Numerically Computing Galois Groups of Minimal Problems

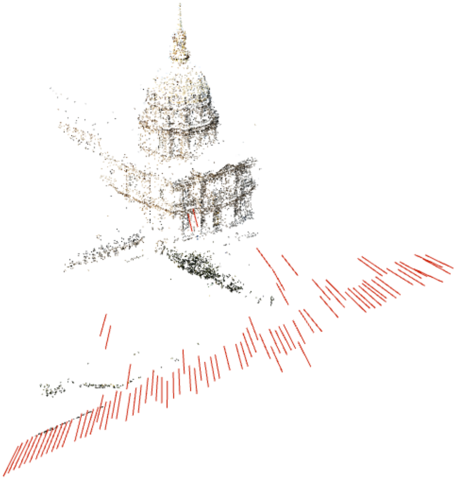

Tim Duff
University of Missouri - Columbia
ISSAC 2025, July 28
CIMAT, Guanajuato Mexico

Overview

The goal of this talk is to explain an unlikely intersection of three subjects:

1. *Computer vision* (mainly, minimal problems)
2. *Galois theory* for polynomial systems
3. *Numerical continuation methods methods* for solving these systems.

I'll also spend a large amount of time motivating the study of *minimal problems*.

For further reading and references: see the short article accompanying this tutorial (https://arxiv.org/abs/2507.10407, to appear in ISSAC 2025 Proceedings.)

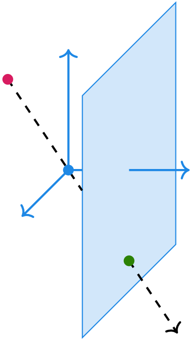

$$\mathbb{R}^3 \dashrightarrow \mathbb{R}^2$$

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} x/z \\ y/z \end{pmatrix}$$

Classical computer vision begins with the pinhole camera in standard coordinates, projecting a 3D point onto the plane $z = 1$.

Not linear, but projective linear. Represented by a 3×4 camera matrix

$$\begin{pmatrix} x/z \\ y/z \\ 1 \end{pmatrix} \sim \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

(Recall: n-dimensional projective space over the real numbers,

$$\mathbb{P}^n = \left(\mathbb{R}^{n+1} \setminus \{\mathbf{0}\}\right) / \left(\mathbf{x} \sim \lambda \mathbf{x} \quad \forall \lambda \in \mathbb{R} \setminus \{\mathbf{0}\}\right).$$

Points in $\mathbb{R}^n$ have homogeneous coordinates in $\mathbb{P}^n$.
If $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{n+1}$ represent the same point in $\mathbb{P}^n$, we write $\mathbf{x} \sim \mathbf{y}$.)

Anatomy of the pinhole camera

In general, we need coordinate systems for the camera ($\mathbf{A} \in \mathbb{P}\left(\mathrm{Hom}(\mathbb{R}^4, \mathbb{R}^3)\right) \cong \mathbb{P}^{11}$), the world points ($\mathbf{q} \in \mathbb{P}^3$), and the image points ($\mathbf{p} \in \mathbb{P}^2$).

A generic linear projection $\mathbf{A} : \mathbb{P}^3 \dashrightarrow \mathbb{P}^2$ has 11 degrees of freedom.
Their physical meaning can be seen from RQ decomposition,

$$\mathbf{A} \sim \begin{pmatrix} \alpha_x & s & x_0 \\ 0 & \alpha_y & y_0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \left(\mathbf{R} \mid \mathbf{t} \right),$$

where $\mathbf{R} \in \mathrm{SO}_3$ is a 3×3 rotation matrix, $\mathbf{t}$ a translation vector.

The camera matrix $\left(\mathbf{R} \mid \mathbf{t}\right)$ is said to be *calibrated*. It describes the position and orientation of the camera in space.

Intrinsic parameters $\alpha_x, s, x_0, \alpha_y, y_0$ determine the camera's pixel width, image center, aspect ratio, skew, and focal length. They are often (not always!) known in practice.

Classical camera geometry problems

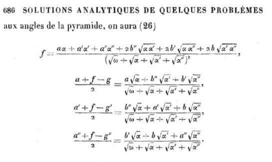

686 SOLUTIONS ANALYTIQUES DE QUELQUES PROBLÈMES

aux angles de la pyramide, on aura (26)

$$f = \frac{a\alpha + a'\alpha' + a''\alpha'' + 2b''\sqrt{\alpha\alpha'} + 2b'\sqrt{\alpha\alpha''} + 2b\sqrt{\alpha'\alpha''}}{(\sqrt{\omega} + \sqrt{\alpha} + \sqrt{\alpha'} + \sqrt{\alpha''})^2},$$

$$\frac{a + f - g}{2} = \frac{a\sqrt{\alpha} + b''\sqrt{\alpha'} + b'\sqrt{\alpha''}}{\sqrt{\omega} + \sqrt{\alpha} + \sqrt{\alpha'} + \sqrt{\alpha''}},$$

$$\frac{a' + f - g'}{2} = \frac{b''\sqrt{\alpha} + a'\sqrt{\alpha'} + b\sqrt{\alpha''}}{\sqrt{\omega} + \sqrt{\alpha} + \sqrt{\alpha'} + \sqrt{\alpha''}},$$

$$\frac{a'' + f - g''}{2} = \frac{b'\sqrt{\alpha} + b\sqrt{\alpha'} + a''\sqrt{\alpha''}}{\sqrt{\omega} + \sqrt{\alpha} + \sqrt{\alpha'} + \sqrt{\alpha''}}.$$

Lagrange (1773)
P3P (Resectioning)

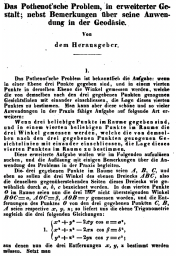

Das Pothenot'sche Problem, in erweiterter Gestalt; nebst Bemerkungen über seine Anwendung in der Geodäsie.

Von

dem Herausgeber.

1.

Das Pothenot'sche Problem ist bekanntlich die Aufgabe: wenn in einer Ebene drei Punkte gegeben sind, und in einem vierten Punkte in derselben Ebene die Winkel gemessen werden, welche die von demselben nach den drei gegebenen Punkten gezogenen Gesichtslinien mit einander einschliessen, die Lage dieses vierten Punktes zu bestimmen. Man kann aber diese schöne und so vieler Anwendungen in der Praxis fähige Aufgabe auf folgende Art erweitern:

Wenn drei beliebige Punkte im Raume gegeben sind, und in einem vierten beliebigen Punkte im Raume die drei Winkel gemessen werden, welche die von demselben nach den drei gegebenen Punkten gezogenen Gesichtslinien mit einander einschliessen, die Lage dieses vierten Punktes im Raume zu bestimmen.

Diese erweiterte Aufgabe wollen wir im Folgenden aufzulösen suchen, und die Auflösung mit einigen Bemerkungen über die Anwendung des Problems in der Praxis begleiten.

Die drei gegebenen Punkte im Raume seien A, B, C, und eben so sollen die drei Winkel des ebenen Dreiecks ABC, also die denselben gegenüberstehenden Seiten dieses Dreiecks wie gewöhnlich durch a, b, c bezeichnet werden. In dem vierten Punkte O im Raume seien nun die drei 180° nicht übersteigenden Winkel $BOC=\alpha$, $AOC=\beta$, $AOB=\gamma$ gemessen worden, und die Entfernungen des Punktes O von den drei gegebenen Punkten C, B, A seien respective x, y, z; so liefert uns die ebene Trigonometrie sogleich die drei folgenden Gleichungen:

$$1. \begin{cases} x^2 + y^2 - 2xy \cos\alpha = a^2, \\ x^2 + z^2 - 2xz \cos\beta = b^2, \\ y^2 + z^2 - 2yz \cos\gamma = c^2; \end{cases}$$

aus denen nun die drei Entfernungen x, y, z bestimmt werden müssen. Setzt man

Grunert (1841)
P3P (Resectioning)

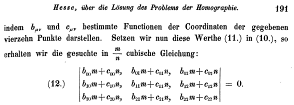

Hesse, *über die Lösung des Problems der Homographie.* 191

indem $b_{\mu\nu}$ und $c_{\mu\nu}$ bestimmte Functionen der Coordinaten der gegebenen vierzehn Punkte darstellen. Setzen wir nun diese Werthe (11.) in (10.), so erhalten wir die gesuchte in $\frac{m}{n}$ cubische Gleichung:

$$(12.)\quad \begin{vmatrix} b_{00}m + c_{00}n, & b_{01}m + c_{01}n, & b_{02}m + c_{02}n \\ b_{10}m + c_{10}n, & b_{11}m + c_{11}n, & b_{12}m + c_{12}n \\ b_{20}m + c_{20}n, & b_{21}m + c_{21}n, & b_{22}m + c_{22}n \end{vmatrix} = 0.$$

Hesse (1863)
7-point method (Reconstruction)

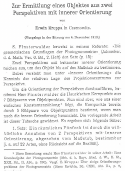

Zur Ermittlung eines Objektes aus zwei Perspektiven mit innerer Orientierung

von

Erwin Kruppa in Czernowitz.

