



PL₁P: Point-Line Minimal Problems under Partial Visibility in Three Views

Timothy Duff¹ · Kathlén Kohn² · Anton Leykin³ · Tomas Pajdla⁴

Received: 13 November 2022 / Accepted: 25 December 2023 / Published online: 10 March 2024
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

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 140,616 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 74,575 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 fewer than 300 solutions for generic data.

Keywords Minimal problems · Calibrated cameras · 3D reconstruction

1 Introduction

Minimal problems (Agarwal et al., 2017; Byröd et al., 2008; Bhayani et al., 2020; Barath et al., 2017; Barath, 2018; Barath & Hajder, 2018; Camposeco et al., 2016; Elqursh & Elgammal, 2011; Hartley & Li, 2012; Kukulova et al.,

2008; Kuang & Åström, 2013a; 2013b; Kneip et al., 2012; Larsson et al., 2017a; 2017b; 2017c; Mirzaei & Roumeliotis, 2011; Miraldo et al., 2018; Nistér, 2004; Ramalingam & Sturm, 2008; Stewenius et al., 2006; Saurer et al., 2015; Salaün et al., 2016; Ventura et al., 2015), 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 (Snavely et al., 2006, 2008; Schönberger & Frahm, 2016), image matching (Rocco et al., 2018), visual odometry and localization (Nistér et al., 2004; Alismail et al., 2011; Sattler et al., 2017; Taira et al., 2018). 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 (Duff et al., 2019). 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

Communicated by Takayuki Okatani.

✉ Timothy Duff
timduff@uw.edu

✉ Kathlén Kohn
kathlen@kth.se

✉ Anton Leykin
leykin@math.gatech.edu

✉ Tomas Pajdla
pajdla@cvut.cz

¹ Department of Mathematics, University of Washington, Padelford Hall, Seattle, WA 98195, USA

² Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

³ School of Mathematics, Georgia Tech, Atlanta, GA, USA

⁴ CIIRC, Czech Technical University in Prague, Jugoslávských partyzánů 3, 16000 Prague, Czechia

that they can be *reduced* to 140,616 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 (Fischler & Bolles, 1981; Raguram et al., 2013), 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 Kileel (2017) are non-minimal yet they all are camera-minimal.

As an example, consider the problem from Kileel (2017) with 3 point-point-point (PPP) and 1 point-point-line (PPL) correspondence; see Fig. 1a. 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. To make the problem minimal, the line would have to be seen in one more view; see Fig. 1b featuring pictogram encoding of problems in three views that we use in this article. In applications, e.g. Hruby et al. (2022), the problem in Fig. 1a is used as a relaxation of the 4-points-in-3-views problem, where an auxiliary line is introduced in one of the images but there is no need to recover the corresponding line in 3D space.

From the minimal problem in the bottom of Fig. 1b, we can remove the line from the second view without losing camera-minimality. We apply the same principle to every minimal problem: We delete superfluous features in images that can be removed without losing camera-minimality to obtain a simplest camera-minimal problem. Those we call *terminal camera-minimal* problems.

We show that, up to relabeling cameras, there are 74,575 such problems. 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 74,575 terminal camera-minimal problems were known—66 problems listed in Kileel (2017) plus 10 additional cases from Duff et al. (2019)—the remaining 74,499, 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 Duff et al. (2019) 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 point¹. 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 (Lowe, 2004) and LAF (Matas et al., 2002) provide quivers (points with one direction attached), which can be interpreted as lines through the points and used to compute relative camera poses (Fabbri et al., 2023).

2 Previous Work


A large number of minimal problems have appeared in the literature. See eg. references above, (Duff et al., 2019; Kukulova et al., 2017; Kileel, 2017; Larsson et al., 2018), and all references therein for work on general minimal problems. 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 (Hartley, 1997). Point-line incidences in the uncalibrated setup are introduced in Johansson et al. (2002) 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 Ma et al. (2004). Non-incident points and lines in three uncalibrated images also appear in Oskarsson et al. (2001; 2004) and Kahl et al., (2001). The cases of four points and three lines, two points and six lines, and nine lines are studied (Larsson et al., 2017); constructs a solver for nine lines. Works (Fabbri et al., 2012, 2020; Fabbri & Kimia, 2016) look at lines incident to points which arise from tangent lines to curves and (Fabbri et al., 2023) presents a solver for that case. Results (Agarwal et al., 2017; Aholt & Oeding, 2014; Aholt et al., 2013; Joswig et al., 2016; Duff et al., 2019; Trager, 2018; Trager et al., 2016) introduced some of the techniques that are useful for classifying classes of minimal problems.

Work (Kileel, 2017) classifies camera-minimal problems in 3 calibrated views that can be formulated with linear constraints on the trifocal tensor (Hartley & Zisserman, 2003). It presents 66 camera-minimal problems, all covered in our classification as terminal camera-minimal problems. Among

¹ Under this restriction, in two cameras, the only reduced (camera-)minimal problem is the five-point problem. See Sect. 10 for an explanation.

Fig. 1 Relaxation of the 4-points-in-3-views problem



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 Duff et al. (2019). It gives 30 minimal problems. Among them, 17 problems include exactly three cameras but only 12 of them (3002₁, 3002₂, 3010₀, 2005₃, 2005₄, 2005₅, 2013₂, 2013₃, 2021₁, 1024₄, 1032₂, 1040₀ in Table 1 of Duff et al. (2019)) meet our restrictions on incidences. These 12 cases are all terminal camera-minimal as well as reduced minimal. Notice that the remaining 5 problems (3100₀, 2103₁, 2103₂, 2103₃, 2111₁ in Table 1 of Duff et al. (2019)) 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 Kileel (2017) and Duff et al. (2019) to a much larger class of problems in three calibrated views under partial multi-view visibility.

3 Glossary of Assumptions, Properties and Concepts

Here we provide a glossary of assumptions, properties and concepts used in the paper. Our aim is to give a concise and intuitive exposition of concepts used.

1. Realizability of $\mathcal{I}$ —the incidence relations are realizable by some point-line arrangement in $\mathbb{R}^3$.
2. Completeness of $\mathcal{I}$ —every incidence which is automatically implied by the incidences in $\mathcal{I}$ must also be contained in $\mathcal{I}$.
3. Completeness of $\mathcal{O}$ —if a camera observes two lines that meet according to $\mathcal{I}$, then it observes their point of intersection.
4. Genericity of $\mathcal{C}_m$ —the points and lines in the views are in generic positions with respect to the specified incidences $\mathcal{I}$, i.e. random noise in image measurements does not

change the number of solutions when the incidences from 3D are not broken by noise in images.

5. Complete visibility—all points and lines are observed in all images and all observed information is used to formulate minimal problems.
6. Partial visibility—some points and lines may be forgotten when projecting into images.
7. PL_kP—each line in 3D is incident to at most k points.
8. Dominant—“almost all” 2D images have a solution (i.e. a 3D arrangement and cameras yielding the images)
9. Balanced—the number of DOF (i.e. the dimensions of varieties) describing cameras and preimages in 3D are equal to the number of DOF describing image measurements.
10. Minimal—balanced and dominant. This is equivalent to that “almost all” images in 2D have a positive finite number of solutions.
11. Camera-Minimal—the number of solutions in camera parameters is finite and positive (despite that infinitely many solutions may exist for 3D structures).
12. Reduced problem—a problem where all viewed features imply nontrivial constraints on the cameras
13. Degree of a minimal problem—the number of point-line configurations and camera poses consistent with generic image data
14. Camera-degree of a camera-minimal problem—the number of camera poses consistent with generic image data

4 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

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

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

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

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

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.

Let us briefly explain why the assumptions on $\mathcal{I}$ are needed. Note that, for instance, the realizability assumption implies that two distinct lines cannot have more than one point in common. In other words, the incidence relation $\mathcal{I} = \{(1, 1), (1, 2), (2, 1), (2, 2)\}$ is non-realizable. We also remark that related notions of realizability (also called representability) are well-studied in matroid theory. The Fano plane is perhaps the best-known example of a non-realizable point-line relation. We refer to Oxley (2022, Chapter 6) for a detailed treatment of these concepts. For an example of an incidence relation that is not complete, consider 4 lines contained in a common plane where we do not specify one of their intersection points:

$$\mathcal{I} = \{(1, 1), (1, 2), (2, 2), (2, 3), (3, 3), (3, 1), (4, 1), (4, 4), (5, 2), (5, 4)\}.$$

Since the fourth line also intersects the third line in a point, two incidences (6,3) and (6,4) are missing.

Consider now our assumption on $\mathcal{O}$ in Definition 6. This is completely 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 $X_1, \dots, X_p$ and l lines $L_1, \dots, L_l$ 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

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

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

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

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

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

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

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 (Hartley & Zisserman, 2003), we set $P_1 = [I \mid 0]$ and the first coordinate of t_2 to 1. So our *camera configurations* are parameterized by

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

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 $\Phi_{p,l,\mathcal{I},\mathcal{O}}^{-1}(x, \ell) \neq \emptyset$, and

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

- the preimage $\Phi_{p,l,\mathcal{I},\mathcal{O}}^{-1}(x, \ell)$ of a generic element (x, ℓ) in $\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ is finite.

Remark 3 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 (Nistér & Schaffalitzky, 2006).

