2020 SPRING No. 871245 projects at the Czech Institute of Informatics, Robotics and Cybernetics of the Czech Technical University in Prague.

References

1. Agarwal, S., Lee, H., Sturmfels, B., Thomas, R.R.: On the existence of epipolar matrices. International Journal of Computer Vision **121**(3), 403–415 (2017). https://doi.org/10.1007/s11263-016-0949-7, `https://doi.org/10.1007/s11263-016-0949-7`
2. Aholt, C., Oeding, L.: The ideal of the trifocal variety. Math. Comput. **83**(289), 2553–2574 (2014)
3. Aholt, C., Sturmfels, B., Thomas, R.: A hilbert scheme in computer vision. Canadian Journal of Mathematics **65**(5), 961–988 (2013)
4. Fabri et al, R.: Trifocal relative pose from lines at points and its efficient solution. Preprint arXiv:1903.09755 (2019)
5. Alismail, H.S., Browning, B., Dias, M.B.: Evaluating pose estimation methods for stereo visual odometry on robots. In: the 11th International Conference on Intelligent Autonomous Systems (IAS-11) (January 2011)
6. Barath, D.: Five-point fundamental matrix estimation for uncalibrated cameras. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 235–243 (2018)
7. Barath, D., Hajder, L.: Efficient recovery of essential matrix from two affine correspondences. IEEE Trans. Image Processing **27**(11), 5328–5337 (2018)
8. Barath, D., Toth, T., Hajder, L.: A minimal solution for two-view focal-length estimation using two affine correspondences. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 2557–2565 (2017)
9. Byröd, M., Josephson, K., Åström, K.: A column-pivoting based strategy for monomial ordering in numerical gröbner basis calculations. In: European Conference on Computer Vision (ECCV), vol. 5305, pp. 130–143. Springer (2008)
10. Camposeco, F., Sattler, T., Pollefeys, M.: Minimal solvers for generalized pose and scale estimation from two rays and one point. In: ECCV – European Conference on Computer Vision. pp. 202–218 (2016)
11. Chen, H.H.: Pose determination from line-to-plane correspondences: existence condition and closed-form solutions. In: ICCV. pp. 374–378 (1990)
12. Dhome, M., Richetin, M., Lapreste, J.., Rives, G.: Determination of the attitude of 3d objects from a single perspective view. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(12), 1265–1278 (1989)
13. Duff, T., Hill, C., Jensen, A., Lee, K., Leykin, A., Sommars, J.: Solving polynomial systems via homotopy continuation and monodromy. IMA Journal of Numerical Analysis (2018)
14. Duff, T., Kohn, K., Leykin, A., Pajdla, T.: PLMP - Point-line minimal problems in complete multi-view visibility. In: International Conference on Computer Vision (ICCV) (2019)
15. Elqursh, A., Elgammal, A.M.: Line-based relative pose estimation. In: cvpr (2011)
16. Fabbri, R., Giblin, P., Kimia, B.: Camera pose estimation using first-order curve differential geometry. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020)
17. Fabbri, R., Giblin, P.J., Kimia, B.B.: Camera pose estimation using first-order curve differential geometry. In: Proceedings of the European Conference in Computer Vision (2012)
18. Fabbri, R., Kimia, B.B.: Multiview differential geometry of curves. International Journal of Computer Vision **120**(3), 324–346 (2016)

19. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24**(6), 381–395 (1981)
20. Grayson, D.R., Stillman, M.E.: Macaulay2, a software system for research in algebraic geometry. Available at http://www.math.uiuc.edu/Macaulay2/
21. Hartley, R., Li, H.: An efficient hidden variable approach to minimal-case camera motion estimation. IEEE PAMI **34**(12), 2303–2314 (2012)
22. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge, 2nd edn. (2003)
23. Hartley, R.I.: Lines and points in three views and the trifocal tensor. International Journal of Computer Vision **22**(2), 125–140 (1997)
24. Johansson, B., Oskarsson, M., Åström, K.: Structure and motion estimation from complex features in three views. In: ICVGIP 2002, Proceedings of the Third Indian Conference on Computer Vision, Graphics & Image Processing, Ahmadabad, India, December 16-18, 2002 (2002)
25. Joswig, M., Kileel, J., Sturmfels, B., Wagner, A.: Rigid multiview varieties. IJAC **26**(4), 775–788 (2016). https://doi.org/10.1142/S021819671650034X, `https://doi.org/10.1142/S021819671650034X`
26. Kahl, F., Heyden, A., Quan, L.: Minimal projective reconstruction including missing data. IEEE Trans. Pattern Anal. Mach. Intell. **23**(4), 418–424 (2001). https://doi.org/10.1109/34.917578, `https://doi.org/10.1109/34.917578`
27. Kileel, J.: Minimal problems for the calibrated trifocal variety. SIAM Journal on Applied Algebra and Geometry **1**(1), 575–598 (2017)
28. Kneip, L., Scaramuzza, D., Siegwart, R.: A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In: CVPR – IEEE Conference on Computer Vision and Pattern Recognition. pp. 2969–2976 (2011)
29. Kneip, L., Siegwart, R., Pollefeys, M.: Finding the exact rotation between two images independently of the translation. In: ECCV – European Conference on Computer Vision. pp. 696–709 (2012)
30. Kuang, Y., Åström, K.: Pose estimation with unknown focal length using points, directions and lines. In: IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013. pp. 529–536 (2013)
31. Kuang, Y., Åström, K.: Stratified sensor network self-calibration from tdoa measurements. In: 21st European Signal Processing Conference (2013)
32. Kukelova, Z., Bujnak, M., Pajdla, T.: Automatic generator of minimal problem solvers. In: European Conference on Computer Vision (ECCV) (2008)
33. Kukelova, Z., Kileel, J., Sturmfels, B., Pajdla, T.: A clever elimination strategy for efficient minimal solvers. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2017)
34. Larsson, V., et. al: Automatic generator of minimal problems. http://www2.maths.lth.se/matematiklth/personal/viktorl/code/basis_selection.zip (2015)
35. Larsson, V., Åström, K., Oskarsson, M.: Efficient solvers for minimal problems by syzygy-based reduction. In: Computer Vision and Pattern Recognition (CVPR) (2017)
36. Larsson, V., Åström, K., Oskarsson, M.: Efficient solvers for minimal problems by syzygy-based reduction. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 2383–2392 (2017)

