

Using Algebraic Geometry for Computer Vision

Joe Kileel
UC Berkeley

CMS @ Caltech
October 27, 2016

Algebraic vision

Multiview geometry studies 3D scene reconstruction from images. Foundations in projective geometry. *Algebraic vision* bridges to algebraic geometry (combinatorial, computational, numerical, ...).



Oct 8–9, 2015, Berlin



May 2–6, 2016, San Jose

Algebraic vision

Multiview geometry studies 3D scene reconstruction from images. Foundations in projective geometry. *Algebraic vision* bridges to algebraic geometry (combinatorial, computational, numerical, ...).



Oct 8–9, 2015, Berlin



May 2–6, 2016, San Jose

Today, I want to tell you about a numerical algebraic geometry project that I have been working on this summer. I welcome your feedback.

First the background: 3D reconstruction

Example from *Building Rome in a Day* (2009) by S. Agarwal et al.

Input: 2106 Flickr images tagged “Colosseum”

First the background: 3D reconstruction

Example from *Building Rome in a Day* (2009) by S. Agarwal et al.

Input: 2106 Flickr images tagged “Colosseum”

Output: configuration of cameras and 819,242 3D points




Figure: 3D model of the Colosseum in Rome from 2106 Flickr images

How does Google do it?

- ▶ Identify pairs or triplets of images that overlap.
- ▶ Do robust reconstruction with pairs or triplets of images.
- ▶ Piece together.

Google solves polynomial systems

- ▶ Identify pairs or triplets of images that overlap.
- ▶ Do robust reconstruction with pairs or triplets of images.
- ▶ Piece together.

For mathematicians: ‘Tiny’ reconstructions are subroutines in large-scale reconstructions. The ‘tiny’ reconstructions rely on super-fast, specialized **polynomial** equation solvers.

[Fischler-Bolles: *Random Sample Consensus: a Paradigm for Model Fitting with Application to Image Analysis and Automated Cartography*, 1981]

[Kúkelová-Bujnak-Pajdla: *Automatic Generator of Minimal Problem Solvers*, 2008]

What is a camera?

A **camera** is a full rank 3×4 real matrix A (up to scale).

Determines a projection $\mathbb{P}^3 \dashrightarrow \mathbb{P}^2; X \mapsto AX$.

What is a camera?

A **camera** is a full rank 3×4 real matrix A (up to scale).

Determines a projection $\mathbb{P}^3 \dashrightarrow \mathbb{P}^2; X \mapsto AX$.

Math	Interpretation
\mathbb{P}^3	world
\mathbb{P}^2	image plane
$\ker(A)$	camera center
K	internal parameters (e.g. focal length)
$[R t]$	external parameters (orientation, center)

Above $A_{3 \times 4} =: K_{3 \times 3} [R_{3 \times 3} | t_{3 \times 1}]$ where K is upper triangular and R is a rotation. If $K = I$, so $A = [R | t]$, then A is **calibrated**.

RQ factorization

What is a camera configuration?

Images alone do not determine absolute position of cameras.



Definition

A **configuration** of n calibrated cameras is an orbit of the group $\mathcal{G} \subset GL(4)$ of appropriate changes of world coordinates acting on:

$$\{(A_1, \dots, A_n) : A_i \text{ is a calibrated camera}\}$$

via simultaneous right multiplication. Here \mathcal{G} consists of composites translations, rotations, central dilations:

$$\mathcal{G} := \{g \in \mathbb{C}^{4 \times 4} \mid (g_{ij})_{1 \leq i, j \leq 3} \in SO(3, \mathbb{C}), g_{41} = g_{42} = g_{43} = 0 \text{ and } g_{44} \neq 0\}.$$

Warmup: Two Calibrated Views

$$A = \begin{bmatrix} \star & \star & \star & \star \\ \star & \star & \star & \star \\ \star & \star & \star & \star \end{bmatrix} \quad B = \begin{bmatrix} \diamond & \diamond & \diamond & \diamond \\ \diamond & \diamond & \diamond & \diamond \\ \diamond & \diamond & \diamond & \diamond \end{bmatrix}$$