Over the complex numbers, the cardinality of the preimage $\Phi_{p,l,\mathcal{I},\mathcal{O}}^{-1}(x, \ell)$ 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 4 Let $\gamma : \mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m \rightarrow \mathcal{C}_m$ 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
- $\gamma(\Phi_{p,l,\mathcal{I},\mathcal{O}}^{-1}(x, \ell))$ 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 $\gamma(\Phi_{p,l,\mathcal{I},\mathcal{O}}^{-1}(x, \ell))$ is the *camera-degree* of $(p, l, \mathcal{I}, \mathcal{O})$.

Remark 5 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 (Duff et al., 2019, Cor. 2).

In Duff et al. (2019), 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 Duff et al. (2019) to all point-line problems with *partial visibility*. A first obstruction is that there are minimal point-line problems for arbitrarily many cameras, whereas the result in Duff et al. (2019) 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*; the case of two views is much simpler and discussed in Sect. 10. We also restrict our attention to point-line problems satisfying the following assumption:

Definition 6 We say that a point-line problem is a *PL_kP* 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 PL₁P 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 PL₁Ps is the following fact.

Lemma 7 *The degree and camera-degree of a minimal PL₁P coincide.*

The assumption of minimality cannot be completely relaxed, as Example 31 shows that Lemma 7 fails for general point-line problems. All proofs of the results above and those to follow may be found in Sect. 12, towards the end of the manuscript.

We will see that there are *infinitely many* (camera-)minimal PL₁Ps in three views. However, we can partition them into finitely many classes such that all PL₁Ps 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* PL₁Ps as the canonical representatives of the finitely many classes of minimal PL₁Ps we aim to find (see Sect. 5).

Step 2: Basic principles from algebraic geometry brought up in Duff et al. (2019) imply

Lemma 8 *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 PL₁Ps in three views that contains all reduced minimal PL₁Ps (see Sect. 6).

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 PL₁Ps in three views (see Sect. 7).

Step 4: For each of the finitely many balanced PL₁Ps 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 PL₁Ps in three views (see Sect. 8).

In addition to the classification, we compute the camera-degrees of all reduced camera-minimal PL₁Ps in three views whose camera-degrees are less than 300 (see Sect. 9 for this and related results on natural subfamilies of PL₁Ps.)

5 Reduced PL₁Ps

From a given PL₁P $(p, l, \mathcal{I}, \mathcal{O})$ we can obtain a new PL₁P by forgetting some points and lines, both in space and in the views. Formally, if $\mathcal{P}' \subset \{1, \dots, p\}$ and $\mathcal{L}' \subset \{1, \dots, l\}$ are the sets of points and lines which are *not* forgotten, the new PL₁P 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 $\mathcal{O}' = ((\mathcal{P}'_1, \mathcal{L}'_1), \dots, (\mathcal{P}'_m, \mathcal{L}'_m))$, 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 9 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 10 If a PL₁P is minimal and reducible to another PL₁P, then both are minimal and have the same degree.

We can partition *all* (infinitely many) minimal PL₁Ps in three views into finitely many classes using this reduction process. Each class is represented by a unique PL₁P that is *reduced*, i.e. not reducible to another PL₁P.

Theorem 11 A minimal PL₁P $(p, l, \mathcal{I}, \mathcal{O})$ in three views is reducible to a unique reduced PL₁P $(p', l', \mathcal{I}', \mathcal{O}')$. 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 
- 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 PL₁P in three views results in a minimal PL₁P.

Hence, it is enough to classify all reduced minimal PL₁Ps. We will see that there are finitely many reduced minimal PL₁Ps in three views. To count them, we need to understand how they look.

Theorem 12 A reduced minimal PL₁P 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.

6 Balanced PL₁Ps








A reduced minimal PL₁P in three views is uniquely determined by a *signature*, a vector consisting of 27 numbers $(c_7, \dots, c_3, c_{2,0}, c_{2,1}^a, \dots, f)$, 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 12, the local features in such a PL₁P are free lines or points with at most seven pins. We denote by f the number of free lines and write $c_3, c_4, \dots, c_7$ for the numbers of points with three, four, ..., seven pins. By Theorem 12, these local features are completely observed by the cameras; moreover at most c_i can equal 1. 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, $c_{2,1}^a, c_{2,1}^b, c_{2,1}^c$ 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 $c_7, \dots, f$ describes a unique PL₁P in three views, which is reduced by construction (see Thm. 11 and 12) but not necessarily minimal.

Due to Lemma 8, every minimal PL₁P $(p, l, \mathcal{I}, \mathcal{O})$ 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

Table 1 How points with two / one / zero pins and free lines can be observed in the three views of a reduced minimal PL₁P (up to permuting the views)

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

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 Sect. 6

PL₁P in three views satisfies

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

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 $k = 3, \dots, 7$ is completely observed by the cameras, it has $3(2 + k)$ degrees of freedom in the three views. Therefore, we have

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

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 PL₁P in three views is

$$11 = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}) - \dim(\mathcal{X}_{p,l,\mathcal{I}}), \text{ i.e.}$$

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

The linear equation (2) has 845161 non-negative integer solutions. Each of these is a signature that represents a PL₁P in three views which is reduced and balanced. Thus, it remains to check which of the 845161 signatures represent *minimal* PL₁Ps. Some of the 845161 solutions of (2) yield *label-equivalent* PL₁Ps, i.e. PL₁Ps which are the same up to relabeling the three views. It turns out that there are **143494** such label-equivalence classes of PL₁Ps given by solutions to (2). So all in all, we have to check 143494 PL₁Ps for minimality, namely one representative for each label-equivalence class. For further details about this enumeration procedure, we refer to Sect. 13.1.

7 Camera-Minimal PL₁Ps


As in the case of minimal problems, we can understand all camera-minimal PL₁Ps from the reduced ones (see also Theorem 20).

Theorem 13 *If a PL₁P is reducible to another PL₁P, 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 PL₁Ps look, in comparison to reduced minimal PL₁Ps as described

in Theorem 12, 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_1P s that are camera-minimal but not minimal.


Theorem 14 *The local features of a reduced camera-minimal PL_1P in three views are viewed as described in Theorem 12 plus as in the following three additional cases:*



Remark 15 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_1P s. By adding more constraints to images, we make configurations in space uniquely reconstructible.

Theorem 16 *The following replacements in images lift a reduced camera-minimal PL_1P in three views to a reduced minimal PL_1P (cf. Thm. 14 and Table 1):*



Moreover, the camera-degrees of both PL_1P s are the same.

This has two important implications for classifying (camera-)minimal PL_1P s. First, reversing the replacements in Theorem 16 transforms each reduced camera-minimal PL_1P in three views into a *terminal* PL_1P of the same camera-degree.

Definition 17 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_1P s in three views, it is enough to find the terminal ones.

Secondly, Theorem 16 implies for minimal PL_1P s the following corollary.

Corollary 18 *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 Sect. 6, we defined two PL_1P s to be label-equivalent if they are the same up to relabeling the views. We note that the swap described in Corollary 18 does *not* preserve the label-equivalence class of a PL_1P . Instead, we say that two PL_1P s 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 18. We conclude that either all PL_1P s 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 16 yields the following.

Corollary 19 *The swap&label-equivalence classes of reduced minimal PL_1P s in three views are in a camera-degree preserving one-to-one correspondence with the label-equivalence classes of terminal camera-minimal PL_1P s in three views.*

Hence, we do not have to check minimality for all 143494 label-equivalence classes of PL_1P s given by solutions to (2), that we found at the end of Sect. 6. 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. So to find *all* (camera-)minimal PL_1P s in three views, we only have to check 76446 PL_1P s for minimality, namely one representative for each of the swap&label-equivalence classes. Finally, we present the analog to Theorem 11 and describe how all camera-minimal PL_1P s are obtained from the reduced camera-minimal ones.

Theorem 20 *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 11
- every line (free or pin) that is not observed in any view
- every free line that is observed in exactly one view to reduce



- 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 .

Table 2 Distribution of camera-degrees of terminal camera-minimal PL_1 Ps in three calibrated views with: (a) camera-degree less than 300, (b) at most one pin per point and camera-degree less than 500

(a)															
deg	64	80	144	160	216	224	240	256	264	272	288				
# pr	13	9	3	547	7	2	159	2	2	11	4				
(b)															
deg	80	160	216	240	256	264	272	288	304	312	320	352	360	368	376
# pr	9	173	4	80	2	2	2	1	5	2	213	3	9	3	1
deg	384	392	400	408	416	424	432	448	456	464	472	480	488	496	
# pr	2	9	14	2	6	10	2	7	11	4	1	96	12	9	

8 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 8, as explained in Duff et al. (2019). 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*³ 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 Duff et al., (2019, Algorithm 1).

The computation described above detects non-minimality for 2878 of the 143,494 label-equivalence classes of PL_1 Ps given by solutions to (2). Among the 76,446 swap&label-equivalence classes, 1871 are not minimal.

Result 21 *In three calibrated views, up to relabeling cameras, there are*

- 140, 616 = 143, 494 – 2878 reduced minimal PL_1 Ps and
- 74, 575 = 76, 446 – 1871 terminal camera-minimal PL_1 Ps.