(Vorgelegt in der Sitzung am 4. Dezember 1913.)

S. Finsterwalder beweist in seinem Referate: »Die geometrischen Grundlagen der Photogrammetrie« (Jahresber. d. d. Math. Ver. 6. Bd., 2. Heft) den Satz (p. 15):

Zwei Perspektiven mit bekannter innerer Orientierung reichen aus, um das Objekt bis auf den Maßstab zu bestimmen.

Dabei versteht man unter »innerer Orientierung« die Kenntnis der relativen Lage des Projektionszentrums zur Perspektive.

Um die Orientierung der Perspektiven durchzuführen, bestimmt Herr Finsterwalder die Hauck'schen Kernpunkte aus 7 Bildpaaren von Objektpunkten. Nun sind aber, wie aus einer einfachen Konstantenzählung[1] folgt, die Kernpunkte bereits durch 5 Bildpaare von Objektpunkten bestimmt, wenn man noch die innere Orientierung heranzieht. Die vorliegende Arbeit ist dieser Tatsache gewidmet. Es ergeben sich folgende Sätze:

1. Satz: Ein räumliches Fünfeck ist durch die willkürliche Annahme von 2 Perspektiven mit innerer Orientierung, abgesehen vom Maßstab, bestimmbar [i. a. auf 22 Arten, ohne Rücksicht auf die Realität].

[1] Diese Bemerkung macht Herr Finsterwalder in seiner Arbeit: Eine Grundaufgabe der Photogrammetrie (Abh. d. k. Bayr. Akad. d. W., II. Kl., XXII. Bd., II. Abt., p. 230). Vergl. E. Kruppa: Über einige Orientierungsprobleme der Photogrammetrie (Diese Berichte, Bd. CXXI, Abt. IIa, p. 9,

Kruppa (1913)
5-point method (Reconstruction)

Perspective 3-Point Problem (aka P3P, Calibrated Resectioning, ...)

Given: 3D points $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3 \in \mathbb{P}^3$ and matching 2D points $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{P}^2$

Unknown: calibrated camera $(\mathbf{R} \;\; \mathbf{t})$ such that $\mathbf{p}_i \sim (\mathbf{R} \;\; \mathbf{t})\mathbf{q}_i$ for $i = 1, 2, 3$.

Perspective 3-Point Problem (aka P3P, Calibrated Resectioning, ...)

Given: 3D points $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3 \in \mathbb{P}^3$ and matching 2D points $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{P}^2$

Unknown: calibrated camera $(\mathbf{R} \ \ \mathbf{t})$ such that $\mathbf{p}_i \sim (\mathbf{R} \ \ \mathbf{t})\mathbf{q}_i$ for $i = 1, 2, 3$.

Choose projective representatives such that $\mathbf{q}_i = \begin{pmatrix} q_{i1} & q_{i2} & q_{i3} & 1 \end{pmatrix}^T$, and $\mathbf{p}_i{}^T \mathbf{p}_i = 1$.

Define (unknown) scalar *depths* $\lambda_1, \lambda_2, \lambda_3$ such that $\lambda_i \mathbf{p}_i = (\mathbf{R} \ \ \mathbf{t})\mathbf{q}_i$.

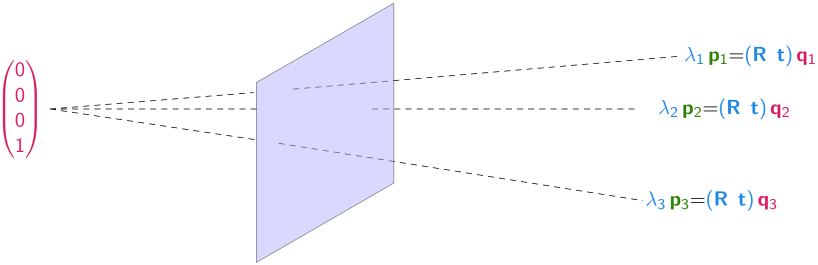

Grunert (1847) derived 3 polynomial equations in 3 unknowns: for $1 \leq i < j \leq 3$,

$$\lambda_i{}^2 + \lambda_j{}^2 - 2\left(\mathbf{p}_i{}^T \mathbf{p}_j\right) \lambda_i \lambda_j = (\mathbf{q}_i - \mathbf{q}_j)^T (\mathbf{q}_i - \mathbf{q}_j).$$

For generic data $(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3)$, this system has 8 (complex-valued) solutions.

Universal problem of algebraic vision: *Let* $\pi : \mathcal{X} \dashrightarrow \mathbb{R}^m$ *be a map of algebraic varieties, where* $\mathcal{X}$ *is a space of unknown states and* $\pi(\mathcal{X}) \subset \mathbb{R}^m$ *is a space of idealized data. Given a measurement* $\widetilde{y} \approx y$ *of "true" data* $y = \pi(x)$*, recover* $\widetilde{x} \in \mathcal{X}$ *with* $\widetilde{x} \approx x$*.*

Universal problem of algebraic vision: *Let* $\pi : \mathcal{X} \dashrightarrow \mathbb{R}^m$ *be a map of algebraic varieties, where* $\mathcal{X}$ *is a space of unknown states and* $\pi(\mathcal{X}) \subset \mathbb{R}^m$ *is a space of idealized data. Given a measurement* $\widetilde{y} \approx y$ *of "true" data* $y = \pi(x)$, *recover* $\widetilde{x} \in \mathcal{X}$ *with* $\widetilde{x} \approx x$.

Example 1: Perspective n-point (PnP / calibrated resectioning)

$$\pi_{\mathrm{PnP}} : \mathrm{SE}_3 \to \left(\mathbb{R}^2\right)^n$$

$$(\mathbf{R} \mid \mathbf{t}) \mapsto \left(\Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_1\right), \dots, \Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_n\right)\right)$$

$$\text{where} \quad \Pi(x, y, z) = (x/z, y/z).$$

Universal problem of algebraic vision: *Let* $\pi : \mathcal{X} \dashrightarrow \mathbb{R}^m$ *be a map of algebraic varieties, where* $\mathcal{X}$ *is a space of unknown states and* $\pi(\mathcal{X}) \subset \mathbb{R}^m$ *is a space of idealized data. Given a measurement* $\widetilde{y} \approx y$ *of "true" data* $y = \pi(x)$*, recover* $\widetilde{x} \in \mathcal{X}$ *with* $\widetilde{x} \approx x$.

