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

— Supplementary Material —

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

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

Here we give proofs of all results from the main paper

Duff, T., Kohn, K., Leykin, A., Pajdla, T.: PL_1P - Point-line minimal problems under partial visibility in three views. In: European Conference on Computer Vision (ECCV), 2020,

some additional results for camera registration, and other details. On the last page, we provide a glossary of assumptions, properties and concepts used through the whole article; see Section 14. Our code is available at `https://github.com/timduff35/PL1P`.

10 Note on partial visibility in two views

Let us recall that we consider reduced minimal and reduced camera-minimal PL_1Ps in *three views*, since there is only one such PL_1P in two views, namely the five-points problem.

To argue this, we first notice that free lines cannot occur in a reduced PL_1P 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_1P (i.e. forgetting the free line satisfies Definition 5 of reducibility).

Secondly, we argue that pins cannot appear in a reduced PL_1P 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_1P 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 if 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_1P 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_1P 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 $\mathrm{PL}_1\mathrm{P}$ 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 $\mathrm{PL}_1\mathrm{P}$ is reducible by forgetting the point together with all of its pins.

Finally, a $\mathrm{PL}_1\mathrm{P}$ 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 $\mathrm{PL}_1\mathrm{P}$s 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 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 4. 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 Theorems 1 and 4

We start by investigating the implications of the two conditions in Definition 5 of reducibility. When obtaining a new $\mathrm{PL}_1\mathrm{P}$ $(p', l', \mathcal{I}', \mathcal{O}')$ from a given $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ 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 Figure 1. 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 5 (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 3. *If the forgetting map Π is feasible, then both Π and π are surjective and have irreducible fibers*[13] *of equal dimension.*

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

[13] A fiber of a map is the preimage over a single point in its codomain.

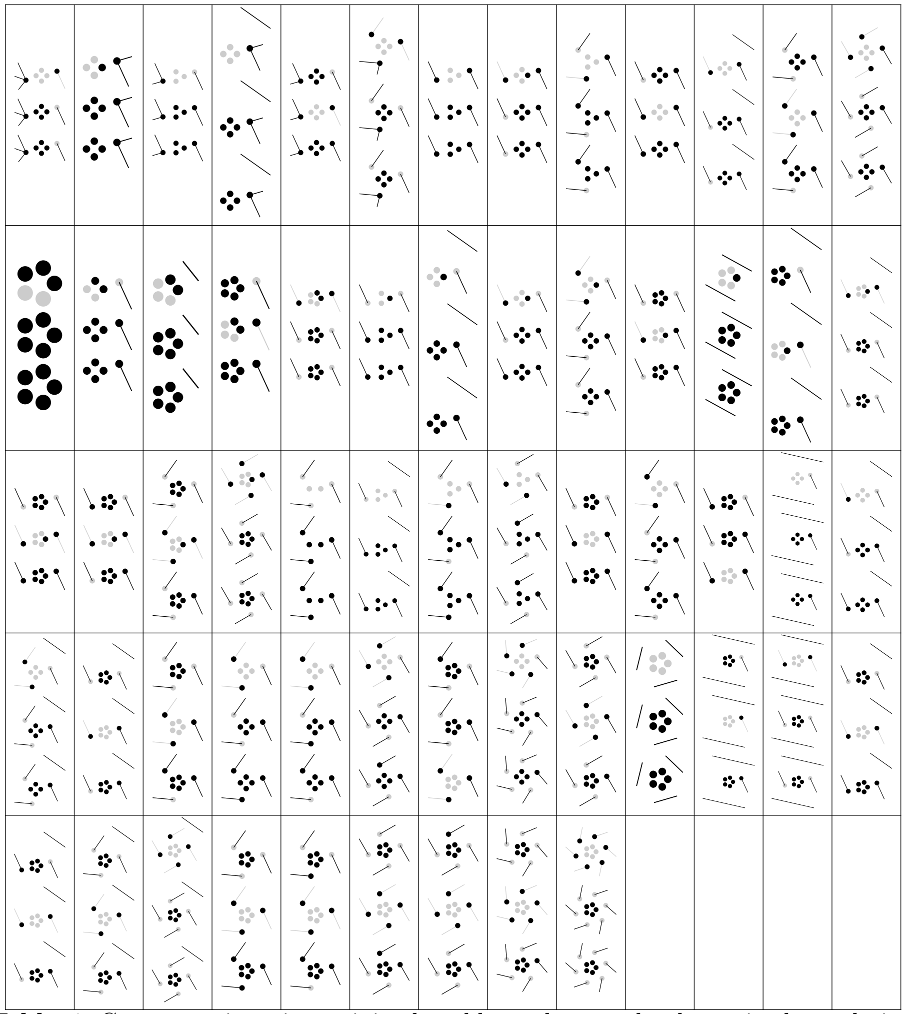

Table 4. Camera registration: minimal problems that can be determined as relative pose problems for two views appended with a minimal registration problem for the third camera.