9 Computing Degrees

From the perspective of solving minimal problems, it is highly desirable to compute all degrees of our minimal PL_1 Ps. In particular, we wish to identify problems with small degrees that may be of practical interest. Since some problems in this list are known from prior work (Kileel, 2017) to have large degrees (> 1000), our main technique is a

monodromy approach based on numerical homotopy continuation. Our implementation in Macaulay2 (Grayson & Stillman, xxx; Duff et al., 2018) is similar to that used in prior work (Duff et al., 2019). The next result shows that there are many interesting problems with small degrees that are solvable with existing solving technology (Fabbri et al., 2023; Larsson et al., 2015).

Result 22 *There are 759 (up to relabeling cameras) terminal camera-minimal PL_1 Ps 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 23 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 23 generalizes Definition 6 of PL_1 Ps to get a hierarchy of subfamilies of point-line problems: $PL_0P \subset PL_1P \subset PL_2P \subset \dots$

For instance, PL_0 Ps consist of free points and free lines only⁴. The family of PL_1 Ps contains many problems involving at most one pin per point—see Result 24 and Table 2(b). Such features are readily provided by the SIFT (Lowe, 2004) detector. PL_1 Ps can also include features with two pins per point, which are readily provided by the LAF (Matas et al., 2002) detector, which can also be used to get PL_2 Ps if all 3 LAF points are used. More complex incidences (e.g. PL_3 Ps) can be obtained from line t-junction detectors (Xia et al., 2014).

Result 24 *There are 9533 (up to relabeling cameras) terminal camera-minimal PL_1 Ps 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).*

Result 25 *There are 51 (up to relabeling cameras) reduced minimal PL_0 Ps in three calibrated views. They are depicted together with their degrees in Table 3.*

³ Since we are testing minimality, being minimal is the positive outcome. See Sect. 13.2 for detailed explanation why false positives cannot occur.

⁴ We note that reduced minimal PL_0 Ps are terminal.

Table 3 Reduced minimal PL_0 Ps and their degrees. Points not visible in a given view are indicated in grey. Five-point subproblems are indicated in red

216	160	384	256	80	416	568	320	320	768	360	512	616	160	528	776	984
320	320	720	1024	1456	400	560	640	1376	920	744	1416	1608	160	800	1480	1656
2232	320	320	1040	1360	2016	2568	400	560	640	1200	1920	2688	400	800	960	2000

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 (Kneip et al., 2011)): 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 (Chen, 1990; Dhome et al., 1989; Miraldo & Araujo, 2015; Miraldo et al., 2018; Ramalingam et al., 2011).

Result 26 We determined all PL_1 Ps 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 Sect. 11)
- 3648 are terminal camera-minimal.

10 Note on Partial Visibility in Two Views

Let us recall that we consider reduced minimal and reduced camera-minimal PL_1 Ps in *three* views, since there is only one such PL_1 P in two views, namely the five-points problem.

To argue this, we first notice that free lines cannot occur in a reduced PL_1 P in two views. Indeed, if there was a free line—no matter if it is observed in both views, one view, or not at all—we could forget the free line both in 3D and in

the images to reduce the PL_1 P (i.e. forgetting the free line satisfies Definition 9 of reducibility).

Secondly, we argue that pins cannot appear in a reduced PL_1 P in two views. To see this we distinguish several cases, depending on how often a hypothetical pin is observed in the views.

- If there exists a pin in 3D that is not observed in any of the two views, then the PL_1 P would be reducible by simply forgetting that pin.
- If a pin in 3D is observed in exactly one of the two views, it could either appear as a pin or as a free line in that view (i.e. depending on whether that view also observes the point of the pin or not).
 1. If the single view seeing the pin also observes its point, the PL_1 P is reducible by forgetting the pin.
 2. If the single view seeing the pin does not observe its point, that view cannot observe any of the other pins of that point either (due to our assumptions on $\mathcal{O}$ at the end of Definition 1). Hence, no matter how the other view observes that point and its other pins, the PL_1 P is reducible by forgetting the point together with all of its pins.
- If a pin in 3D is observed in both views, it could appear as either a pin or a free line in either of the views.
 1. If both views also observe the point of the pin, the PL_1 P is reducible by forgetting the pin.
 2. If at least one of the two views does not observe the point of the pin, then we may argue as in case 2 above that such a view cannot observe any of the other pins of that point either. Once again, the PL_1 P is reducible by forgetting the point together with all of its pins.

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

Fig. 2 Two PL₁Ps and their joint camera maps related by a forgetting map Π . The upper PL₁P is reducible to the lower PL₁P iff the forgetting map Π is feasible and satisfies the lifting property

Finally, a PL₁P which observes a point without pins in at most one of its views is reducible by forgetting that point. All in all, we conclude that reduced PL₁Ps in two views can only consist of points without pins which are observed in both views. The only such (camera-)minimal problem is the five-points problem.

11 Note on Camera Registration Problems

Among the minimal problems we discovered that there are extensions of the classical five-point problem in two views that can be interpreted as camera registration problems; see Table 4.

Remark 27 Extension of reduced camera-minimal problems by means of camera registration problems is one simple construction for (camera-)minimal problems in arbitrarily many views.

12 Proofs

12.1 Proofs of Theorems 10 and 13

We start by investigating the implications of the two conditions in Definition 9 of reducibility. When obtaining a new PL₁P $(p', l', \mathcal{I}', \mathcal{O}')$ from a given PL₁P $(p, l, \mathcal{I}, \mathcal{O})$ by forgetting some of its points and lines, the induced projections Π and π between the domains and codomains of the joint camera maps yield the commutative diagram in Fig. 2. We call Π the *forgetting map*, and note that the map π between the image varieties is completely determined by the forgetting map Π . If the forgetting map Π satisfies the first condition in Definition 9 (i.e. for each forgotten point, at most one of its pins is kept), we say that it is *feasible*. The second condition in Definition 5 we refer to as the *lifting property*.

Lemma 28 *If the forgetting map Π is feasible, then both Π and π are surjective and have irreducible fibers⁵ of equal dimension.*

⁵ A fiber of a map is the preimage over a single point in its codomain.

Proof We first show that Π is surjective and has irreducible fibers of the same dimension. We may assume that Π forgets either a single point or a single line, as a every feasible forgetting map is the composition of several feasible forgetting maps which forget either one point or one line. Moreover, the composition of several surjective maps with irreducible fibers of equal dimension is again surjective and has irreducible fibers of the same dimension.

Let us first assume that Π forgets a single line. If this line is the pin of a point, each fiber of Π is the set of lines in $\mathbb{P}^3$ through that point; this set is isomorphic to $\mathbb{P}^2$. If Π forgets a free line, its fibers are isomorphic to $\mathbb{G}_{1,3}$. In both cases, the forgetting map Π is surjective.

Now let us assume that Π forgets a single point. As Π is feasible, this point is incident to at most one line in space. Depending on if it is incident to one or zero lines, the fibers of Π are isomorphic to either $\mathbb{P}^1$ or $\mathbb{P}^3$. In either case, Π is surjective.

To show that π is surjective with irreducible fibers of equal dimension, we apply a nearly-identical proof as above to the m factors of π corresponding to the different views. $\square$

For a feasible forgetting map Π , we write $\dim(\text{fiber}(\Pi))$, resp. $\dim(\text{fiber}(\pi))$, for the dimensions of the fibers of Π , resp. π . Moreover, we denote by $\text{cdeg}(p, l, \mathcal{I}, \mathcal{O})$ the camera-degree of a point-line problem $(p, l, \mathcal{I}, \mathcal{O})$. We will also use this notation for point-line problems which are not camera-minimal: in that case, the camera-degree is either zero (if the joint camera map is not dominant) or ∞ .

Lemma 29 *Consider Fig. 2 with a feasible forgetting map Π .*

1. *If the upper joint camera map Φ is dominant, so is the lower one Φ' .*
2. $\text{cdeg}(p', l', \mathcal{I}', \mathcal{O}') \geq \text{cdeg}(p, l, \mathcal{I}, \mathcal{O})$.
3. *If $(p, l, \mathcal{I}, \mathcal{O})$ is minimal, then $\dim(\text{fiber}(\Pi)) \leq \dim(\text{fiber}(\pi))$.*

Proof For a generic image $(x', \ell') \in \mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}$ of the lower PL₁P, we consider a generic image $(x, \ell) \in \pi^{-1}(x', \ell')$ of the upper PL₁P. Every solution $S \in \Phi^{-1}(x, \ell)$ of the upper PL₁P yields a solution $\Pi(S)$ of the lower PL₁P with the same cameras. This shows the first two parts of the assertion.

For the third part, since the joint camera map Φ of a minimal PL₁P is dominant, we use part 1 to see that the lower joint camera map Φ' is also dominant. This implies:

$$\begin{array}{ccc}
\dim(\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_m) & \geq & \dim(\mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}) \\
\parallel & & \parallel \\
\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m) - \dim(\text{fiber}(\Pi)) & & \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}) - \dim(\text{fiber}(\pi))
\end{array}$$

Since each minimal PL₁P is balanced, i.e. $\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m) = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}})$, this concludes the proof. $\square$

Table 4 Camera registration: minimal problems that can be determined as relative pose problems for two views appended with a minimal registration problem for the third camera

Lemma 30 Consider Fig. 2 with a feasible forgetting map Π that satisfies the lifting property.

$$\dim(\pi^{-1}(\Phi'(S'))) = \dim(\text{fiber}(\pi)).$$

□

1. Φ is dominant $\Leftrightarrow \Phi'$ is dominant.
2. $\text{cdeg}(p', l', \mathcal{I}', \mathcal{O}') = \text{cdeg}(p, l, \mathcal{I}, \mathcal{O})$.
3. $\dim(\text{fiber}(\Pi)) \geq \dim(\text{fiber}(\pi))$.