Example 1: Perspective n-point (PnP / calibrated resectioning)

$$\pi_{\mathrm{PnP}} : \mathrm{SE}_3 \to \left(\mathbb{R}^2\right)^n$$
$$(\mathbf{R} \mid \mathbf{t}) \mapsto \left(\Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_1\right), \ldots, \Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_n\right)\right)$$
$$\text{where} \quad \Pi(x, y, z) = (x/z, y/z).$$

Three regimes:

Universal problem of algebraic vision: *Let* $\pi : \mathcal{X} \dashrightarrow \mathbb{R}^m$ *be a map of algebraic varieties, where* $\mathcal{X}$ *is a space of unknown states and* $\pi(\mathcal{X}) \subset \mathbb{R}^m$ *is a space of idealized data. Given a measurement* $\widetilde{y} \approx y$ *of "true" data* $y = \pi(x)$*, recover* $\widetilde{x} \in \mathcal{X}$ *with* $\widetilde{x} \approx x$*.*

Example 1: Perspective n-point (PnP / calibrated resectioning)

$$\begin{aligned} \pi_{\mathrm{PnP}} : \mathrm{SE}_3 &\to \left(\mathbb{R}^2\right)^n \\ (\mathbf{R} \mid \mathbf{t}) &\mapsto (\Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_1\right), \ldots, \Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_n\right)) \\ &\text{where} \quad \Pi(x, y, z) = (x/z, y/z). \end{aligned}$$

Three regimes:

- **Underconstrained**, $\dim(\mathcal{X}) > \dim(\pi(\mathcal{X})) = m$: eg. P1P / P2P

Universal problem of algebraic vision: *Let* $\pi : \mathcal{X} \dashrightarrow \mathbb{R}^m$ *be a map of algebraic varieties, where* $\mathcal{X}$ *is a space of unknown states and* $\pi(\mathcal{X}) \subset \mathbb{R}^m$ *is a space of idealized data. Given a measurement* $\widetilde{y} \approx y$ *of "true" data* $y = \pi(x)$*, recover* $\widetilde{x} \in \mathcal{X}$ *with* $\widetilde{x} \approx x$*.*

Example 1: Perspective n-point (PnP / calibrated resectioning)

$$\pi_{\mathrm{PnP}} : \mathrm{SE}_3 \to \left(\mathbb{R}^2\right)^n$$
$$(\mathbf{R} \mid \mathbf{t}) \mapsto (\Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_1\right), \dots, \Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_n\right))$$
$$\text{where} \quad \Pi(x, y, z) = (x/z, y/z).$$

Three regimes:

- **Underconstrained**, $\dim(\mathcal{X}) > \dim(\pi(\mathcal{X})) = m$: eg. P1P / P2P
- **Overconstrained**, $\dim(\pi(\mathcal{X})) < m$: eg. P4P, P5P, ...

Universal problem of algebraic vision: *Let* $\pi : \mathcal{X} \dashrightarrow \mathbb{R}^m$ *be a map of algebraic varieties, where* $\mathcal{X}$ *is a space of unknown states and* $\pi(\mathcal{X}) \subset \mathbb{R}^m$ *is a space of idealized data. Given a measurement* $\widetilde{y} \approx y$ *of "true" data* $y = \pi(x)$*, recover* $\widetilde{x} \in \mathcal{X}$ *with* $\widetilde{x} \approx x$*.*

Example 1: Perspective n-point (PnP / calibrated resectioning)

$$\pi_{\mathrm{PnP}} : \mathrm{SE}_3 \to \left(\mathbb{R}^2\right)^n$$
$$(\mathbf{R} \mid \mathbf{t}) \mapsto \left(\Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_1\right), \ldots, \Pi\left((\mathbf{R} \mid \mathbf{t})\mathbf{q}_n\right)\right)$$
$$\text{where} \quad \Pi(x, y, z) = (x/z, y/z).$$

Three regimes:

- **Underconstrained**, $\dim(\mathcal{X}) > \dim(\pi(\mathcal{X})) = m$: eg. P1P / P2P
- **Overconstrained**, $\dim(\pi(\mathcal{X})) < m$: eg. P4P, P5P, ...
- **Minimal / well-constrained**, $\dim(\mathcal{X}) = \dim(\pi(\mathcal{X})) = m$: eg. P3P

Underconstrained regime: infinitely-many (complex) solutions, so exact recovery of $(\mathbf{R} \mid \mathbf{t})$ is hopeless. Still, there may be *some* constraints worth studying.

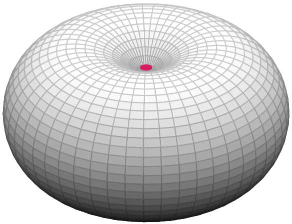

Exercise: For two generic 3D-2D point matches, $(\mathbf{q}_1, \mathbf{p}_1), (\mathbf{q}_2, \mathbf{p}_2)$, viewed by some calibrated camera $\mathbf{A} = (\mathbf{R} \mid \mathbf{t})$, show that the *camera center* $[\ker \mathbf{A}] \in \mathbb{P}^3$, where

$$\ker \mathbf{A} = \operatorname{span} \left\{ \begin{bmatrix} -\mathbf{R}^T \mathbf{t} \\ 1 \end{bmatrix} \right\}$$

must lie on a quartic surface with two singular points at the $\mathbf{q}_i$, and the equation of this surface is *independent* of the camera orientation $\mathbf{R}$.

Overconstrained regime: again, exact recovery is hopeless. One must choose some objective to minimize, eg. the *reprojection error*,

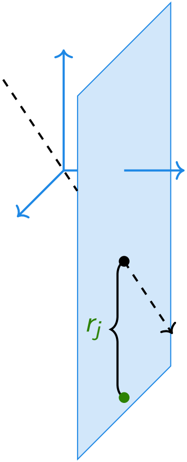

$$\min_{\mathbf{A}} \left(\sum_{1 \le j \le n} \underbrace{\left(\frac{\mathbf{p}_{ij}[1]}{\mathbf{p}_{ij}[3]} - \frac{\mathbf{A}[1,:]\,\mathbf{q}_j}{\mathbf{A}[3,:]\,\mathbf{q}_j} \right)^2 + \left(\frac{\mathbf{p}_{ij}[2]}{\mathbf{p}_j[3]} - \frac{\mathbf{A}[2,:]\,\mathbf{q}_j}{\mathbf{A}[3,:]\,\mathbf{q}_j} \right)^2}_{r_j} \right) \tag{1}$$

Local optimization requires an initial guess. Minimal solvers help here!

Overconstrained regime: again, exact recovery is hopeless. One must choose some objective to minimize, eg. the *reprojection error*,

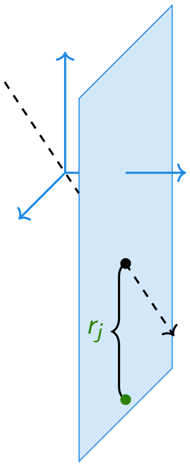

$$\min_{\mathbf{A}} \left(\sum_{1 \le j \le n} \underbrace{\left(\frac{\mathbf{p}_{ij}[1]}{\mathbf{p}_{ij}[3]} - \frac{\mathbf{A}[1,:]\,\mathbf{q}_j}{\mathbf{A}[3,:]\,\mathbf{q}_j} \right)^2 + \left(\frac{\mathbf{p}_{ij}[2]}{\mathbf{p}_{j}[3]} - \frac{\mathbf{A}[2,:]\,\mathbf{q}_j}{\mathbf{A}[3,:]\,\mathbf{q}_j} \right)^2}_{r_j} \right) \tag{1}$$

Local optimization requires an initial guess. Minimal solvers help here!

Global optimization is oftentimes impractical...

Conjecture (Connelly-D- Loucks-Tavitas, Math. Comp. '24)

For uncalibrated cameras with $n \geq 6$ measurements, the number of complex-valued regular critical points of (1) equals

$$(80/3)n^3 - 368n^2 + (5068/3)n - 2580.$$

(Open) How many critical points if $\mathbf{A}$ is constrained to be calibrated?

Minimal regime (P3P cont.)

Given a solution $(\lambda_1, \lambda_2, \lambda_3)$ to Grunert's equations,

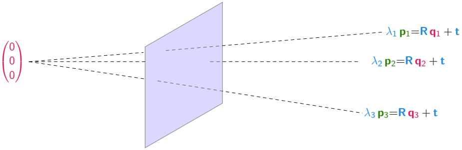

$$\lambda_i{}^2 + \lambda_j{}^2 - 2\left(\mathbf{p}_i{}^T \mathbf{p}_j\right) \lambda_i \lambda_j = (\mathbf{q}_i - \mathbf{q}_j)^T (\mathbf{q}_i - \mathbf{q}_j), \quad 1 \le i < j \le 3, \quad \underline{\mathbf{q}_i \in \mathbb{R}^3},$$

we can recover the calibrated camera $(\mathbf{R} \mid \mathbf{t}) \in SE_3$ as follows:

$$\begin{aligned}
\mathbf{R} &= \mathbf{P}\mathbf{Q}^{-1}, \quad \text{where} \\
\mathbf{P} &= (\lambda_1 \mathbf{p}_1 - \lambda_2 \mathbf{p}_2 \mid \lambda_1 \mathbf{p}_1 - \lambda_3 \mathbf{p}_3 \mid (\lambda_1 \mathbf{p}_1 - \lambda_2 \mathbf{p}_2) \times (\lambda_1 \mathbf{p}_1 - \lambda_3 \mathbf{p}_3)) \\
\mathbf{Q} &= (\mathbf{q}_1 - \mathbf{q}_2 \mid \mathbf{q}_1 - \mathbf{q}_3 \mid (\mathbf{q}_1 - \mathbf{q}_2) \times (\mathbf{q}_1 - \mathbf{q}_3)), \text{ and} \\
\mathbf{t} &= \lambda_i \mathbf{p}_i - \mathbf{R}\mathbf{q}_i \quad (\text{any } i.)
\end{aligned}$$

Minimal solvers (eg. Ding et al., CVPR '23) can run in less than a microsecond!