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

$$\begin{tikzcd}
\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_m \arrow[r, dashed, "\Phi = \Phi_{p,l,\mathcal{I},\mathcal{O}}"] \arrow[d, "\Pi"'] & \mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}} \arrow[d, "\pi"] \\
\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_m \arrow[r, dashed, "\Phi' = \Phi_{p',l',\mathcal{I}',\mathcal{O}'}"'] & \mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}
\end{tikzcd}$$

Fig. 1. Two $\mathrm{PL_1P}$s and their joint camera maps related by a forgetting map Π. The upper $\mathrm{PL_1P}$ is reducible to the lower $\mathrm{PL_1P}$ iff the forgetting map Π is feasible and satisfies the lifting property.

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. □

For a feasible forgetting map Π, we write $\dim(\mathrm{fiber}(\Pi))$, resp. $\dim(\mathrm{fiber}(\pi))$, for the dimensions of the fibers of Π, resp. π. Moreover, we denote by $\mathrm{cdeg}(p,l,\mathcal{I},\mathcal{O})$ 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 4. *Consider Figure 1 with a feasible forgetting map Π.*

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

Proof. For a generic image $(x',\ell') \in \mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}$ of the lower $\mathrm{PL_1P}$, we consider a generic image $(x,\ell) \in \pi^{-1}(x',\ell')$ of the upper $\mathrm{PL_1P}$. Every solution $S \in \Phi^{-1}(x,\ell)$ of the upper $\mathrm{PL_1P}$ yields a solution $\Pi(S)$ of the lower $\mathrm{PL_1P}$ 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 $\mathrm{PL_1P}$ is dominant, we use part 1 to see that the lower joint camera map Φ' is also dominant. This implies:

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

Since each minimal $\mathrm{PL_1P}$ 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. □

Lemma 5. *Consider Figure 1 with a feasible forgetting map Π that satisfies the lifting property.*

1. *Φ is dominant $\Leftrightarrow$ Φ' is dominant.*
2. $\mathrm{cdeg}(p',l',\mathcal{I}',\mathcal{O}') = \mathrm{cdeg}(p,l,\mathcal{I},\mathcal{O})$.

3. $\dim(\mathrm{fiber}(\Pi)) \geq \dim(\mathrm{fiber}(\pi))$.

Proof. We can pick a generic image of the upper $\mathrm{PL}_1\mathrm{P}$, by first choosing a generic image $(x', \ell') \in \mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}$ of the lower $\mathrm{PL}_1\mathrm{P}$ 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 $\mathrm{PL}_1\mathrm{P}$ can be lifted to a solution $S \in \Phi^{-1}(x, \ell)$ of the upper $\mathrm{PL}_1\mathrm{P}$ with *the same* cameras. Together with Lemma 4, 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

$$\dim(\mathrm{fiber}(\Pi)) = \dim(\Pi^{-1}(S')) \geq \dim(\pi^{-1}(\Phi'(S'))) = \dim(\mathrm{fiber}(\pi)). \qquad \square$$

Proof of Lemma 1. Let $(p, l, \mathcal{I}, \mathcal{O})$ be a minimal $\mathrm{PL}_1\mathrm{P}$ with camera-degree d. Fix $(x, \ell) \in \mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}$ generic and let $\gamma\left(\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell)\right) = \{P_1, \ldots, P_d\}$. We may assume that $\Phi^{-1}_{p,l,\mathcal{I},\mathcal{O}}(x, \ell)$ is a nonempty finite set; we wish to show it contains d elements. We begin by lifting $P_1, \ldots, P_d$ to partial solutions in $\mathbb{G}^l_{1,3} \times \mathcal{C}_m$, 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 though 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 $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ 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 $(L_i, P_i) \in \mathbb{G}^l_{1,3} \times \mathcal{C}_m$. To conclude that there are exactly d lifts

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

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 $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ is minimal, the latter cannot hold. $\square$

Example 1. Lemma 1 does *not* hold for point-line problems in general. Here we present a minimal $\mathrm{PL}_4\mathrm{P}$ 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 Figure 2) 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 $\mathrm{PL}_4\mathrm{P}$ 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 Figure 2.

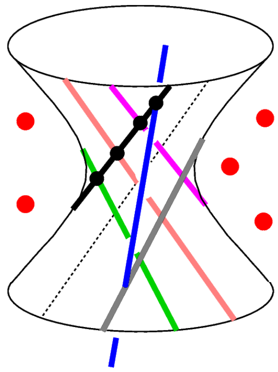

An arrangement of 5 lines and 9 points in 3D space. The grey line is *not* part of the arrangement.

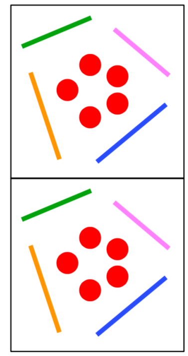

A projection to two images that forgets all black features.