Proof We can pick a generic image of the upper PL_1P , by first choosing a generic image $(x', \ell') \in \mathcal{Y}_{p', l', \mathcal{I}', \mathcal{O}'}$ of the lower PL_1P and then considering a generic image $(x, \ell) \in \pi^{-1}(x', \ell')$ in its fiber. If Φ' is dominant, a generic solution $S' \in \Phi'^{-1}(x', \ell')$ of the lower PL_1P can be lifted to a solution $S \in \Phi^{-1}(x, \ell)$ of the upper PL_1P with the same cameras. Together with Lemma 29, this shows the first two parts of the assertion. For the third part, we consider a generic element $S' \in \mathcal{X}_{p', l', \mathcal{I}'} \times \mathcal{C}_m$. By the lifting property, we have that

$$\dim(\text{fiber}(\Pi)) = \dim(\Pi^{-1}(S')) \geq$$

Proof of Lemma 7 Let $(p, l, \mathcal{I}, \mathcal{O})$ be a minimal PL_1P with camera-degree d . Fix $(x, \ell) \in \mathcal{Y}_{p, l, \mathcal{I}, \mathcal{O}}$ generic and let $\gamma(\Phi_{p, l, \mathcal{I}, \mathcal{O}}^{-1}(x, \ell)) = \{P_1, \dots, P_d\}$. We may assume that $\Phi_{p, l, \mathcal{I}, \mathcal{O}}^{-1}(x, \ell)$ is a nonempty finite set; we wish to show it contains d elements. We begin by lifting $P_1, \dots, P_d$ to partial solutions in $\mathbb{G}_{1,3}^l \times \mathcal{C}_m$, i.e. by reconstructing lines in 3D. Our first observation is that each of the l world lines must be viewed at least twice. Indeed, if a line was observed at most once, then even if it is a pin (i.e. it passes through a point in 3D) there would be an at least one-dimensional family of world lines yielding the same view (and possibly passing through the pin point). This contradicts our assumption that the PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ is minimal.

Now since each world line is observed in at least two views, it can be uniquely recovered from fixed cameras; in other words, each camera solution P_i extends uniquely to a partial solution $(L_i, P_i) \in \mathbb{G}_{1,3}^l \times \mathcal{C}_m$. To conclude that there are exactly d lifts

$$((X_1, L_1), P_1), \dots, ((X_d, L_d), P_d) \in \Phi_{p,l,\mathcal{I},\mathcal{O}}^{-1}(x, \ell),$$

we note that the condition that each X_i is a lifted solution of the partial solution (L_i, P_i) is given by linear equations depending on P_i and L_i . Thus, there is either a unique or infinitely many such X_i , but since the PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ is minimal, the latter cannot hold. $\square$

Example 31 Lemma 7 does not hold for point-line problems in general. Here we present a minimal PL_4P whose degree is double its camera-degree.

Consider the variety $\mathcal{X}_{9,5,\mathcal{I}}$ of point-line arrangements consisting of five free points (red in Fig. 3) and five lines (and four additional black points) where one of the lines (in black) intersects the other four (colorful) lines (in the four black extra points), i.e.

$$\mathcal{I} = \{(1, 1), (2, 2), (3, 3), (4, 4), (1, 5), (2, 5), (3, 5), (4, 5)\}.$$

We are interested in the PL_4P in two views where both views observe the five free points and the first four (colorful) lines, but not the fifth (black) line and none of the (black) extra points; see Fig. 3.

Clearly this point-line problem contains the classical five-point problem as a subproblem. In fact, it is reducible to the five-point problem by forgetting all five lines plus the four extra points. To see this we have to check the lifting property: for fixed camera poses, adding the first four lines to their views uniquely recovers these four lines in space. For generic four lines $L_1, \dots, L_4$ (shown in blue, green, orange and pink in Fig. 3) in 3D, there are two lines (the black one and the grey one) which intersect $L_1, \dots, L_4$ (Sottile, 2006). Hence, every solution to the five-point problem can be lifted to a solution of the PL_4P in Fig. 3 in two ways.

By Lemma 30, we see that the PL_4P is camera-minimal and has the same camera-degree as the five-point problem, namely 20. Moreover, the PL_4P is balanced since its reduction to the balanced five-point problem removes 16 degrees of freedom both in 3D and in the views. Hence, the PL_4P we constructed is indeed minimal. Since each of the 20 solutions of the five-points problem can be lifted to two solutions of the PL_4P , we have shown that its degree is 40.

Proof of Theorems 10 and 13 Theorem 13 immediately follows from parts 1 and 2 of Lemma 30. To complete the proof of Theorem 10, suppose the upper PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ is minimal. From parts 3 of lemmas 29 and 30, we have that Π and

π have equal fiber dimensions. It follows that the lower PL_1P $(p', l', \mathcal{I}', \mathcal{O}')$ is balanced, since

$$\begin{aligned} \dim(\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_m) &= \dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m) - \dim(\text{fiber}(\Pi)) \\ &= \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}) - \dim(\text{fiber}(\pi)) \\ &= \dim(\mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}). \end{aligned}$$

Moreover, from part 1 of Lemma 29 we have that Φ' is dominant; so $(p', l', \mathcal{I}', \mathcal{O}')$ is minimal.

For generic image data $(x', \ell') \in \mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}$ and $(x, \ell) \in \pi^{-1}(x', \ell')$, it remains to show that the finite sets $\Phi^{-1}(x, \ell)$ and $(\Phi')^{-1}(x', \ell')$ have the same cardinality. By Theorem 13, we already know this for the camera solutions $\gamma(\Phi^{-1}(x, \ell))$ and $\gamma((\Phi')^{-1}(x', \ell'))$. Since both problems are minimal, we are done by Lemma 7. $\square$

12.2 Theorems 11, 12, 14 and 20

In this subsection, we give full details on how the local features appearing in Theorems 11, 12, 14 and 20 are derived. The first Lemma 32 explains why unobserved features cannot occur in neither minimal nor reduced problems. The next two Lemmas 33 and 34 give reduction rules for pins.

Lemma 32 Consider a PL_1P in three views. If some point or line in space is not observed in any view, then the PL_1P is reducible and not minimal.


- For an unobserved line, the reduction forgets the line.
- For an unobserved point with at most one pin, the reduction forgets the point.
- For an unobserved point with at least two pins, the reduction forgets the point and all its pins.

Proof For each bullet we check that the associated forgetting map Π satisfies the conditions in the definition of reducibility. Feasibility holds vacuously for the first two bullets and is also easily seen for the third. Lifting solutions in the first two bullets is also trivial since Π only forgets what is not observed at all. This argument also shows that $\dim(\text{fiber}(\Pi)) > 0 = \dim(\text{fiber}(\pi))$, so by part 3 of Lemma 29 the PL_1P cannot be minimal in the first two bullets.


To verify the lifting property for the third bullet, we note that at most one of the pins can occur in each of the three views. By the first bullet, we may assume that all pins are visible. With these assumptions, it follows that the point has at most three pins. Hence, we are left with the following three cases:

Now we can easily check that the lifting property is satisfied in each of these three cases. For instance, for the last row (i.e. a point with three pins), the preimage of any three pins viewed like  under fixed cameras is three planes

Fig. 3 The five-point problem in two views can be combined with the classical Schubert four-line problem to obtain a minimal PL_4P with degree 40 and camera-degree 20



An arrangement of 5 lines and 9 points in 3D space. The grey line is *not* part of the arrangement.



A projection to two images that forgets all black features.

3-space	views	dim(fiber(Π))	dim(fiber(π))
point with 2 pins		7	4
point with 2 pins		7	6
point with 3 pins		9	6

in space which necessarily intersect in a point. Thus we can pick that point plus any three lines passing through that point and contained in the respective planes to lift solutions. This shows that the PL_1P is reducible as described in the third bullet. Moreover, we see that $\dim(\text{fiber}(\Pi)) > \dim(\text{fiber}(\pi))$ holds in all three cases depicted above, so the PL_1P cannot be minimal by part 3 of Lemma 29. $\square$

Lemma 33 Consider a PL_1P in three views. If some pin is observed in exactly one view such that the view also observes the point of the pin, then the PL_1P is not minimal and it is reducible by forgetting the pin.

Proof The forgetting map Π which forgets the pin is clearly feasible. It also satisfies the lifting property: for a fixed arrangement of cameras that view the point X of the pin, no matter how the pin is viewed in its single view, there is in fact a pencil of lines through X yielding that view. This shows that the PL_1P is reducible. Moreover, this reduction satisfies that

$$\dim(\text{fiber}(\Pi)) = 2 > 1 = \dim(\text{fiber}(\pi));$$

so by part 3 of Lemma 29 the PL_1P cannot be minimal. $\square$

Lemma 34 Consider a PL_1P in three views. If some pin is observed in exactly two views such that both views also observe the point of the pin, then the PL_1P is reducible by forgetting the pin.