Why does anybody care?

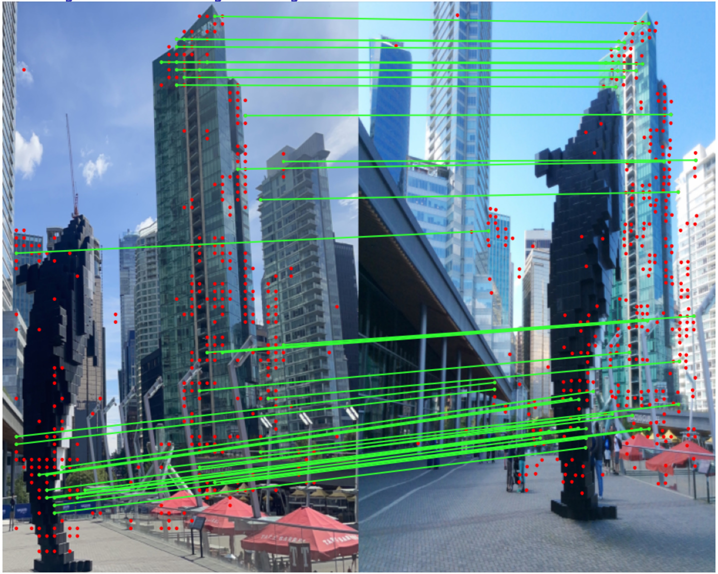

retrieved image (known 3D) vs query image (only 2D)

PnP assumes 3D-2D matches.
Is that reasonable?

Why does anybody care?

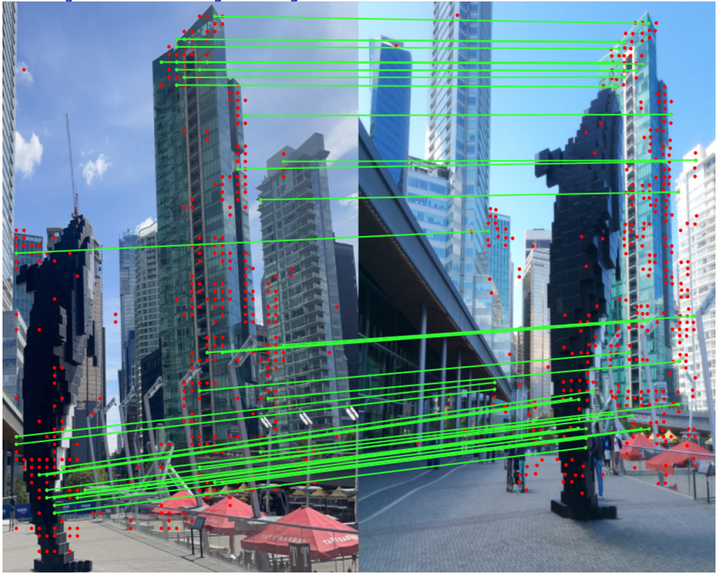

retrieved image (known 3D) vs query image (only 2D)

PnP assumes 3D-2D matches.
Is that reasonable?

Random **Sa**mpling and **C**onsensus, aka **RANSAC**, is used in computer vision for outlier-robust model-fitting.

Why does anybody care?

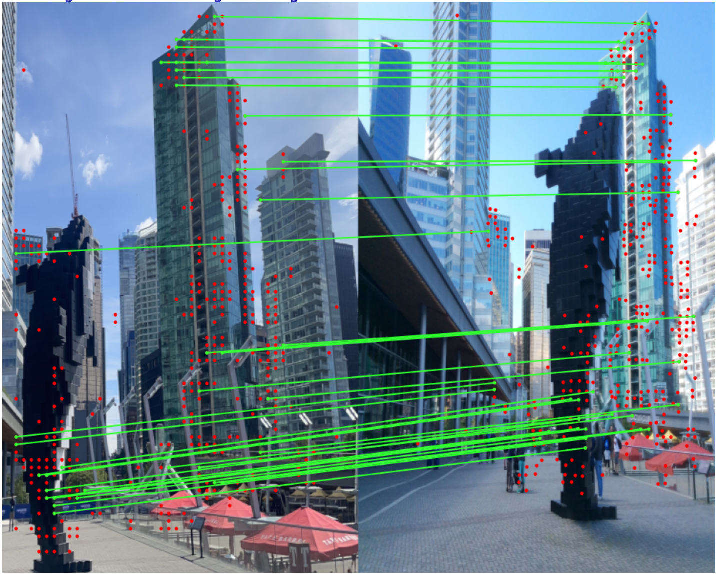

retrieved image (known 3D) vs query image (only 2D)

PnP assumes 3D-2D matches.
Is that reasonable?

Random **Sa**mpling and **C**onsensus, aka **RANSAC**, is used in computer vision for outlier-robust model-fitting.

Over N trials:

Why does anybody care?

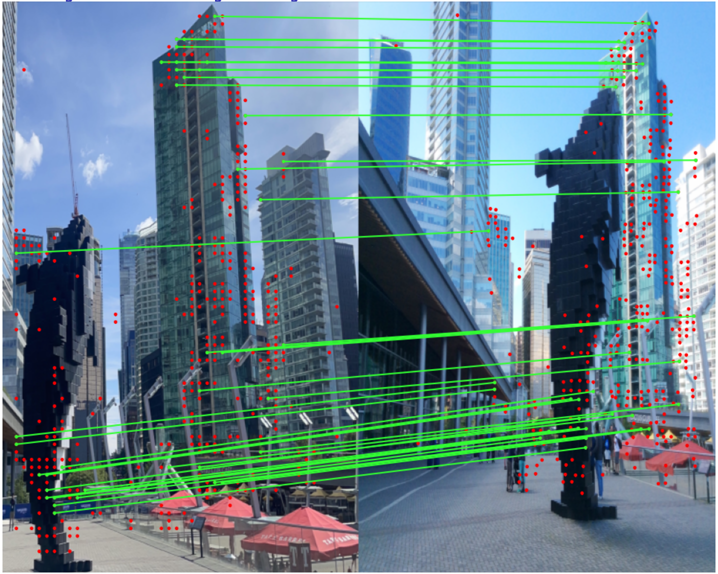

retrieved image (known 3D) vs query image (only 2D)

PnP assumes 3D-2D matches.
Is that reasonable?

Random **Sa**mpling and **C**onsensus, aka **RANSAC**, is used in computer vision for outlier-robust model-fitting.

Over N trials:

1. Sample k out of n 3D-2D matches, uniformly at random

Why does anybody care?

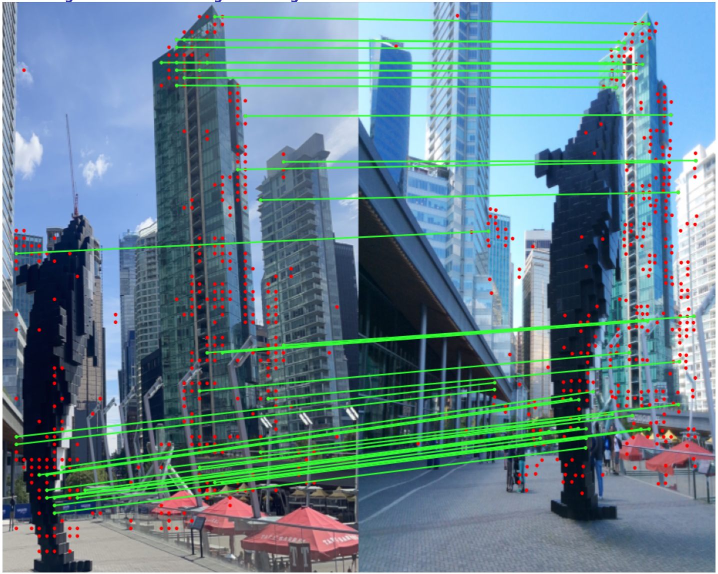

retrieved image (known 3D) vs query image (only 2D)

PnP assumes 3D-2D matches.
Is that reasonable?