37. Larsson, V., Åström, K., Oskarsson, M.: Polynomial solvers for saturated ideals. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. pp. 2307–2316 (2017)
38. Larsson, V., Kukelova, Z., Zheng, Y.: Making minimal solvers for absolute pose estimation compact and robust. In: International Conference on Computer Vision (ICCV) (2017)
39. Larsson, V., Oskarsson, M., Åström, K., Wallis, A., Kukelova, Z., Pajdla, T.: Beyond grobner bases: Basis selection for minimal solvers. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 3945–3954 (2018), http://openaccess.thecvf.com/content_cvpr_2018/html/Larsson_Beyond_Grobner_Bases_CVPR_2018_paper.html
40. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2), 91–110 (2004)
41. Ma, Y., Huang, K., Vidal, R., Kosecka, J., Sastry, S.: Rank conditions on the multiple-view matrix. International Journal of Computer Vision **59**(2), 115–137 (2004)
42. Matas, J., Obdrzálek, S., Chum, O.: Local affine frames for wide-baseline stereo. In: 16th International Conference on Pattern Recognition, ICPR 2002, Quebec, Canada, August 11-15, 2002. pp. 363–366. IEEE Computer Society (2002)
43. Matas, J., Obdrzálek, S., Chum, O.: Local affine frames for wide-baseline stereo. In: 16th International Conference on Pattern Recognition, ICPR 2002, Quebec, Canada, August 11-15, 2002. pp. 363–366 (2002)
44. Miraldo, P., Araujo, H.: Direct solution to the minimal generalized pose. Cybernetics, IEEE Transactions on **45**(3), 418–429 (2015)
45. Miraldo, P., Dias, T., Ramalingam, S.: A minimal closed-form solution for multi-perspective pose estimation using points and lines. In: Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XVI. pp. 490–507 (2018)
46. Mirzaei, F.M., Roumeliotis, S.I.: Optimal estimation of vanishing points in a manhattan world. In: International Conference on Computer Vision (ICCV) (2011)
47. Nistér, D.: An efficient solution to the five-point relative pose problem. IEEE Transactions on Pattern Analysis and Machine Intelligence **26**(6), 756–770 (Jun 2004)
48. Nistér, D., Naroditsky, O., Bergen, J.: Visual odometry. In: Computer Vision and Pattern Recognition (CVPR). pp. 652–659 (2004)
49. Nistér, D., Schaffalitzky, F.: Four points in two or three calibrated views: Theory and practice. International Journal of Computer Vision **67**(2), 211–231 (2006)
50. Oskarsson, M., Åström, K., Overgaard, N.C.: Classifying and solving minimal structure and motion problems with missing data. In: International Conference on Computer Vision (ICCV). pp. 628–634. IEEE Computer Society (2001). https://doi.org/10.1109/ICCV.2001.10072, http://doi.ieeecomputersociety.org/10.1109/ICCV.2001.10072
51. Oskarsson, M., Zisserman, A., Åström, K.: Minimal projective reconstruction for combinations of points and lines in three views. Image Vision Comput. **22**(10), 777–785 (2004)
52. Raguram, R., Chum, O., Pollefeys, M., Matas, J., Frahm, J.: USAC: A universal framework for random sample consensus. IEEE Transactions on Pattern Analysis Machine Intelligence **35**(8), 2022–2038 (2013)
53. Ramalingam, S., Bouaziz, S., Sturm, P.: Pose estimation using both points and lines for geo-localization. In: ICRA. pp. 4716–4723 (2011)

54. Ramalingam, S., Sturm, P.F.: Minimal solutions for generic imaging models. In: CVPR – IEEE Conference on Computer Vision and Pattern Recognition (2008)
55. Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Pajdla, T., Sivic, J.: Neighbourhood consensus networks (2018)
56. Salaün, Y., Marlet, R., Monasse, P.: Robust and accurate line- and/or point-based pose estimation without manhattan assumptions. In: European Conference on Computer Vision (ECCV) (2016)
57. Sattler, T., Leibe, B., Kobbelt, L.: Efficient & effective prioritized matching for large-scale image-based localization. IEEE Trans. Pattern Anal. Mach. Intell. **39**(9), 1744–1756 (2017)
58. Saurer, O., Pollefeys, M., Lee, G.H.: A minimal solution to the rolling shutter pose estimation problem. In: Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on. pp. 1328–1334. IEEE (2015)
59. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
60. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3D. In: ACM SIGGRAPH (2006)
61. Snavely, N., Seitz, S.M., Szeliski, R.: Modeling the world from internet photo collections. International Journal of Computer Vision (IJCV) **80**(2), 189–210 (2008)
62. Stewenius, H., Engels, C., Nistér, D.: Recent developments on direct relative orientation. ISPRS J. of Photogrammetry and Remote Sensing **60**, 284–294 (2006)
63. Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: InLoc: Indoor visual localization with dense matching and view synthesis. In: CVPR (2018)
64. Trager, M., Ponce, J., Hebert, M.: Trinocular geometry revisited. International Journal Computer Vision pp. 1–19 (March 2016)
65. Trager, M.: Cameras, Shapes, and Contours: Geometric Models in Computer Vision. (Caméras, formes et contours: modèles géométriques en vision par ordinateur). Ph.D. thesis, École Normale Supérieure, Paris, France (2018)
66. Ventura, J., Arth, C., Lepetit, V.: An efficient minimal solution for multi-camera motion. In: International Conference on Computer Vision (ICCV). pp. 747–755 (2015)
67. Xia, G., Delon, J., Gousseau, Y.: Accurate junction detection and characterization in natural images. International Journal of Computer Vision **106**(1), 31–56 (2014)