Fig. 2. The five-point problem in two views can be combined with the classical Schubert four-line problem to obtain a minimal $\mathrm{PL}_4\mathrm{P}$ with degree 40 and camera-degree 20.

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 Figure 2) in 3D, there are two lines (the black one and the grey one) which intersect $L_1, \dots, L_4$.[14] Hence, every solution to the five-point problem can be lifted to a solution of the $\mathrm{PL}_4\mathrm{P}$ in Figure 2 *in two ways*.

By Lemma 5, we see that the $\mathrm{PL}_4\mathrm{P}$ is camera-minimal and has the same camera-degree as the five-point problem, namely 20. Moreover, the $\mathrm{PL}_4\mathrm{P}$ 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 $\mathrm{PL}_4\mathrm{P}$ we constructed is indeed minimal. Since each of the 20 solutions of the five-points problem can be lifted to *two* solutions of the $\mathrm{PL}_4\mathrm{P}$, we have shown that its degree is 40.

Proof of Theorems 1 and 4. Theorem 4 immediately follows from parts 1 and 2 of Lemma 5. To complete the proof of Theorem 1, suppose the upper $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ is minimal. From parts 3 of lemmas 4 and 5, we have that Π and π have equal fiber dimensions. It follows that the lower $\mathrm{PL}_1\mathrm{P}$ $(p', l', \mathcal{I}', \mathcal{O}')$ is balanced, since

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

[14] Sottile, F.: Enumerative real algebraic geometry. In: Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science. (2001)

Moreover, from part 1 of Lemma 4 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 4, we already know this for the camera solutions $\gamma\left(\Phi^{-1}(x, \ell)\right)$ and $\gamma\left((\Phi')^{-1}(x', \ell')\right)$. Since both problems are minimal, we are done by Lemma 1. □

12.2 Theorems 2, 3, 5, and 7

In this subsection, we give full details on how the local features appearing in Theorems 2, 3, 5, and 7 are derived. The first Lemma 6 explains why unobserved features cannot occur in neither minimal nor reduced problems. The next two Lemmas 7 and 8 give reduction rules for pins.

Lemma 6. *Consider a PL$_1$P in three views. If some point or line in space is not observed in any view, then the PL$_1$P 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 4 the PL$_1$P 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:

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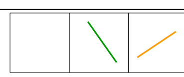	7	4
point with 2 pins	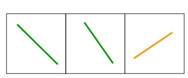	7	6
point with 3 pins	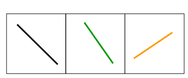	9	6

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 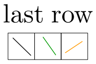 under fixed cameras is three planes in space which necessarily intersect in a point. Thus we can pick that point plus any three lines passing through that point and contained in the respective planes to lift solutions. This shows that the $\mathrm{PL_1P}$ is reducible as described in the third bullet. Moreover, we see that $\dim(\mathrm{fiber}(\Pi)) > \dim(\mathrm{fiber}(\pi))$ holds in all three cases depicted above, so the $\mathrm{PL_1P}$ cannot be minimal by part 3 of Lemma 4. □

Lemma 7. *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 $\mathrm{PL_1P}$ is reducible. Moreover, this reduction satisfies that

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

so by part 3 of Lemma 4 the $\mathrm{PL_1P}$ cannot be minimal. □

Lemma 8. *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. Is 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 $\mathrm{PL_1P}$ is reducible. □

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 $\mathrm{PL_1P}$ in three views as well as a finite list of reduction rules to obtain such a $\mathrm{PL_1P}$. 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 9 below dispenses with cases involving local features with free lines and points with up to two pins that are already handled by Lemma 6, 7, or 8. 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:

– By Part 3 of Lemma 4, if $\dim(\mathrm{fiber}(\Pi)) > \dim(\mathrm{fiber}(\pi))$ 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_1Ps. We mark this in Table 6 with a "yes" in the column "minimal"[15].

Finally, depending on the outcome of the reducibilty and minimality checks, we assign each observed local feature listed in Table 6 to one of the Theorems 2,3,5, or 7. This assignment is summarized in Table 5.

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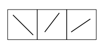	Theorem 5 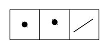
reducible	Theorem 2 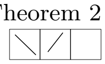	Theorem 7 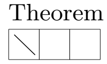

Table 5. Local features observed in three views pertaining to Theorems 2, 3, 5, and 7 with examples.

Lemma 9. *All possibilities of how free lines and points with at most two pins are observed in three views are either treated by Lemmas 6, 7, 8 or appear in Table 6.*

Proof. By Lemma 6, 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 $\boldsymbol{\lambda : \rho}$ to denote that the pin resp. its point is viewed λ resp. ρ times. By Lemma 6, we have that $\lambda, \rho \in \{1, 2, 3\}$.

1:1 By Lemma 7, 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 7, 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.

[15] Our computations described in Section 7 verify that actually all observed features marked as minimal in Table 6 do appear in some minimal PL_1Ps.