Proof The forgetting map Π which forgets the pin is clearly feasible. It also satisfies the lifting property: for a fixed arrangement of cameras that view the point X of the pin, no matter how the pin is viewed in its two views, since it also passes through X in both views, there is a unique line in 3D through X yielding these two views. This shows that the PL_1P is reducible. $\square$

With the above lemmas in hand, we are able to enumerate a finite list of local features that may appear in a reduced (camera-)minimal PL_1P in three views as well as a finite list of reduction rules to obtain such a PL_1P . To aid in this task, we list all possible ways in which free lines and points with 0, 1 or 2 pins are viewed, and classify them according to whether or not a) they may appear in a minimal problem and b) they are reducible.

Lemma 35 below dispenses with cases involving local features with free lines and points with up to two pins that are already handled by Lemma 32, 33, or 34. The remaining cases are shown in Table 6. For each of these local features, we may forget either a point and/or some number of lines. Table 6 lists all feasible forgetting maps Π for each observed local feature. From this, we classify which observations of local features a) may appear in minimal problems and b) are reducible.

To determine reducibility of an observed local feature, we simply have to check if one of the listed feasible forgetting maps satisfies the lifting property. Finding out if a local feature can be observed in a certain way in a minimal problem, is more subtle. We use the following two rules to exclude observed local features from appearing in minimal problems:

Table 5 Local features observed in three views pertaining to Theorems 2, 3, 5, and 7 with examples

	can occur in a minimal problem?	
	YES	NO
reduced	Theorem 3 	Theorem 5
reducible	Theorem 2 	Theorem 7

- By Part 3 of Lemma 29, if $\dim(\text{fiber}(\Pi)) > \dim(\text{fiber}(\pi))$ for some feasible forgetting map Π , then the observed local feature cannot occur in any minimal PL_1P .
- A dangling pin (i.e. a pin in 3D which is observed in a single view) cannot occur in any minimal PL_1P .

All observations of local features which we cannot exclude from minimal problems using the two rules described above, we allow a priori to be part of minimal PL_1P s. We mark this in Table 6 with a “yes” in the column “minimal”⁶.

Finally, depending on the outcomes reducibility and minimality checks, we assign each observed local feature listed in Table 6 to one of the Theorems 11, 12, 14 and 20. The possible assignments are illustrated in Table 5.

Lemma 35 *All possibilities of how free lines and points with at most two pins are observed in three views are either treated by Lemmas 32, 33, 34 or appear in Table 6.*

Proof By Lemma 32, we only have to record the cases where each point and line is observed at least once. All such cases for free points and free lines are depicted in rows 1–6 of Table 6. So we are left to discuss points with one or two pins.

Points with one pin We are distinguishing the different cases by how often the pin and its point are observed in the three views. We use the short notation $\lambda : \rho$ to denote that the pin resp. its point is viewed λ resp. ρ times. By Lemma 32, we have that $\lambda, \rho \in \{1, 2, 3\}$.

- 1:1 By Lemma 33, we are left with the case where the pin and its point do not appear in the same view; see row 7 of Table 6.
- 1:2 By Lemma 33, we can exclude the cases where one view sees the pin together with its point. Hence, we get that one view sees the pin and the other two views observe its point; see row 8 of Table 6.
- 1:3 Here all three views observe the point and one of them also sees the pin. This is already handled by Lemma 33 and thus does not appear in Table 6.

- 2:1 Both such cases are shown in rows 9 and 10 of Table 6.
- 2:2 By Lemma 34, we may assume that the two views which see the pin do not both observe its point; see row 11 of Table 6.
- 2:3 Here all three views observe the point and two of them also see the pin. This is already handled by Lemma 34 and thus does not appear in Table 6.
- 3:1 —3:3 These cases are depicted in rows 12, 13 and 14 of Table 6.

Points with two pins We are distinguishing the different cases by how often each pin is observed in the three views. We use the short notation $\lambda_1 : \lambda_2$ to denote that the first resp. second pin is viewed λ_1 resp. λ_2 times. By Lemma 32, we have that $\lambda_1, \lambda_2 \in \{1, 2, 3\}$.

- 1:1 By Lemma 33, we can exclude the cases where a view that sees one of the pins also observes the point. So we get that each view observing one of the pins does neither see the point nor the other pin; see row 15 of Table 6.
- 1:2 By Lemma 33, we are left with the cases where the view observing the first pin does not see the point. So that view cannot see the other pin either. The other two views both observe the second pin, and at least one of them has to view the point. By Lemma 34, we may assume that the point is not observed by both of these views; see row 16 of Table 6.
- 1:3 Here all three views observe the second pin and one of them also sees the first pin, so also the point. This is already handled by Lemma 33 and does not appear in Table 6.
- 2:2 By Lemma 34, we can exclude the cases where two views observe both pins. Hence, we are left with the situation where one view sees both pins (and their point), and the other two views see one pin each. Again by Lemma 34, we may assume that none of the latter two views observes the point; see row 17 of Table 6.
- 2:3 There are exactly two views where both pins, and hence the point are seen. We apply Lemma 34 to see that this does not appear in Table 6.
- 3:3 See row 18 of Table 6.

□

Next, we prove Lemmas 36, 38, and 39, which address the cases involving three or more pins. Along the way, we prove Lemma 37, which will also be useful in establishing Theorem 16.

Lemma 36 *Consider a PL_1P in three views. If a point with at least three pins is not completely observed in the views (i.e. at least one view does not see at least one pin), then the PL_1P is reducible by one of the cases in Lemmas 32, 33, 34.*

⁶ Our computations described in Sect. 8 verify that actually each of the observed features marked as minimal in Table 6 does appear in some minimal PL_1P .


Proof If the point X or one of its pins is not observed in any view, then we are in the setting of Lemma 32. Hence, we assume that the point X and each of its pins is viewed at least once. Let us first assume that the point X is viewed by exactly one camera. The other two views can each see at most one of its pins. Since X has at least three pins in 3D, at least one of its pins must be only viewed by the same camera which sees the point X . This situation is handled by Lemma 33.

It is left to consider the cases when X is viewed by at least two cameras. Let us assume that the point X is observed in exactly two views. As above, at least one of its pins is not observed by the view not seeing X . So this pin has to be observed in either one or both of the views which observe the point X . This shows that the PL_1P is reducible by either Lemmas 33 or 34.

Finally, we assume that the point X is observed in all three views. Since the point X and its pins are (by our assumption in Lemma 36) not completely observed, one of its pins is seen in either one or two views. This is reducible by either Lemmas 33 or 34. $\square$

Lemma 38 shows that a point with 8 pins cannot occur in any reduced camera-minimal problem; for minimal problems in complete visibility, this is already a result in Duff et al. (2019). In the generality of camera-minimal problems, this result does not follow immediately from Lemma 36. However, the next lemma lets us get around this, and also proves part of Theorem 16.

Lemma 37 Applying any of the following replacements in images of a PL_1P in three views preserves camera-minimality, camera-degrees, and reducedness:



Proof The three replacements have in common that they fixate a dangling pin (when read from left to right) or create a dangling pin (when read from right to left). We prove the assertions in Lemma 37 for all three replacement rules at once. For this, we consider a PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ in three views that views one of its local features as depicted on the left of any of the three replacement rules. After applying that replacement rule (from left to right) once, we obtain a new PL_1P $(p', l', \mathcal{I}', \mathcal{O}')$. Clearly, every solution of $(p', l', \mathcal{I}', \mathcal{O}')$ is also a solution of $(p, l, \mathcal{I}, \mathcal{O})$. Moreover, for every solution (X, P) of $(p, l, \mathcal{I}, \mathcal{O})$, there is in fact a one-dimensional set $\{(\tilde{X}, P)\}$ of solutions of $(p, l, \mathcal{I}, \mathcal{O})$ with the same camera poses P , where the 3D arrangements $\tilde{X}$ differ from the fixed 3D arrangement X exactly by the dangling pin involved in the replacement rule. Hence, one of these solutions $(\tilde{X}, P)$ is also a solution of $(p', l', \mathcal{I}', \mathcal{O}')$. This shows

$$\text{cdeg}(p, l, \mathcal{I}, \mathcal{O}) = \text{cdeg}(p', l', \mathcal{I}', \mathcal{O}');$$

in particular, $(p, l, \mathcal{I}, \mathcal{O})$ is camera-minimal if and only if $(p', l', \mathcal{I}', \mathcal{O}')$ is camera-minimal.

Furthermore, local features, which are observed in one of the five different ways present in the three replacement rules, are not reducible; see also rows 8, 11, 15, 16, and 17 in Table 6. This means that the reducibility / reducedness of a point-line problem in three views does not depend on the appearance of local features observed as in the three replacement rules. More precisely, $(p, l, \mathcal{I}, \mathcal{O})$ is reduced if and only if $(p', l', \mathcal{I}', \mathcal{O}')$ is reduced. $\square$

Lemma 38 A reduced camera-minimal PL_1P in three views cannot have a point in 3D with eight pins.