Random **S**ampling and **C**onsensus, aka **RANSAC**, is used in computer vision for outlier-robust model-fitting.

Over N trials:

1. Sample k out of n 3D-2D matches, uniformly at random
2. (Somehow) solve PkP

Why does anybody care?

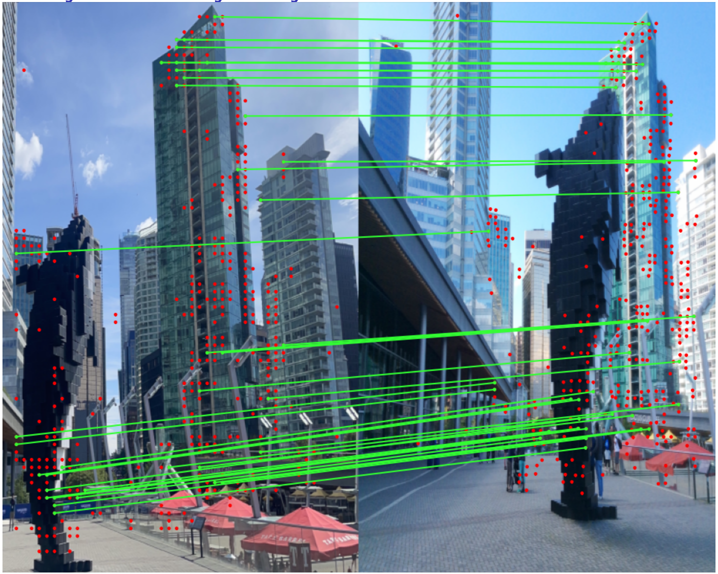

retrieved image (known 3D) vs query image (only 2D)

PnP assumes 3D-2D matches.
Is that reasonable?

Random **Sa**mpling and **C**onsensus, aka **RANSAC**, is used in computer vision for outlier-robust model-fitting.

Over N trials:

1. Sample k out of n 3D-2D matches, uniformly at random
2. (Somehow) solve PkP
3. Measure *consensus* of solutions on remaining samples, and keep the *maximum-consensus solution*

RanSaC, analyzed

Assume we have the following:

RanSaC, analyzed

Assume we have the following:

1. N, the number of trials,

RanSaC, analyzed

Assume we have the following:

1. N, the number of trials,
2. a PkP solver,

RanSaC, analyzed

Assume we have the following:

1. N, the number of trials,
2. a PkP solver,
3. a consensus criterion (depends on some threshold),

RanSaC, analyzed

Assume we have the following:

1. N, the number of trials,
2. a PkP solver,
3. a consensus criterion (depends on some threshold),
4. a confidence $s \in (0, 1)$, the desired probability of obtaining an outlier-free sample of k matches after N trials.

RanSaC, analyzed

Assume we have the following:

1. N, the number of trials,
2. a PkP solver,
3. a consensus criterion (depends on some threshold),
4. a confidence $s \in (0,1)$, the desired probability of obtaining an outlier-free sample of k matches after N trials.

Let $p \in (0,1)$ denote the fraction of erroneous matches, so that the probability of drawing an all-inlier sample in one trial is

$$P = \frac{\binom{pn}{k}}{\binom{n}{k}}. \tag{2}$$

From our specification above, we should have

$$(1-P)^N \leq 1-s \quad \Rightarrow \quad N \geq \frac{\log(1-s)}{\log(1-P)}. \tag{3}$$

Minimal Solvers Maximize Sample Efficiency!

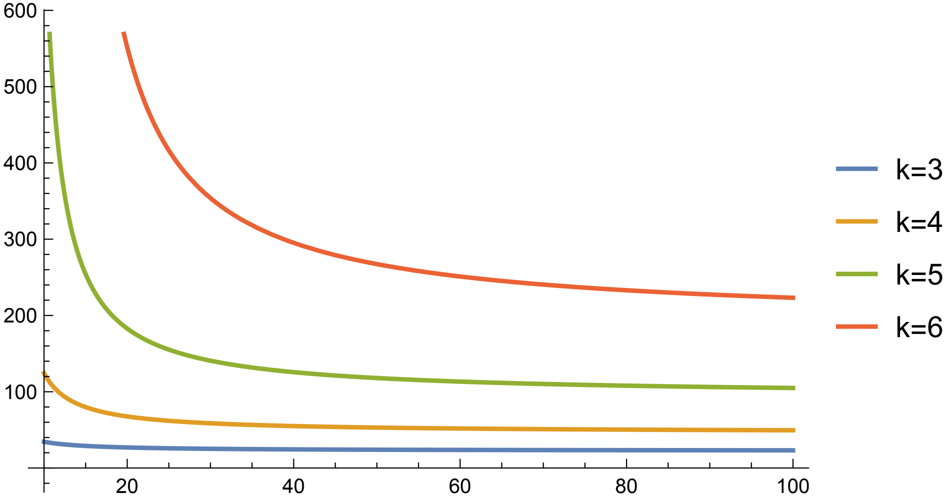

Figure: RanSaC trials N needed to find an outlier-free subsample of size k with confidence $s = .95$ from $n \in [10, 100]$ total matches, and with $p = .5$ outlier probability.

Back to the P3P

$$
\begin{aligned}
{\lambda_1}^2 + {\lambda_2}^2 - 2\left(\mathbf{p}_1{}^T\mathbf{p}_2\right)\lambda_1\lambda_2 &= (\mathbf{q}_1 - \mathbf{q}_2)^T(\mathbf{q}_1 - \mathbf{q}_2),\\
{\lambda_1}^2 + {\lambda_3}^2 - 2\left(\mathbf{p}_1{}^T\mathbf{p}_3\right)\lambda_1\lambda_3 &= (\mathbf{q}_1 - \mathbf{q}_3)^T(\mathbf{q}_1 - \mathbf{q}_3),\\
{\lambda_2}^2 + {\lambda_3}^2 - 2\left(\mathbf{p}_2{}^T\mathbf{p}_3\right)\lambda_2\lambda_3 &= (\mathbf{q}_2 - \mathbf{q}_3)^T(\mathbf{q}_2 - \mathbf{q}_3).
\end{aligned}
\tag{4}
$$

Three quadrics in three unknowns generally have $2^3 = 8$ solutions.

This is indeed the case for the system above.

P3P is an *easy* minimal problem—in fact, it is even easier than solving 3 quadrics in 3 unknowns, because of a $\mathbb{Z}/2\mathbb{Z}$-*symmetry*.

Back to the P3P

$$
\begin{aligned}
{\lambda_1}^2 + {\lambda_2}^2 - 2\left({\mathbf{p}_1}^T\mathbf{p}_2\right)\lambda_1\lambda_2 &= (\mathbf{q}_1 - \mathbf{q}_2)^T(\mathbf{q}_1 - \mathbf{q}_2), \\
{\lambda_1}^2 + {\lambda_3}^2 - 2\left({\mathbf{p}_1}^T\mathbf{p}_3\right)\lambda_1\lambda_3 &= (\mathbf{q}_1 - \mathbf{q}_3)^T(\mathbf{q}_1 - \mathbf{q}_3), \\
{\lambda_2}^2 + {\lambda_3}^2 - 2\left({\mathbf{p}_2}^T\mathbf{p}_3\right)\lambda_2\lambda_3 &= (\mathbf{q}_2 - \mathbf{q}_3)^T(\mathbf{q}_2 - \mathbf{q}_3).
\end{aligned} \tag{4}
$$

Three quadrics in three unknowns generally have $2^3 = 8$ solutions.

This is indeed the case for the system above.