1:3 Here all three views observe the point and one of them also sees the pin. This is already handled by Lemma 7 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 8, 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 8 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 $\boldsymbol{\lambda_1 : \lambda_2}$ to denote that the first resp. second pin is viewed λ_1 resp. λ_2 times. By Lemma 6, we have that $\lambda_1, \lambda_2 \in \{1, 2, 3\}$.

1:1 By Lemma 7, 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 7, 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 8, 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 7 and does not appear in Table 6.

2:2 By Lemma 8, 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 8, 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 8 to see that this does not appear in Table 6.

3:3 See row 18 of Table 6. □

Next, we prove Lemmas 10, 12, and 13, which address the cases involving three or more pins. Along the way, we prove Lemma 11, which will also be useful in establishing Theorem 6.

Lemma 10. *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 6, 7, 8.*

Proof. If the point X or one of its pins is not observed in any view, then we are in the setting of Lemma 6. 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 7.

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$_1$P is reducible by either Lemma 7 or Lemma 8.

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 10) not completely observed, one of its pins is seen in either one or two views. This is reducible by either Lemma 7 or Lemma 8. $\square$

Lemma 12 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 [14]. In the generality of camera-minimal problems, this result does not follow immediately from Lemma 10. However, the next lemma lets us get around this, and also proves part of Theorem 6.

Lemma 11. *Applying any of the following replacements in images of a PL$_1$P in three views preserves camera-minimality, camera-degrees, and reducedness:*
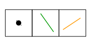 $\leftrightarrow$ 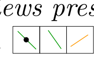, 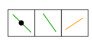 $\leftrightarrow$ 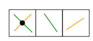, 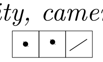 $\leftrightarrow$ 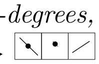

Proof. The three replacements have in common that they fixate a dangling pin (when read from left to right) or create a dangling pin (when read from right to left). We prove the assertions in Lemma 11 for all three replacement rules at once. For this, we consider a PL$_1$P $(p, l, \mathcal{I}, \mathcal{O})$ 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$_1$P $(p', l', \mathcal{I}', \mathcal{O}')$. 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

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

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 12. *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 $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ in three views which has a point with eight pins in its 3D arrangement. By Lemma 10, 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.

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

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 $\mathrm{PL}_1\mathrm{P}$ completely observed by three calibrated views. However, in [14] it is shown that this point-line problem is *not* minimal, a contradiction.

Hence, we see that the $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ has to contain at least one other local feature. By Lemma 9, 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 10, 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 $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ 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 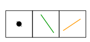 (row 15 in Table 6) or 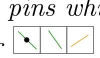 (row 16 in Table 6). $(\star)$

In particular, the $\mathrm{PL}_1\mathrm{P}$ $(p, l, \mathcal{I}, \mathcal{O})$ has to contain a point with two pins viewed like 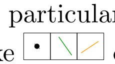 or 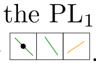. We apply the replacements in Lemma 11 to obtain a new reduced camera-minimal $\mathrm{PL}_1\mathrm{P}$ in three views containing a point with eight pins and a point with two pins viewed like 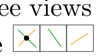, such that none of its points with two pins is observed like 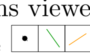 nor 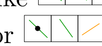. This contradicts $(\star)$. □

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 $\mathrm{PL}_1\mathrm{P}$s in three views. Afterwards we are ready to summarize our findings to provide proofs for Theorems 2, 3, 5, and 7.

Lemma 13. *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 10, 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 Figure 1. Since this forgetting map Π is feasible and the given (upper) PL_1P is camera-minimal, Lemma 4 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 12, we have that $\lambda \leq 7$, which concludes the proof. □

Proof of Theorems 2 and 7. 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 9 and 10, this assertion follows from examining the reducible cases in Table 6 together with Lemmas 6,7,8.

Moreover, we obtain all reduction rules listed in Theorem 7 by collecting the forgetting maps described in Lemmas 6,7,8, as well as the reducible observed features in Table 6. Among those, Lemma 8 as well as rows 2, 10, and 12 of Table 6 are applicable for minimal problems; thus these reduction rules are listed in Theorem 2.

Finally, we address the last assertion in Theorem 2. Considering a commutative diagram as in Figure 1 which is obtained from one of the four reduction rules listed in Theorem 2, 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 5, the joint camera map of the upper PL_1P is also dominant. Furthermore, the four forgetting maps Π listed in Theorem 2 satisfy that $\dim(\text{fiber}(\Pi)) = \dim(\text{fiber}(\pi))$. Since the lower PL_1P is balanced, we see from Lemma 3 that

$$\dim(\mathcal{X}_{p,l,\mathcal{I}} \times \mathcal{C}_3) = \dim(\mathcal{X}_{p',l',\mathcal{I}'} \times \mathcal{C}_3) + \dim(\text{fiber}(\Pi))$$

$$= \dim(\mathcal{Y}_{p',l',\mathcal{I}',\mathcal{O}'}) + \dim(\text{fiber}(\pi)) = \dim(\mathcal{Y}_{p,l,\mathcal{I},\mathcal{O}}).$$