Proof We assume by contradiction that there is a reduced camera-minimal PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ in three views which has a point with eight pins in its 3D arrangement. By Lemma 36, this point and all its pins are observed in all views. Since $(p, l, \mathcal{I}, \mathcal{O})$ is camera-minimal, its joint camera map is dominant. In particular, we have that $\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_3) \geq \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}})$, i.e.

$$11 \geq \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}) - \dim(\mathcal{X}_{p,l,\mathcal{I}}). \quad (3)$$

If the 3D arrangement of $(p, l, \mathcal{I}, \mathcal{O})$ consists only of the point with its eight pins, then (3) is actually an equality, so $(p, l, \mathcal{I}, \mathcal{O})$ is a minimal PL_1P completely observed by three calibrated views. However, in Duff et al. (2019) it is shown that this point-line problem is *not* minimal, a contradiction.

Hence, we see that the PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ has to contain at least one other local feature. By Lemma 35, the only possible local features are either points with at least three pins or the local features listed in Table 6 which are not reducible (rows 3,5,6,8,9,11,13–18). Since points with $k \geq 3$ pins have to be completely observed by Lemma 36, they have more degrees of freedom in the 2D images ($= 3(2+k)$) than in the 3D arrangement ($= 3+2k$). Similarly, all non-reducible features in Table 6, except rows 15 and 16, have more degrees of freedom in the 2D images than in the 3D arrangement. Since the inequality (3) has to hold for the PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ and the point with eight pins already makes this inequality tight, we have shown the following:

(★) If the 3D arrangement of a reduced camera-minimal PL_1P in three views contains a point with eight pins and at least one additional local feature, then it contains a point

with two pins which is either viewed like

(row 15 in Table 6) or (row 16 in Table 6).

In particular, the PL_1P $(p, l, \mathcal{I}, \mathcal{O})$ has to contain a point with two pins viewed like or . We apply the replacements in Lemma 37 to obtain a new reduced

camera-minimal PL_1P in three views containing a point with eight pins and a point with two pins viewed like , such that none of its points with two pins is observed like  nor . This contradicts $(\star)$. $\square$

Finally, we combine everything we have learned so far about pins to bound the maximum number of pins per point and the maximum number of points with many pins in reduced camera-minimal PL_1P s in three views. Afterwards we are ready to summarize our findings to provide proofs for Theorems 2, 3, 5, and 7.

Lemma 39 *A reduced camera-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.*

Proof By Lemma 36, every point with three or more pins has to be completely observed in a reduced PL_1P in three views. Hence, it is left to show that 1) at most one such point with many pins exists, and that 2) it has at most seven pins.

For the first assertion, we denote by ρ the number of points in 3D which have three or more pins. We consider the forgetting map Π which forgets everything, except these ρ points with exactly three of their pins each. We obtain a diagram as in Fig. 2. Since this forgetting map Π is feasible and the given (upper) PL_1P is camera-minimal, Lemma 29 implies that the joint camera map of the resulting (lower) PL_1P is dominant. In particular, the resulting PL_1P satisfies

$$9\rho + 11 = \dim(\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_3) \geq \dim(\mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}') = 15\rho,$$

i.e. $11 \geq 6\rho$. Thus, we see that $\rho \leq 1$, which means that the given (upper) PL_1P has at most one point in 3D with three or more pins.

Finally, we show that such a point has at most seven pins, if it exists. We assume $\rho = 1$ and denote by $\lambda \geq 3$ the number of pins at that point. We consider the forgetting map Π which forgets everything, except that single point with its λ pins. As before, we see that the joint camera map of the resulting lower PL_1P is dominant, which yields that

$$3 + 2\lambda + 11 = \dim(\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_3) \geq \dim(\mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}') = 3(2 + \lambda),$$

so $8 \geq \lambda$. By Lemma 38, we have that $\lambda \leq 7$, which concludes the proof. $\square$

Proof of Theorems 11 and 20 We first show that each PL_1P in three views is reducible to a unique reduced PL_1P . For

this, we notice that a PL_1P is reducible if and only if one of the local features in its 3D arrangement is reducible. Hence, we only have to check that each possible local feature is reducible to a unique reduced local feature. By Lemmas 35 and 36, this assertion follows from examining the reducible cases in Table 6 together with Lemmas 32 to 34.

Moreover, we obtain all reduction rules listed in Theorem 20 by collecting the forgetting maps described in Lemmas 32, to 34, as well as the reducible observed features in Table 6. Among those, Lemma 34 as well as rows 2, 10, and 12 of Table 6 are applicable for minimal problems; thus these reduction rules are listed in Theorem 11.

Finally, we address the last assertion in Theorem 11. Considering a commutative diagram as in Fig. 2 which is obtained from one of the four reduction rules listed in Theorem 11, we assume that the lower PL_1P is minimal and aim to prove that the upper PL_1P is minimal as well. Since the lower PL_1P is minimal, it is balanced and its joint camera map is dominant. By Lemma 30, the joint camera map of the upper PL_1P is also dominant. Furthermore, the four forgetting maps Π listed in Theorem 11 satisfy that $\dim(\text{fiber}(\Pi)) = \dim(\text{fiber}(\pi))$. Since the lower PL_1P is balanced, we see from Lemma 28 that

$$\begin{aligned} \dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_3) &= \dim(\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_3) + \dim(\text{fiber}(\Pi)) \\ &= \dim(\mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}') + \dim(\text{fiber}(\pi)) \\ &= \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}). \end{aligned}$$

So the upper PL_1P is also balanced, hence minimal. $\square$



Proof of Theorems 12 and 14 The parts of Theorem 12 (and Theorem 14) addressing points with three or more pins have already been proven in Lemma 39. To find all ways of how free lines and points with at most two pins can be observed by reduced camera-minimal PL_1P s in three views, it is enough (by Lemma 35) to gather the non-reducible observed features in Table 6. Among those, the ones marked as minimal are listed in Theorem 12, the others in Theorem 14. $\square$


12.3 Theorem 16 and its corollaries


Proof of Theorem 16 The nature of the four replacements is to introduce another image of the dangling pin into one of the views. This makes this pin reconstructable in 3D in the resulting PL_1P , but produces no additional constraint on the cameras.

Formally, since the initial problem is reduced and camera-minimal, Lemma 37 tells us that the resulting PL_1P is also reduced and camera-minimal, with the same camera-degree. It is left to show that the resulting PL_1P is indeed minimal. By Theorem 14, we see that the resulting PL_1P only has local features as listed in Theorem 12, i.e. points with three or more pins completely observed in all views plus some of

the features shown in Table 1. All of these features can be uniquely recovered in 3D for a fixed camera solution. $\square$

Proof of Corollary 18 By Lemma 37, replacing a single occurrence of  in a minimal problem with , yields a camera-minimal problem with the same camera-degree. Now we can replace  with  to obtain another camera-minimal PL₁P of the same camera-degree. As argued above, this resulting PL₁P is actually minimal. By Lemma 7, it has the same degree as the initial minimal problem.

Proof of Corollary 19 Starting from a terminal camera-minimal PL₁P in three views, the lift in Theorem 16 produces possibly several reduced minimal PL₁Ps: these only differ depending on whether the third or the fourth replacement in Theorem 16 is applied to each point with one dangling pin viewed like  (up to relabeling the views). So all of the resulting minimal problems lie in the same swap&label-equivalence class.

Starting from a reduced minimal PL₁P in three views, we apply the replacements in Lemma 37, read from right to left, until , , and  do not appear any longer. This, by Definition 17, results in a terminal problem. Moreover, all reduced minimal PL₁Ps which are related via the swaps in Corollary 18 yield the same resulting terminal problem.





This explains the one-to-one correspondence in Corollary 19. It preserves camera-degrees by Lemma 37. $\square$

13 Computations

13.1 Swap&Label-Equivalence Classes of Signatures

As noted in Sect. 6, each reduced minimal PL₁P in three views can be encoded as a *signature*, which is an integer solution to the dimension-count equation (2). Thus, a signature is simply an integer vector of length 27. In general, we can enumerate all nonnegative integer solutions $(k_1, \dots, k_l)$ to $a_1 k_1 + \dots + a_l k_l = n$ by recursively solving $a_2 k_2 + \dots + a_l k_l = n - j a_1$ for $j = \lfloor n/a_1 \rfloor, \dots, 0$. The result of this enumeration procedure is a list of solutions that is sorted decreasingly with respect to the usual lexicographic order on $\mathbb{Z}^{27}$. This is how we computed all 845161 solutions of (2).

For each signature, we find all PL₁Ps which are the same up to relabeling of the views—that is, we mod out the action of the permutation group S_3 that permutes the three views. Note that if we take one representative of each orbit of this action in the order they appear in the solution list from above, this ensures that the lex order from this step is preserved.

The procedure above gives us 143494 label-equivalence classes of reduced PL₁Ps. It remains to extract a representative of each swap&label-equivalence class from this list. Recall that a *swap* operation exchanges ordered local features of the form  and , and that two PL₁Ps are *swap&label-equivalent* if they differ only by some sequence of swaps and S_3 -permutations of the views. In our implementation, the coordinates of the signature vectors which participate in swaps are indexed from 12 to 17 inclusive. They are arranged such that 12 resp. 13 index the ordered local features  resp. —similarly for 14, 15 and 16, 17. Thus, we may restrict attention to those signatures whose coordinates 13, 15, and 17 equal zero; if any of these coordinates is nonzero, then we may find a swap&label-equivalent representative earlier in the list. To see this, note that if we swap 13 to 12, 15 to 14, and 17 to 16, we get a lexicographically larger signature; moreover, the signature is also lex-maximal in its label-equivalence class, and hence occurs in the list of 143494 representatives.