P3P is an *easy* minimal problem—in fact, it is even easier than solving 3 quadrics in 3 unknowns, because of a $\mathbb{Z}/2\mathbb{Z}$-*symmetry*.

$$(\lambda_1, \lambda_2, \lambda_3) \text{ solves (4)} \quad \Leftrightarrow \quad (-\lambda_1, -\lambda_2, -\lambda_3) \text{ solves (4)}$$

Decomposing P3P

Consider the rational $\mathbb{Z}/2\mathbb{Z}$-invariants

$$(\rho_1, \rho_2, \rho_3) = (\lambda_1/\lambda_3, \lambda_2/\lambda_3, {\lambda_3}^2),$$

and set, for $1 \leq i < j \leq 3$,

$$c_{ij} = -2{\mathbf{p}_i}^T \mathbf{p}_j, \quad d_{ij} = (\mathbf{q}_i - \mathbf{q}_j)^T(\mathbf{q}_i - \mathbf{q}_j).$$

Our system:

$$\begin{aligned}
{\lambda_1}^2 + {\lambda_2}^2 - 2\left({\mathbf{p}_1}^T \mathbf{p}_2\right) \lambda_1\lambda_2 &= (\mathbf{q}_1 - \mathbf{q}_2)^T(\mathbf{q}_1 - \mathbf{q}_2),\\
{\lambda_1}^2 + {\lambda_3}^2 - 2\left({\mathbf{p}_1}^T \mathbf{p}_3\right) \lambda_1\lambda_3 &= (\mathbf{q}_1 - \mathbf{q}_3)^T(\mathbf{q}_1 - \mathbf{q}_3),\\
{\lambda_2}^2 + {\lambda_3}^2 - 2\left({\mathbf{p}_2}^T \mathbf{p}_3\right) \lambda_2\lambda_3 &= (\mathbf{q}_2 - \mathbf{q}_3)^T(\mathbf{q}_2 - \mathbf{q}_3).
\end{aligned}$$

Decomposing P3P

Consider the rational $\mathbb{Z}/2\mathbb{Z}$-invariants

$$(\rho_1, \rho_2, \rho_3) = (\lambda_1/\lambda_3, \lambda_2/\lambda_3, {\lambda_3}^2),$$

and set for $1 \leq i < j \leq 3$,

$$c_{ij} = -2{\mathbf{p}_i}^T \mathbf{p}_j, \quad d_{ij} = (\mathbf{q}_i - \mathbf{q}_j)^T (\mathbf{q}_i - \mathbf{q}_j).$$

Our system:

$$\begin{aligned} {\lambda_1}^2 + {\lambda_2}^2 + c_{12}\lambda_1\lambda_2 &= d_{12}, \\ {\lambda_1}^2 + {\lambda_3}^2 + c_{13}\lambda_1\lambda_3 &= d_{13}, \\ {\lambda_2}^2 + {\lambda_3}^2 + c_{23}\lambda_2\lambda_3 &= d_{23}. \end{aligned}$$

Decomposing P3P

Consider the rational $\mathbb{Z}/2\mathbb{Z}$-invariants

$$(\rho_1, \rho_2, \rho_3) = (\lambda_1/\lambda_3, \lambda_2/\lambda_3, {\lambda_3}^2),$$

and set for $1 \leq i < j \leq 3$,

$$c_{ij} = -2{\mathbf{p}_i}^T \mathbf{p}_j, \quad d_{ij} = (\mathbf{q}_i - \mathbf{q}_j)^T(\mathbf{q}_i - \mathbf{q}_j).$$

Our system:

$$\begin{aligned} {\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2 &= d_{12}\rho_3, \\ 1 + {\rho_1}^2 + c_{13}\rho_1 &= d_{13}\rho_3, \\ 1 + {\rho_2}^2 + c_{23}\rho_2 &= d_{23}\rho_3. \end{aligned}$$

Decomposing P3P

Consider the rational $\mathbb{Z}/2\mathbb{Z}$-invariants

$$(\rho_1, \rho_2, \rho_3) = (\lambda_1/\lambda_3, \lambda_2/\lambda_3, {\lambda_3}^2),$$

and set for $1 \leq i < j \leq 3$,

$$c_{ij} = -2{\mathbf{p}_i}^T\mathbf{p}_j, \quad d_{ij} = (\mathbf{q}_i - \mathbf{q}_j)^T(\mathbf{q}_i - \mathbf{q}_j).$$

Our system:

$$\begin{aligned} {\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2 &= d_{12}\rho_3, \\ d_{13}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_1}^2 + c_{13}\rho_1\right), \\ d_{23}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_2}^2 + c_{23}\rho_2\right). \end{aligned}$$

Decomposing P3P

Consider the rational $\mathbb{Z}/2\mathbb{Z}$-invariants

$$(\rho_1, \rho_2, \rho_3) = (\lambda_1/\lambda_3, \lambda_2/\lambda_3, {\lambda_3}^2),$$

and set for $1 \leq i < j \leq 3$,

$$c_{ij} = -2{\mathbf{p}_i}^T \mathbf{p}_j, \quad d_{ij} = (\mathbf{q}_i - \mathbf{q}_j)^T(\mathbf{q}_i - \mathbf{q}_j).$$

Our system:

$$\left\{ {\rho_1}^2 + {\rho_2}^2 - 2c_{12}\rho_1\rho_2 = d_{12}\rho_3, \right.$$

$\rightarrow$ linear in ρ_3

$$\begin{cases} d_{13}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) & = d_{12}\left(1 + {\rho_1}^2 + c_{13}\rho_1\right), \\ d_{23}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) & = d_{12}\left(1 + {\rho_2}^2 + c_{23}\rho_2\right) \end{cases}$$

$\rightarrow$ intersection of two plane conics!

Decomposing P3P — Summary

Decomposing P3P — Summary

Instead of solving a degree-8 problem, we can:

1. First, solve a degree-4 problem,

$$\begin{aligned}
{\rho_1}^2 + {\rho_2}^2 - 2c_{12}\rho_1\rho_2 &= d_{12}\rho_3, \\
d_{13}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_1}^2 + c_{13}\rho_1\right), \\
d_{23}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_2}^2 + c_{23}\rho_2\right).
\end{aligned}$$

Decomposing P3P — Summary

Instead of solving a degree-8 problem, we can:

1. First, solve a degree-4 problem,

$$\begin{aligned} {\rho_1}^2 + {\rho_2}^2 - 2c_{12}\rho_1\rho_2 &= d_{12}\rho_3, \\ d_{13}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_1}^2 + c_{13}\rho_1\right), \\ d_{23}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_2}^2 + c_{23}\rho_2\right). \end{aligned}$$

2. Then, recover depths by solving a degree-2 problem,

$$\begin{aligned} \lambda_1 &= \rho_1\lambda_3, \\ \lambda_2 &= \rho_2\lambda_3, \\ {\lambda_3}^2 &= \rho_3. \end{aligned}$$

Decomposing P3P — Summary

Instead of solving a degree-8 problem, we can:

1. First, solve a degree-4 problem,

$$\begin{aligned} {\rho_1}^2 + {\rho_2}^2 - 2c_{12}\rho_1\rho_2 &= d_{12}\rho_3, \\ d_{13}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_1}^2 + c_{13}\rho_1\right), \\ d_{23}\left({\rho_1}^2 + {\rho_2}^2 + c_{12}\rho_1\rho_2\right) &= d_{12}\left(1 + {\rho_2}^2 + c_{23}\rho_2\right). \end{aligned}$$

2. Then, recover depths by solving a degree-2 problem,

$$\begin{aligned} \lambda_1 &= \rho_1\lambda_3, \\ \lambda_2 &= \rho_2\lambda_3, \\ {\lambda_3}^2 &= \rho_3. \end{aligned}$$

3. Finally, recover $(\mathbf{R} \mid \mathbf{t})$ from depths by solving a degree-1 problem.

Some natural next questions:

- Do other minimal problems have such decompositions?
- If yes, how to find / analyze them?
- Do problems have *optimal* decompositions?

All of these questions can be addressed by computing the problem's *Galois group*.

Galois Groups

Let $\pi : X \dashrightarrow Z$ be a *minimal problem*—by this, I mean a rational, dominant map between irreducible complex algebraic varieties of the same dimension.

Let $\mathbb{C}(X)$ and $\mathbb{C}(Z)$ denote the fields of rational functions on X and Z, respectively.

The algebraic extension $\mathbb{C}(X)/\mathbb{C}(Z)$ is finite of degree $d > 0$, which equals the generic fiber size or *degree* of π (ie. $\deg(\pi) = d$.)

Let $\overline{\mathbb{C}(X)}$ denote a *normal closure* of the extension $\mathbb{C}(X)/\mathbb{C}(Z)$.

Definition: $\mathrm{Gal}(\pi) = \mathrm{Gal}\left(\overline{\mathbb{C}(X)}/\mathbb{C}(Z)\right)$ is the Galois group of π.