So the upper $\mathrm{PL}_1\mathrm{P}$ is also balanced, hence minimal. □

Proof of Theorems 3 and 5. The parts of Theorem 3 (and Theorem 5) addressing points with three or more pins have already been proven in Lemma 13. To find all ways of how free lines and points with at most two pins can be observed by reduced camera-minimal $\mathrm{PL}_1\mathrm{P}$s in three views, it is enough (by Lemma 9) to gather the non-reducible observed features in Table 6. Among those, the ones marked as minimal are listed in Theorem 3, the others in Theorem 5. □

12.3 Theorem 6 and its corollaries

Proof of Theorem 6. 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 $\mathrm{PL}_1\mathrm{P}$, but produces no additional constraint on the cameras.

Formally, since the initial problem is reduced and camera-minimal, Lemma 11 tells us that the resulting $\mathrm{PL}_1\mathrm{P}$ is also reduced and camera-minimal, with the same camera-degree. It is left to show that the resulting $\mathrm{PL}_1\mathrm{P}$ is indeed minimal. By Theorem 5, we see that the resulting $\mathrm{PL}_1\mathrm{P}$ only has local features as listed in Theorem 3, 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. □

Proof of Corollary 1. By Lemma 11, replacing a single occurrence of 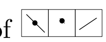 in a minimal problem with 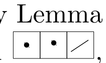, yields a camera-minimal problem with the same camera-degree. Now we can replace 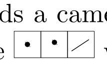 with 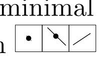 to obtain another camera-minimal $\mathrm{PL}_1\mathrm{P}$ of the same camera-degree. As argued above, this resulting $\mathrm{PL}_1\mathrm{P}$ is actually minimal. By Lemma 1, it has the same degree as the initial minimal problem. □

Proof of Corollary 2. Starting from a terminal camera-minimal $\mathrm{PL}_1\mathrm{P}$ in three views, the lift in Theorem 6 produces possibly several reduced minimal $\mathrm{PL}_1\mathrm{P}$s: these only differ depending on whether the third or the fourth replacement in Theorem 6 is applied to each point with one dangling pin viewed like 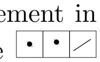 (up to relabeling the views). So all of the resulting minimal problems lie in the same swap&label-equivalence class.

Starting from a reduced minimal $\mathrm{PL}_1\mathrm{P}$ in three views, we apply the replacements in Lemma 11, read from right to left, until 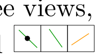, 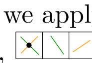, and 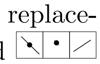 do not appear any longer. This, by Definition 6, results in a terminal problem. Moreover, all reduced minimal $\mathrm{PL}_1\mathrm{P}$s which are related via the swaps in Corollary 1 yield the same resulting terminal problem.

This explains the one-to-one correspondence in Corollary 2. It preserves camera-degrees by Lemma 11. □

13 Computations

13.1 Swap&label-equivalence classes of signatures

As noted in Section 5, each reduced minimal PL$_1$P 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 $(k_1, \ldots, k_l)$ to $a_1k_1 + \cdots + a_lk_l = n$ by recursively solving $a_2k_2 + \ldots + a_lk_l = n - ja_1$ for $j = \lfloor n/a_1 \rfloor, \ldots, 0$. 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 PL$_1$Ps 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 PL$_1$Ps. It remains to extract a representative of each swap&label-equivalence class from this list. Recall that a *swap* operation exchanges orderered local features of the form 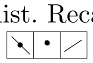 and 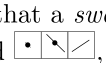, and that two PL$_1$Ps are *swap&label-equivalent* if they differ only by some sequence of swaps and S_3-permutations of the views. In our implementation, the coordinates of the signature vectors which participate in swaps are indexed from 12 to 17 inclusive. They are arranged such that 12 resp. 13 index the ordered local features 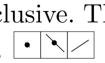 resp. 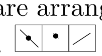 —similarly for $14, 15$ and $16, 17$. Thus, we may restrict attention to those signatures whose coordinates $13, 15$, and 17 equal zero; if any of these coordinates is nonzero, then we may find a swap&label-equivalent representative earlier in the list. To see this, note that if we swap 13 to 12, 15 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 Section 7, 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 $P : \mathbb{C}^{11} \to \mathcal{C}_3$, 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 av 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.

13.3 Computing degrees

The main results of Section 8, namely Results 9, 10 and 11, 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 [13] 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 [14, 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 9) 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 [41,24] but are generally much more expensive. For problems of interest (eg. those appearing in [28,14] 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 Section 8, there are several subfamilies of minimal and camera-minimal PL_1Ps that are of interest: namely, the PL_0Ps occuring in Result 11, the PL_1Ps with at most one pin per point in Result 10, and the extensions of the five-point problem from Result 12. We point out that the signature vectors described in Section 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.