The previous paragraph identifies $\approx 40,000$ swap&label-equivalent pairs, but does not yet yield a unique representative for each class. To do this, we iterate over the list of remaining signatures in order, maintaining a single representative per swap&label-equivalence class encountered so far. For each signature, we enumerate its S_3 -orbit and perform the swaps 13 to 12, 15 to 14, and 17 to 16 to each element in this orbit. This operation produces 5 additional signatures, and we must delete any that appear later in the list. In the end, we are left with 76446 signatures.

13.2 Checking Minimality

As mentioned in Sect. 8, our rank check over the finite field $\mathbb{F}_q$ may be susceptible to *false negatives*—in other words, it is possible that we may incorrectly conclude that a problem is not minimal due to unlucky random choices made during the computation. On the other hand, *false positives* are impossible—we now explain this in detail. Let $(p, l, \mathcal{I}, \mathcal{O})$ be any balanced point-line problem. If we parametrize cameras by a rational map $P : \mathbb{C}^{11} \rightarrow \mathcal{C}_3$, then the exceptional set of complex (X, t) such that the Jacobian of $\Phi_{p,l,\mathcal{I},\mathcal{O}}$ at $(X, P(t))$ drops rank is Zariski-closed. It follows that there is a point (X_0, t_0) with integer coordinates outside of this exceptional set. Passing to residues modulo some prime q , the rank of the Jacobian at this point can only drop. Thus if we find a point with a Jacobian of full rank modulo q , then we may conclude that the joint camera map $\Phi_{p,l,\mathcal{I},\mathcal{O}}$ is dominant, i.e. that the balanced $(p, l, \mathcal{I}, \mathcal{O})$ is minimal.

Table 6 All ways to observe free lines and points with ≤ 2 pins in 3 views (modulo cases treated by Lemma 32, 33, 34) and their appearance in Theorems 2, 3, 5 or 7. For each observed feature, all possible feasible forgetting maps Π are listed, distinguished by the column “forget”: “P” (respectively “L”) denotes a forgotten point (respectively line.)

3D feature		2D views	Forget	Dim. of fiber(Π)	Dim. of fiber(π)	Lifting property	Reducible	Minimal	Thm
free line	1)		L	4	2	yes	yes	no	20
	2)		L	4	4	yes	yes	yes	11
	3)		L	4	6	no	no	yes	12
point	4)		P	3	2	yes	yes	no	20
	5)		P	3	4	no	no	yes	12
	6)		P	3	6	no	no	yes	12
point + 1 pin	7)		L	2	2	no			
			P	1	2	No	Yes	No	20
			P+L	5	4	Yes			
	8)		L	2	2	No			
			P	1	4	No	No	No	14
			P+L	5	6	No			
	9)		L	2	4	No			
			P	1	2	No	No	Yes	12
			P+L	5	6	No			
	10)		L	2	3	No			
			P	1	1	Yes	Yes	Yes	11
			P+L	5	5	Yes			
	11)		L	2	3	No			
			P	1	3	No	No	Yes	12
			P+L	5	7	No			
	12)		L	2	5	No			
			P	1	1	Yes	Yes	Yes	11
			P+L	5	7	No			
	13)		L	2	4	No			
			P	1	2	No	No	Yes	12
			P+L	5	8	No			
	14)		L	2	3	No			
			P	1	3	No	No	Yes	12
			P+L	5	9	No			
point + 2pins	15)		L	2	2	no			
			L+L	4	4	No	No	No	14
			P+L	3	4	No			
			P+L+L	7	6	No			
	16)		L	2	2	No			
			L	2	3	No			
			L+L	4	5	No	No	No	14
			P+L	3	3	No			
			P+L	3	5				No
			P+L+L	7	7	No			
	17)		L	2	3	No			
			L+L	4	6	No	No	Yes	12
			P+L	3	4	No			
			P+L+L	7	8	No			
	18)		L	2	3	No			
			L+L	4	6	No	No	Yes	12
			P+L	3	6	No			
			P+L+L	7	12	No			

13.3 Computing Degrees

The main results of Sect. 9, namely Results 22, 24 and 25, are based on our computation of degrees for minimal and camera-minimal problems using monodromy. In this context, the term “monodromy” refers to a paradigm of numerical continuation methods for solving parametrized polynomial systems that collect solutions for random parameter values in a one-by-one manner. We refer to Duff et al. (2018) for a detailed description of the ideas involved and discussion of implementation issues. Most relevant for our purposes is that this computation can be aborted early; ignoring numerical subtleties, we can then say that the number of solutions obtained at this point is a lower bound on the degree.

Our implementation of degree computation has the option of using one of two formulations. In the first, we solve explicitly for world features as well as camera matrices. In the second, *eliminated formulation*, we only solve for camera matrices using determinantal constraints as in Duff et al., (2019, Sec 6). Our computational results use the eliminated formulation for efficiency. As a sanity check for some problems of interest, we ran monodromy until it stabilized for both world and eliminated formulations to confirm the conclusion of Lemma 1.

Since our degree computations are randomized and susceptible to the possible failures of numerical continuation methods, our operational definition of “success” requires that all algebraic constraints are satisfied by the solutions collected and that either some target degree was exceeded (eg. 300 in Result 22) or that monodromy has *stabilized* in the sense that the random problem instances have collected the same number of solutions with no further progress after several iterations. We note that more sophisticated stopping criteria are available for monodromy (Hauenstein & Rodriguez, 2019; Leykin et al., 2018) but are generally much more expensive. For problems of interest (eg. those appearing in Kileel (2017); Duff et al. (2019) and Table 3), the stabilization heuristic allowed us to recover all previously known degrees, though some problems needed to be run more than once due to numerical failures. We also reran several cases appearing in Table 2 according to this criteria in order to gain more confidence in the reported degree.

13.4 Finding Subfamilies

As noted in Sect. 9, there are several subfamilies of minimal and camera-minimal PL_1 Ps that are of interest: namely, the PL_0 Ps occurring in Result 25, the PL_1 Ps with at most one pin per point in Result 24, and the extensions of the five-point problem from Result 26. We point out that the signature vectors described in Sect. 13.1 give a complete combinatorial description of each problem; in particular, in which views certain points and lines are seen and which incidences they

have. Thus, it is straightforward to enumerate these subfamilies starting from the lists of reduced minimal and terminal camera-minimal problems.

14 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 3 and 2 in this article and Table 4 in Sect. 11.)

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.

Acknowledgements We thank ICERM (NSF DMS-1439786 and the Simons Foundation grant 507536). We acknowledge: T. Duff and A. Leykin - NSF DMS-1719968 and DMS-2001267, 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, and EU H2020 SPRING No. 871245 projects at the Czech Institute of Informatics, Robotics and Cybernetics of the Czech Technical University in Prague.

Data Availability Our code is available at <https://github.com/timduff/35/PL1P>.

References

- Agarwal, S., Lee, H., Sturm, B., & Thomas, R. R. (2017). On the existence of epipolar matrices. *International Journal of Computer Vision*, 121(3), 403–415. <https://doi.org/10.1007/s11263-016-0949-7>
- Aholt, C., & Oeding, L. (2014). The ideal of the trifocal variety. *Mathematics of computation*, 83(289), 2553–2574.
- Aholt, C., Sturm, B., & Thomas, R. (2013). A Hilbert scheme in computer vision. *Canadian Journal of Mathematics*, 65(5), 961–988.
- Alismail, H. S., Browning, B., & Dias, M. B. (2011). Evaluating pose estimation methods for stereo visual odometry on robots. In *The 11th International Conference on Intelligent Autonomous Systems (IAS-11)*.
- Barath, D. (2018). 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.