Galois Groups

Let $\pi : X \dashrightarrow Z$ be a *minimal problem*—by this, I mean a rational, dominant map between irreducible complex algebraic varieties of the same dimension.

Let $\mathbb{C}(X)$ and $\mathbb{C}(Z)$ denote the fields of rational functions on X and Z, respectively.

The algebraic extension $\mathbb{C}(X)/\mathbb{C}(Z)$ is finite of degree $d > 0$, which equals the generic fiber size or *degree* of π (ie. $\deg(\pi) = d$.)

Let $\overline{\mathbb{C}(X)}$ denote a *normal closure* of the extension $\mathbb{C}(X)/\mathbb{C}(Z)$.

Definition: $\mathrm{Gal}(\pi) = \mathrm{Gal}\left(\overline{\mathbb{C}(X)}/\mathbb{C}(Z)\right)$ is the Galois group of π.

If you don't like this definition, you can instead think about *monodromy*.

Monodromy Groups

Let $\pi : X \dashrightarrow Z$ be a *minimal problem* of degree d.

There exists a dense, Zariski open $U \subset Z$ over which the *restricted map*

$$\pi\big|_{\pi^{-1}(U)} : \pi^{-1}(U) \to U$$

is a d-sheeted covering of complex manifolds.

For any $p \in U$, the *monodromy group* $\mathrm{Mon}(\pi, U; p)$ acts on the fiber $\pi^{-1}(p)$.

Monodromy Groups

Let $\pi : X \dashrightarrow Z$ be a *minimal problem* of degree d.

There exists a dense, Zariski open $U \subset Z$ over which the *restricted map*

$$\pi\Big|_{\pi^{-1}(U)} : \pi^{-1}(U) \to U$$

is a d-sheeted covering of complex manifolds.

For any $p \in U$, the *monodromy group* $\mathrm{Mon}(\pi, U; p)$ acts on the fiber $\pi^{-1}(p)$.

Example: Recall that P3P is a minimal problem of degree 8.

Let us describe the monodromy action...

Consider a 1-parameter family of systems

$$H(\lambda; t) = \begin{pmatrix} {\lambda_1}^2 + {\lambda_2}^2 + c_{12}(t)\lambda_1\lambda_2 - d_{12}(t) \\ {\lambda_1}^2 + {\lambda_3}^2 + c_{13}(t)\lambda_1\lambda_3 - d_{13}(t) \\ {\lambda_2}^2 + {\lambda_3}^2 + c_{23}(t)\lambda_2\lambda_3 - d_{23}(t) \end{pmatrix}.$$

where $p = (c_{12}(0), \ldots, d_{23}(0)) = (c_{12}(1), \ldots, d_{23}(1)) \in U \subset \mathbb{C}^3 \times \mathbb{C}^3$.

Suppose we have a solution $\lambda_0 \in \mathbb{C}^3$ at $t = 0$, ie. $H(\lambda_0; 0) = 0$.
The implicit function theorem constructs a *local solution function*,

$$\lambda(t) = \lambda_0 - \int_0^t \left(\frac{\partial H}{\partial \lambda}(\lambda(s); s) \right)^{-1} \cdot \frac{\partial H}{\partial t}(\lambda(s); s)\, ds$$

In general, $\lambda(0) \neq \lambda(1)$—because more than one solution exists!

The monodromy group $\mathrm{Mon}(\pi, U; p)$ consists of all permutations on the solution set $\pi^{-1}(p)$ that send $\lambda(0)$ to $\lambda(1)$ for each local solution function λ.

P3P monodromy

If $\lambda(t)$ is a local solution function, so is $\lambda'(t) = -\lambda(t)$. Since

$$\lambda(t) = -\lambda'(t) \quad \forall t \in [0,1],$$

the monodromy group is not the full-symmetric group S_8. Any monodromy permutation must preserve a nontrivial partition of the solution set

$$\pi^{-1}(p) = \{\lambda_1, -\lambda_1\} \sqcup \{\lambda_2, -\lambda_2\} \sqcup \{\lambda_3, -\lambda_3\} \sqcup \{\lambda_4, -\lambda_4\}.$$

The group of all such permutations is the *wreath product* $S_2 \wr S_4$. Thus,

$$\mathrm{Mon}(\pi, U; p) \subset S_2 \wr S_4$$

Surprisingly, this containment is strict!

Calibrated Resectioning with p Points and l Lines

$$\pi_{p,l} : (\mathbb{P}^3)^{\times p} \times (\mathrm{Gr}(\mathbb{P}^1, \mathbb{P}^3))^{\times l} \times \mathrm{SE}_3 \dashrightarrow (\mathbb{P}^3 \times \mathbb{P}^2)^{\times p} \times (\mathrm{Gr}(\mathbb{P}^1, \mathbb{P}^3) \times \mathrm{Gr}(\mathbb{P}^1, \mathbb{P}^2))^{\times l}$$

$$(\mathbf{q}_1, \ldots, \mathbf{q}_i, \ell_1, \ldots \ell_l, (\mathbf{R}\ \ \mathbf{t})) \mapsto (\mathbf{q}_i, (\mathbf{R}\ \ \mathbf{t})\mathbf{q}_i, 1 \le i \le 3,\ \ell_j, \wedge^2(\mathbf{R}\ \ \mathbf{t})\ell_j, 1 \le j \le l)$$

We get a branched cover when $p + l = 3$.
Previously **(D., Korotynskiy, Pajdla, Regan, SIAM J. Appl. Alg. Geom., 2023)**, we numerically computed the following table of Galois/monodromy groups:

Problem	p	l	$\deg(\pi_{p,l})$	$\mathrm{Gal}(\pi_{p,l})$	$\mathrm{Deck}(\pi_{p,l})$
P3P	3	0	8	$\mathbf{S_2 \wr S_4 \cap A_8}$	$\mathbf{S_2}$
P2P1L	2	1	4	$\mathbf{S_2 \wr S_2}$	$\mathbf{S_2}$
P1P2L	1	2	8	$\mathbf{S_2 \wr S_4 \cap A_8}$	$\mathbf{S_2}$
P3L	0	3	8	$\mathbf{S_8}$	**trivial**

Resectioning with Points and Lines (cont.)

Prior work (Kukelova et al., CVPR '16), (Ramalingam et al., ICRA '11) proposed degree-4 / degree-8 solvers for P2P1L / P1P2L. *Although we do not theoretically prove that our solutions are of the lowest possible degrees, we believe...*

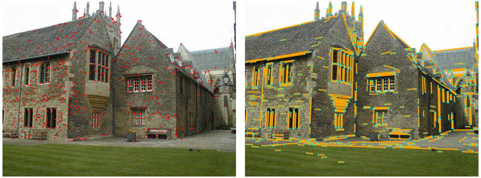

Can the theoretical improvements suggested by Galois groups also be made practical? **Yes—(D. Hruby, Pollefeys, CVPR 2024, arXiv 2404.16552).**

Method	Avg.	Min	Max
P2P1L Ours	**314**	**231**	**3061**
P2P1L Kuk.	1861	1439	10102
P2P1L Ram.	8898	5805	49984
P1P2L Ours	**504**	**364**	**4554**
P1P2L Kuk.	1967	1484	12931

Table: Solver timings in nanoseconds

Method	Avg. $\mathbf{R}_{err}$	Avg. $\mathbf{t}_{rel}$
P2P1L Ours	**5.3e-12**	**3.7e-10**
P2P1L Kuk.	2.8e-05	2.0e-05
P2P1L Ram.	4.7e-07	2.3e-05
PP1P2L Ours	**1.2e-07**	**2.0e-06**
P1P2L Kuk.	3.3e-05	3.4e-05

Table: Average solver errors ($\mathbf{R}_{err}$ in radians.)

A non-toy problem

A *radial camera* is a surjective projective linear map $A : \mathbb{P}^3 \dashrightarrow \mathbb{P}^1$.