3D feature		2D views	forget	dim. of fiber(Π)	dim. of fiber(π)	lifting property	reducible	minimal	Thm.
free line	1)	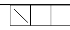	L	4	2	yes	yes	no	7
	2)	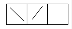	L	4	4	yes	yes	yes	2
	3)	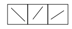	L	4	6	no	no	yes	3
point	4)	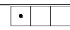	P	3	2	yes	yes	no	7
	5)	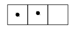	P	3	4	no	no	yes	3
	6)	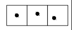	P	3	6	no	no	yes	3
point+ 1 pin	7)	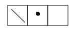	L	2	2	no	yes	no	7
			P	1	2	no			
			P+L	5	4	yes			
	8)	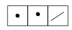	L	2	2	no	no	no	5
			P	1	4	no			
			P+L	5	6	no			
	9)	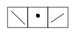	L	2	4	no	no	yes	3
			P	1	2	no			
			P+L	5	6	no			
	10)	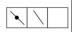	L	2	3	no	yes	yes	2
			P	1	1	yes			
			P+L	5	5	yes			
	11)	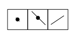	L	2	3	no	no	yes	3
			P	1	3	no			
			P+L	5	7	no			
	12)	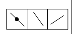	L	2	5	no	yes	yes	2
			P	1	1	yes			
			P+L	5	7	no			
	13)	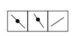	L	2	4	no	no	yes	3
			P	1	2	no			
			P+L	5	8	no			
	14)	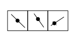	L	2	3	no	no	yes	3
			P	1	3	no			
			P+L	5	9	no			
point+ 2 pins	15)		L	2	2	no	no	no	5
			L+L	4	4	no			
			P+L	3	4	no			
			P+L+L	7	6	no			
	16)		L	2	2	no	no	no	5
			L	2	3	no			
			L+L	4	5	no			
			P+L	3	3	no			
			P+L	3	5	no			
			P+L+L	7	7	no			
	17)	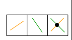	L	2	3	no	no	yes	3
			L+L	4	6	no			
			P+L	3	4	no			
			P+L+L	7	8	no			
	18)	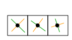	L	2	3	no	no	yes	3
			L+L	4	6	no			
			P+L	3	6	no			
			P+L+L	7	12	no			

Table 6. All ways to observe free lines and points with ≤ 2 pins in 3 views (modulo cases treated by Lemmas 6,7,8) 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" resp. "L" denotes a forgotten point resp. line.

14 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

References

1. Agarwal, S., Lee, H., Sturmfels, B., Thomas, R.R.: On the existence of epipolar matrices. International Journal of Computer Vision **121**(3), 403–415 (2017). https://doi.org/10.1007/s11263-016-0949-7, `https://doi.org/10.1007/s11263-016-0949-7`
2. Aholt, C., Oeding, L.: The ideal of the trifocal variety. Math. Comput. **83**(289), 2553–2574 (2014)
3. Aholt, C., Sturmfels, B., Thomas, R.: A hilbert scheme in computer vision. Canadian Journal of Mathematics **65**(5), 961–988 (2013)

4. Fabri et al, R.: Trifocal relative pose from lines at points and its efficient solution. Preprint arXiv:1903.09755 (2019)
5. Alismail, H.S., Browning, B., Dias, M.B.: Evaluating pose estimation methods for stereo visual odometry on robots. In: the 11th International Conference on Intelligent Autonomous Systems (IAS-11) (January 2011)
6. Barath, D.: Five-point fundamental matrix estimation for uncalibrated cameras. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 235–243 (2018)
7. Barath, D., Hajder, L.: Efficient recovery of essential matrix from two affine correspondences. IEEE Trans. Image Processing **27**(11), 5328–5337 (2018)
8. Barath, D., Toth, T., Hajder, L.: A minimal solution for two-view focal-length estimation using two affine correspondences. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 2557–2565 (2017)
9. Byröd, M., Josephson, K., Åström, K.: A column-pivoting based strategy for monomial ordering in numerical gröbner basis calculations. In: European Conference on Computer Vision (ECCV), vol. 5305, pp. 130–143. Springer (2008)
10. Camposeco, F., Sattler, T., Pollefeys, M.: Minimal solvers for generalized pose and scale estimation from two rays and one point. In: ECCV – European Conference on Computer Vision. pp. 202–218 (2016)
11. Chen, H.H.: Pose determination from line-to-plane correspondences: existence condition and closed-form solutions. In: ICCV. pp. 374–378 (1990)
12. Dhome, M., Richetin, M., Lapreste, J.., Rives, G.: Determination of the attitude of 3d objects from a single perspective view. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(12), 1265–1278 (1989)
13. Duff, T., Hill, C., Jensen, A., Lee, K., Leykin, A., Sommars, J.: Solving polynomial systems via homotopy continuation and monodromy. IMA Journal of Numerical Analysis (2018)
14. Duff, T., Kohn, K., Leykin, A., Pajdla, T.: PLMP - Point-line minimal problems in complete multi-view visibility. In: International Conference on Computer Vision (ICCV) (2019)
15. Elqursh, A., Elgammal, A.M.: Line-based relative pose estimation. In: cvpr (2011)
16. Fabbri, R., Giblin, P., Kimia, B.: Camera pose estimation using first-order curve differential geometry. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020)
17. Fabbri, R., Giblin, P.J., Kimia, B.B.: Camera pose estimation using first-order curve differential geometry. In: Proceedings of the European Conference in Computer Vision (2012)
18. Fabbri, R., Kimia, B.B.: Multiview differential geometry of curves. International Journal of Computer Vision **120**(3), 324–346 (2016)
19. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24**(6), 381–395 (1981)
20. Grayson, D.R., Stillman, M.E.: Macaulay2, a software system for research in algebraic geometry. Available at http://www.math.uiuc.edu/Macaulay2/
21. Hartley, R., Li, H.: An efficient hidden variable approach to minimal-case camera motion estimation. IEEE PAMI **34**(12), 2303–2314 (2012)
22. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge, 2nd edn. (2003)
23. Hartley, R.I.: Lines and points in three views and the trifocal tensor. International Journal of Computer Vision **22**(2), 125–140 (1997)