- Barath, D., Toth, T., & Hajder, L. (2017). 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.
- Barath, D., & Hajder, L. (2018). Efficient recovery of essential matrix from two affine correspondences. *IEEE Transactions on Image Processing*, 27(11), 5328–5337.
- Bhayani, S., Kukulova, Z., & Heikkilä, J. (2020). A sparse resultant based method for efficient minimal solvers. In *Computer Vision and Pattern Recognition (CVPR)*.
- Byröd, M., Josephson, K., & Åström, K. (2008). 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.
- Camposco, F., Sattler, T., & Pollefeys, M. (2016). Minimal solvers for generalized pose and scale estimation from two rays and one point. In *ECCV: European Conference on Computer Vision*, pp. 202–218.
- Chen, H. H. (1990). Pose determination from line-to-plane correspondences: existence condition and closed-form solutions. In *ICCV*, pp. 374–378.
- Dhome, M., Richetin, M., Lapreste, J., & Rives, G. (1989). Determination of the attitude of 3D objects from a single perspective view. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(12), 1265–1278.
- Duff, T., Kohn, K., Leykin, A., & Pajdla, T. (2019). PLMP: Point-line minimal problems in complete multi-view visibility. In *International Conference on Computer Vision (ICCV)*.
- Duff, T., Hill, C., Jensen, A., Lee, K., Leykin, A., & Sommars, J. (2018). Solving polynomial systems via homotopy continuation and monodromy. *IMA Journal of Numerical Analysis*, 39(3), 1421–1446.
- Elqursh, A., & Elgammal, A. M. (2011). Line-based relative pose estimation. In *Cvpr*.
- Fabbri, R., Giblin, P. J., & Kimia, B. B. (2012). Camera pose estimation using first-order curve differential geometry. In *Proceedings of the European Conference in Computer Vision*.
- Fabbri, R., Duff, T., Fan, H., Regan, M., da de Costa Pinho, D., Tsigras, E., Wampler, C., Hauenstein, J., Giblin, P. J., & Kimia, B. B. (2023). Trifocal relative pose from lines at points. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6), 7870–7884.
- Fabbri, R., Giblin, P., & Kimia, B. (2020). Camera pose estimation using first-order curve differential geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10), 3321–3332.
- Fabbri, R., & Kimia, B. B. (2016). Multiview differential geometry of curves. *International Journal of Computer Vision*, 120(3), 324–346.
- Fischler, M. A., & Bolles, R. C. (1981). Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 381–395.
- Grayson, D. R., Stillman, M. E. Macaulay2, a software system for research in algebraic geometry. <http://www.math.uiuc.edu/Macaulay2/>
- Hartley, R. I. (1997). Lines and points in three views and the trifocal tensor. *International Journal of Computer Vision*, 22(2), 125–140.
- Hartley, R., & Li, H. (2012). An efficient hidden variable approach to minimal-case camera motion estimation. *IEEE PAMI*, 34(12), 2303–2314.
- Hartley, R., & Zisserman, A. (2003). *Multiple view geometry in computer vision*. Cambridge.
- Hauenstein, J. D., & Rodriguez, J. I. (2019). Multiprojective witness sets and a trace test. To appear in *Advances in Geometry*. arXiv preprint [arXiv:1507.07069](https://arxiv.org/abs/1507.07069)
- Hruby, P., Duff, T., Leykin, A., & Pajdla, T. (2022). Learning to solve hard minimal problems. In *Computer vision and pattern recognition (CVPR)*.
- Johansson, B., Oskarsson, M., & Åström, K. (2002). 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).
- Joswig, M., Kileel, J., Sturmfels, B., & Wagner, A. (2016). Rigid multi-view varieties. *IJAC*, 26(4), 775–788. <https://doi.org/10.1142/S021819671650034X>
- Kahl, F., Heyden, A., & Quan, L. (2001). Minimal projective reconstruction including missing data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(4), 418–424. <https://doi.org/10.1109/34.917578>
- Kileel, J. (2017). Minimal problems for the calibrated trifocal variety. *SIAM Journal on Applied Algebra and Geometry*, 1(1), 575–598.
- Kneip, L., Scaramuzza, D., & Siegwart, R. (2011). 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.
- Kneip, L., Siegwart, R., & Pollefeys, M. (2012). Finding the exact rotation between two images independently of the translation. In *ECCV: European Conference on Computer Vision*, pp. 696–709.
- Kuang, Y., & Åström, K. (2013a). Stratified sensor network self-calibration from TDOA measurements. In *21st European signal processing conference*.
- Kuang, Y., & Åström, K. (2013b). 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.
- Kukulova, Z., Bujnak, M., & Pajdla, T. (2008). Automatic generator of minimal problem solvers. In *European conference on computer vision (ECCV)*.
- Kukulova, Z., Kileel, J., Sturmfels, B., & Pajdla, T. (2017). A clever elimination strategy for efficient minimal solvers. In *Computer vision and pattern recognition (CVPR) IEEE*.
- Larsson, V., Åström, K., & Oskarsson, M. (2017). 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.
- Larsson, V., Åström, K., & Oskarsson, M. (2017a). Efficient solvers for minimal problems by syzygy-based reduction. In *Computer vision and pattern recognition (CVPR)*.
- Larsson, V., Åström, K., & Oskarsson, M. (2017b). Polynomial solvers for saturated ideals. In *IEEE international conference on computer vision, ICCV 2017*, Venice, Italy, October 22–29, 2017, pp. 2307–2316.
- Larsson, V., Kukulova, Z., & Zheng, Y. (2017c). Making minimal solvers for absolute pose estimation compact and robust. In *International conference on computer vision (ICCV)*.
- Larsson, V., et. al. (2015). Automatic generator of minimal problems. http://www2.maths.lth.se/matematiklth/personal/viktorl/code/basis_selection.zip
- Larsson, V., Oskarsson, M., Åström, K., Wallis, A., Kukulova, Z., & Pajdla, T. (2018). Beyond Gröbner 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. http://openaccess.thecvf.com/content_cvpr_2018/html/Larsson_Beyond_Grobner_Bases_CVPR_2018_paper.html
- Leykin, A., Rodriguez, J. I., & Sottile, F. (2018). Trace test. *Arnold Mathematical Journal*, 4(1), 113–125. <https://doi.org/10.1007/s40598-018-0084-3>

- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2), 91–110.
- Ma, Y., Huang, K., Vidal, R., Kosecka, J., & Sastry, S. (2004). Rank conditions on the multiple-view matrix. *International Journal of Computer Vision*, 59(2), 115–137.
- Matas, J., Obdržálek, S., & Chum, O. (2002). 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.
- Miraldo, P., Dias, T., & Ramalingam, S. (2018). 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.
- Miraldo, P., & Araujo, H. (2015). Direct solution to the minimal generalized pose. *Cybernetics, IEEE Transactions on*, 45(3), 418–429.
- Mirzaei, F. M., & Roumeliotis, S. I. (2011). Optimal estimation of vanishing points in a manhattan world. In *International Conference on Computer Vision (ICCV)*.
- Nistér, D., Naroditsky, O., & Bergen, J. (2004). Visual odometry. In *Computer vision and pattern recognition (CVPR)*, pp. 652–659.
- Nistér, D. (2004). An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6), 756–770.
- Nistér, D., & Schaffalitzky, F. (2006). Four points in two or three calibrated views: Theory and practice. *International Journal of Computer Vision*, 67(2), 211–231.
- Oskarsson, M., Åström, K., & Overgaard, N. C. (2001). Classifying and solving minimal structure and motion problems with missing data. In: *International Conference on Computer Vision (ICCV)*, 628–634. IEEE Computer Society. <https://doi.org/10.1109/ICCV.2001.10072>.
- Oskarsson, M., Zisserman, A., & Åström, K. (2004). Minimal projective reconstruction for combinations of points and lines in three views. *Image and Vision Computing*, 22(10), 777–785.
- Oxley, J. (2022). Matroid theory. In *Handbook of the Tutte polynomial and related topics*, pp. 44–85. Chapman and Hall/CRC, London.
- Raguram, R., Chum, O., Pollefeys, M., Matas, J., & Frahm, J. (2013). USAC: A universal framework for random sample consensus. *IEEE Transactions on Pattern Analysis Machine Intelligence*, 35(8), 2022–2038.
- Ramalingam, S., & Sturm, P. F. (2008). Minimal solutions for generic imaging models. In *CVPR: IEEE Conference on Computer Vision and Pattern Recognition*.
- Ramalingam, S., Bouaziz, S., & Sturm, P. (2011). Pose estimation using both points and lines for geo-localization. In *ICRA*, pp. 4716–4723.
- Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Pajdla, T., & Sivic, J. (2018). Neighbourhood consensus networks.
- Salaün, Y., Marlet, R., & Monasse, P. (2016). Robust and accurate line-and/or point-based pose estimation without manhattan assumptions. In *European Conference on Computer Vision (ECCV)*.
- Sattler, T., Leibe, B., & Kobbelt, L. (2017). Efficient & effective prioritized matching for large-scale image-based localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9), 1744–1756.
- Saurer, O., Pollefeys, M., & Lee, G. H. (2015). A minimal solution to the rolling shutter pose estimation problem. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ international conference on*, IEEE, pp. 1328–1334.
- Schönberger, J. L., & Frahm, J.-M. (2016). Structure-from-motion revisited. In *Conference on computer vision and pattern recognition (CVPR)*.
- Snavely, N., Seitz, S. M., & Szeliski, R. (2006). Photo tourism: Exploring photo collections in 3D. In *ACM SIGGRAPH*.
- Snavely, N., Seitz, S. M., & Szeliski, R. (2008). Modeling the world from internet photo collections. *International Journal of Computer Vision (IJCV)*, 80(2), 189–210.
- Sottile, F. (2001). Enumerative real algebraic geometry. In: *Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science*.
- Stewenius, H., Engels, C., & Nistér, D. (2006). Recent developments on direct relative orientation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 60, 284–294.
- Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A. (2018). InLoc: Indoor visual localization with dense matching and view synthesis. In *CVPR*.
- Trager, M. (2018). Cameras, shapes, and contours: Geometric models in computer vision. (caméras, formes et contours: modèles géométriques en vision par ordinateur). PhD thesis, École Normale Supérieure, Paris, France.
- Trager, M., Ponce, J., & Hebert, M. (2016). Trinocular geometry revisited. *International Journal Computer Vision*, 120, 134–152.
- Ventura, J., Arth, C., & Lepetit, V. (2015). An efficient minimal solution for multi-camera motion. In *International Conference on Computer Vision (ICCV)*, pp. 747–755.
- Xia, G., Delon, J., & Gousseau, Y. (2014). Accurate junction detection and characterization in natural images. *International Journal of Computer Vision*, 106(1), 31–56.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.