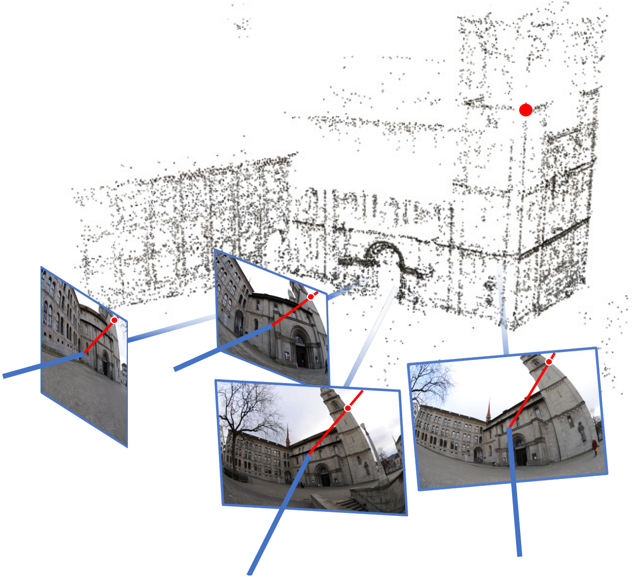

Intuition: The usual pinhole camera $\mathbb{P}^3 \dashrightarrow \mathbb{P}^2$ does not account for lens distortions. Usually, the distortion is radially-symmetric.

However, the space of radial lines through the image center forms a $\mathbb{P}^1$, and each radial line is invariant under distortion.

Goal: recover unknown 3D scene and cameras from matched 1D projections.

To obtain a metrically accurate reconstruction, we need 4 radial cameras $A_1, \ldots, A_4 \in \mathbb{P}(\mathbb{R}^{2\times 4})$, and to assume they are *calibrated*, which means

$$A_i = \begin{bmatrix} \mathbf{r}_{i1}^T & t_{i,1} \\ \mathbf{r}_{i2}^T & t_{i,2} \end{bmatrix} \quad \text{where} \quad \|\mathbf{r}_{i1}\| = \|\mathbf{r}_{i2}\| = 1,\ \mathbf{r}_{i1}^T \mathbf{r}_{i2} = 0.$$

Up to similarity transformation in $\mathbb{R}^3$, we may assume

$$A_1 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}, \quad A_2 = \begin{bmatrix} \mathbf{r}_{21}^T & 0 \\ \mathbf{r}_{22}^T & 1 \end{bmatrix}.$$

This leaves $0+3+5+5=13$ unknowns, and a minimal problem: given matches $\mathbf{p}_{i1}, \mathbf{p}_{i2}, \mathbf{p}_{i3}, \mathbf{p}_{i4} \in \mathbb{P}^1$, find $\mathbf{q}_1, \ldots, \mathbf{q}_{13} \in \mathbb{P}^3$ and A_i as above such that

$$A_i \mathbf{q}_j \sim \mathbf{p}_{ij} \quad \forall\, 1 \le i \le 4,\ 1 \le j \le 13.$$

The number of complex solutions is 3584.

Can we do better?

We can do better!

(Hruby, Korotynskiy, D., et. al, CVPR '23) reconstructs radial cameras by decomposing into subproblems of degree at most **28**.

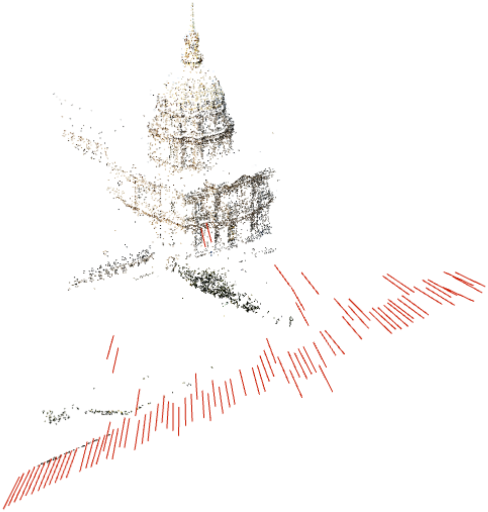

The number **28** is the **Galois width** of a finite group / branched cover naturally associated to this reconstruction problem.

I will explain what this means.

Galois width

Here is a simple model for an exact, algebraic algorithm:

1. Initialize $\mathbb{F}_0 \leftarrow \mathbb{Q}$

Galois width

Here is a simple model for an exact, algebraic algorithm:

1. Initialize $\mathbb{F}_0 \leftarrow \mathbb{Q}$
2. For $i = 1, \ldots, m$, either
   (i) do arithmetic in $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}$, OR

Galois width

Here is a simple model for an exact, algebraic algorithm:

1. Initialize $\mathbb{F}_0 \leftarrow \mathbb{Q}$
2. For $i = 1, \ldots, m$, either
   (i) do arithmetic in $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}$, OR
   (ii) compute a root of a polynomial: formally, extend the working field $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}(\beta)$

Galois width

Here is a simple model for an exact, algebraic algorithm:

1. Initialize $\mathbb{F}_0 \leftarrow \mathbb{Q}$
2. For $i = 1, \ldots, m$, either
   (i) do arithmetic in $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}$, OR
   (ii) compute a root of a polynomial: formally, extend the working field $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}(\beta)$
3. Output: $\alpha \in \mathbb{F}_m$

Galois width

Here is a simple model for an exact, algebraic algorithm:

1. Initialize $\mathbb{F}_0 \leftarrow \mathbb{Q}$
2. For $i = 1, \ldots, m$, either
   (i) do arithmetic in $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}$, OR
   (ii) compute a root of a polynomial: formally, extend the working field $\mathbb{F}_i \leftarrow \mathbb{F}_{i-1}(\beta)$
3. Output: $\alpha \in \mathbb{F}_m$

Definition

An *algorithm* computing an algebraic number $\alpha \in \overline{\mathbb{Q}}$ is a finite tower of fields

$$\mathbb{Q} = \mathbb{F}_0 \subset \mathbb{F}_1 \subset \cdots \subset \mathbb{F}_m \ni \alpha.$$

Definition

The *Galois width* of an algebraic number $\alpha \in \overline{\mathbb{Q}}$ is the quantity

$$\mathrm{gw}(\alpha) = \min_{\substack{\text{algorithms} \\ \mathbb{Q}=\mathbb{F}_0\subset\cdots\subset\mathbb{F}_m\ni\alpha}} \left(\max_{0\le i<m} [\mathbb{F}_{i+1} : \mathbb{F}_i] \right).$$

Definition

The *Galois width* of an algebraic number $\alpha \in \overline{\mathbb{Q}}$ is the quantity

$$\mathrm{gw}(\alpha) = \min_{\substack{\text{algorithms} \\ \mathbb{Q}=\mathbb{F}_0 \subset \cdots \subset \mathbb{F}_m \ni \alpha}} \left(\max_{0 \le i < m} [\mathbb{F}_{i+1} : \mathbb{F}_i] \right).$$

There is a parallel notion in the world of finite groups.

Definition

The *Galois width* of a finite group G is the quantity

$$\mathrm{gw}(G) = \min_{\substack{\text{subgroup chains} \\ id=H_m \le \ldots \le H_0 = G}} \left(\max_{0 \le i < m} [H_i : H_{i+1}] \right).$$

Theorem (D 25)

For any algebraic number $\alpha \in \overline{\mathbb{Q}}$ with minimal polynomial $p(x) \in \mathbb{Q}[x]$ and Galois group $G = Gal(p(x))$, we have $\mathrm{gw}(\alpha) = \mathrm{gw}(G)$.

Theorem (Properties of the Galois width (D '25))

Let G be any finite group.

1. $\mathrm{gw}(H) \leq \mathrm{gw}(G)$ *for any subgroup* $H \leq G$.
2. $\mathrm{gw}(G) = \max(\mathrm{gw}(N), \mathrm{gw}(G/N))$ *for any normal subgroup* $N \trianglelefteq G$.
3. *For any composition series* $id = N_m \trianglelefteq N_{m-1} \trianglelefteq \cdots \trianglelefteq N_0 = G$, *we have*
$$\mathrm{gw}(G) = \max_{0 \leq i < m} \mathrm{gw}\left(N_i / N_{i+1}\right).$$
4. *If G is simple, then*
$$\mathrm{gw}(G) = \min_{H < G} [G : H].$$
5. *For any prime* p, *we have* $\mathrm{gw}(\mathbb{Z}/p\mathbb{Z}) = p$.
6. *For any* $n \geq 1$, *we have* $\mathrm{gw}(S_n) = \mathrm{gw}(A_n) = \begin{cases} 3 & \text{if } n = 4, \\ n & \text{else.} \end{cases}$.