24. Hauenstein, J.D., Rodriguez, J.I.: Multiprojective witness sets and a trace test. To appear in Advances in Geometry. arXiv preprint arXiv:1507.07069 (2019)
25. Johansson, B., Oskarsson, M., Åström, K.: Structure and motion estimation from complex features in three views. In: ICVGIP 2002, Proceedings of the Third Indian Conference on Computer Vision, Graphics & Image Processing, Ahmadabad, India, December 16-18, 2002 (2002)
26. Joswig, M., Kileel, J., Sturmfels, B., Wagner, A.: Rigid multiview varieties. IJAC **26**(4), 775–788 (2016). https://doi.org/10.1142/S021819671650034X, `https://doi.org/10.1142/S021819671650034X`
27. Kahl, F., Heyden, A., Quan, L.: Minimal projective reconstruction including missing data. IEEE Trans. Pattern Anal. Mach. Intell. **23**(4), 418–424 (2001). https://doi.org/10.1109/34.917578, `https://doi.org/10.1109/34.917578`
28. Kileel, J.: Minimal problems for the calibrated trifocal variety. SIAM Journal on Applied Algebra and Geometry **1**(1), 575–598 (2017)
29. Kneip, L., Scaramuzza, D., Siegwart, R.: A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In: CVPR – IEEE Conference on Computer Vision and Pattern Recognition. pp. 2969–2976 (2011)
30. Kneip, L., Siegwart, R., Pollefeys, M.: Finding the exact rotation between two images independently of the translation. In: ECCV – European Conference on Computer Vision. pp. 696–709 (2012)
31. Kuang, Y., Åström, K.: Pose estimation with unknown focal length using points, directions and lines. In: IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013. pp. 529–536 (2013)
32. Kuang, Y., Åström, K.: Stratified sensor network self-calibration from tdoa measurements. In: 21st European Signal Processing Conference (2013)
33. Kukelova, Z., Bujnak, M., Pajdla, T.: Automatic generator of minimal problem solvers. In: European Conference on Computer Vision (ECCV) (2008)
34. Kukelova, Z., Kileel, J., Sturmfels, B., Pajdla, T.: A clever elimination strategy for efficient minimal solvers. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2017)
35. Larsson, V., et. al: Automatic generator of minimal problems. http://www2.maths.lth.se/matematiklth/personal/viktorl/code/basis_selection.zip (2015)
36. Larsson, V., Åström, K., Oskarsson, M.: Efficient solvers for minimal problems by syzygy-based reduction. In: Computer Vision and Pattern Recognition (CVPR) (2017)
37. Larsson, V., Åström, K., Oskarsson, M.: Efficient solvers for minimal problems by syzygy-based reduction. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 2383–2392 (2017)
38. Larsson, V., Åström, K., Oskarsson, M.: Polynomial solvers for saturated ideals. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. pp. 2307–2316 (2017)
39. Larsson, V., Kukelova, Z., Zheng, Y.: Making minimal solvers for absolute pose estimation compact and robust. In: International Conference on Computer Vision (ICCV) (2017)
40. Larsson, V., Oskarsson, M., Åström, K., Wallis, A., Kukelova, Z., Pajdla, T.: Beyond grobner bases: Basis selection for minimal solvers. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp.

3945–3954 (2018), http://openaccess.thecvf.com/content_cvpr_2018/html/Larsson_Beyond_Grobner_Bases_CVPR_2018_paper.html
41. Leykin, A., Rodriguez, J.I., Sottile, F.: Trace test. Arnold Mathematical Journal **4**(1), 113–125 (2018)
42. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2), 91–110 (2004)
43. Ma, Y., Huang, K., Vidal, R., Kosecka, J., Sastry, S.: Rank conditions on the multiple-view matrix. International Journal of Computer Vision **59**(2), 115–137 (2004)
44. Matas, J., Obdrzálek, S., Chum, O.: Local affine frames for wide-baseline stereo. In: 16th International Conference on Pattern Recognition, ICPR 2002, Quebec, Canada, August 11-15, 2002. pp. 363–366. IEEE Computer Society (2002)
45. Matas, J., Obdrzálek, S., Chum, O.: Local affine frames for wide-baseline stereo. In: 16th International Conference on Pattern Recognition, ICPR 2002, Quebec, Canada, August 11-15, 2002. pp. 363–366 (2002)
46. Miraldo, P., Araujo, H.: Direct solution to the minimal generalized pose. Cybernetics, IEEE Transactions on **45**(3), 418–429 (2015)
47. Miraldo, P., Dias, T., Ramalingam, S.: A minimal closed-form solution for multi-perspective pose estimation using points and lines. In: Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XVI. pp. 490–507 (2018)
48. Mirzaei, F.M., Roumeliotis, S.I.: Optimal estimation of vanishing points in a manhattan world. In: International Conference on Computer Vision (ICCV) (2011)
49. Nistér, D.: An efficient solution to the five-point relative pose problem. IEEE Transactions on Pattern Analysis and Machine Intelligence **26**(6), 756–770 (Jun 2004)
50. Nistér, D., Naroditsky, O., Bergen, J.: Visual odometry. In: Computer Vision and Pattern Recognition (CVPR). pp. 652–659 (2004)
51. Nistér, D., Schaffalitzky, F.: Four points in two or three calibrated views: Theory and practice. International Journal of Computer Vision **67**(2), 211–231 (2006)
52. Oskarsson, M., Åström, K., Overgaard, N.C.: Classifying and solving minimal structure and motion problems with missing data. In: International Conference on Computer Vision (ICCV). pp. 628–634. IEEE Computer Society (2001). https://doi.org/10.1109/ICCV.2001.10072, http://doi.ieeecomputersociety.org/10.1109/ICCV.2001.10072
53. Oskarsson, M., Zisserman, A., Åström, K.: Minimal projective reconstruction for combinations of points and lines in three views. Image Vision Comput. **22**(10), 777–785 (2004)
54. Raguram, R., Chum, O., Pollefeys, M., Matas, J., Frahm, J.: USAC: A universal framework for random sample consensus. IEEE Transactions on Pattern Analysis Machine Intelligence **35**(8), 2022–2038 (2013)
55. Ramalingam, S., Bouaziz, S., Sturm, P.: Pose estimation using both points and lines for geo-localization. In: ICRA. pp. 4716–4723 (2011)
56. Ramalingam, S., Sturm, P.F.: Minimal solutions for generic imaging models. In: CVPR – IEEE Conference on Computer Vision and Pattern Recognition (2008)
57. Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Pajdla, T., Sivic, J.: Neighbourhood consensus networks (2018)
58. Salaün, Y., Marlet, R., Monasse, P.: Robust and accurate line- and/or point-based pose estimation without manhattan assumptions. In: European Conference on Computer Vision (ECCV) (2016)

59. Sattler, T., Leibe, B., Kobbelt, L.: Efficient & effective prioritized matching for large-scale image-based localization. IEEE Trans. Pattern Anal. Mach. Intell. **39**(9), 1744–1756 (2017)
60. Saurer, O., Pollefeys, M., Lee, G.H.: A minimal solution to the rolling shutter pose estimation problem. In: Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on. pp. 1328–1334. IEEE (2015)
61. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
62. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3D. In: ACM SIGGRAPH (2006)
63. Snavely, N., Seitz, S.M., Szeliski, R.: Modeling the world from internet photo collections. International Journal of Computer Vision (IJCV) **80**(2), 189–210 (2008)
64. Stewenius, H., Engels, C., Nistér, D.: Recent developments on direct relative orientation. ISPRS J. of Photogrammetry and Remote Sensing **60**, 284–294 (2006)
65. Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: InLoc: Indoor visual localization with dense matching and view synthesis. In: CVPR (2018)
66. Trager, M., Ponce, J., Hebert, M.: Trinocular geometry revisited. International Journal Computer Vision pp. 1–19 (March 2016)
67. Trager, M.: Cameras, Shapes, and Contours: Geometric Models in Computer Vision. (Caméras, formes et contours: modèles géométriques en vision par ordinateur). Ph.D. thesis, École Normale Supérieure, Paris, France (2018)
68. Ventura, J., Arth, C., Lepetit, V.: An efficient minimal solution for multi-camera motion. In: International Conference on Computer Vision (ICCV). pp. 747–755 (2015)
69. Xia, G., Delon, J., Gousseau, Y.: Accurate junction detection and characterization in natural images. International Journal of Computer Vision **106**(1), 31–56 (2014)