The image of $(A, B) : \mathbb{P}^3 \dashrightarrow \mathbb{P}^2 \times \mathbb{P}^2$ is the hypersurface defined by:

$$f(x, x') = \det \begin{bmatrix} \star & \star & \star & \star & x_1 & 0 \\ \star & \star & \star & \star & x_2 & 0 \\ \star & \star & \star & \star & x_3 & 0 \\ \diamond & \diamond & \diamond & \diamond & 0 & x'_1 \\ \diamond & \diamond & \diamond & \diamond & 0 & x'_2 \\ \diamond & \diamond & \diamond & \diamond & 0 & x'_3 \end{bmatrix} \quad (\text{multiview variety})$$

This polynomial is bilinear:

$$f(x, x') = \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix} \cdot \begin{bmatrix} \square & \square & \square \\ \square & \square & \square \\ \square & \square & \square \end{bmatrix} \cdot \begin{bmatrix} x'_1 \\ x'_2 \\ x'_3 \end{bmatrix}$$

This 3×3 -matrix is the *essential matrix* of the two cameras.

It is determined by the configuration of (A, B) .

What is its SVD?

How to write \square in terms of \star and \diamond ?

Warmup cont'd: Nister's 5 point algorithm

- ▶ The set of all essential matrices forms a variety $\mathcal{E} \subset \mathbb{P}(\mathbb{C}^{3 \times 3})$. \mathcal{E} has **dimension 5** and **degree 10**.
- ▶ Using Gröbner bases, D. Nister built an efficient solver that recovers **10** essential matrices from **5** *image point pairs*.
- ▶ Given a pair $(x, x') \in \mathbb{P}^2 \times \mathbb{P}^2$ of points in the first and second images that are pictures of the same world point. Then the essential matrix E for the views must satisfy $x^T E x' = 0$. Nister intersects 5 of these hyperplanes with \mathcal{E} . **Minimal problem**
- ▶ From essential matrix, camera configuration is easy to get.
- ▶ Nister's solver is used **alot** for RANSAC *3D* reconstruction.

[D. Nistér: *An efficient solution to the five-point relative pose problem*, 2004]

Warmup cont'd: Nister's 5 point algorithm


- ▶ The set of all essential matrices forms a variety $\mathcal{E} \subset \mathbb{P}(\mathbb{C}^{3 \times 3})$. \mathcal{E} has **dimension 5** and **degree 10**.
- ▶ Using Gröbner bases, D. Nister built an efficient solver that recovers **10** essential matrices from **5 image point pairs**.
- ▶ Given a pair $(x, x') \in \mathbb{P}^2 \times \mathbb{P}^2$ of points in the first and second images that are pictures of the same world point. Then the essential matrix E for the views must satisfy $x^T E x' = 0$. Nister intersects 5 of these hyperplanes with \mathcal{E} . **Minimal problem**
- ▶ From essential matrix, camera configuration is easy to get.
- ▶ Nister's solver is used **alot** for RANSAC 3D reconstruction.

[D. Nistér: *An efficient solution to the five-point relative pose problem*, 2004]

- ▶ Effort for 3 calibrated view solver so far has proven elusive.

Minimal Problems for the Calibrated Trifocal Variety

- ▶ **Determines:** algebraic degree for various parametrized polynomial systems, for recovery of configurations of **three** calibrated cameras
- ▶ **Method:** take special linear sections of fixed $\mathcal{T}_{\text{cal}} \subset \mathbb{P}(\mathbb{C}^{3 \times 3 \times 3})$



- ▶ **Relies on:** numerical algebraic geometry software

Point/line correspondences

- ▶ $3D$ reconstruction uses points/lines in the photos that match,
- ▶ Call elements of \mathbb{P}^2 **image points**, and elements of $(\mathbb{P}^2)^\vee$ **image lines**.
- ▶ An element of $(\mathbb{P}^2 \sqcup (\mathbb{P}^2)^\vee)^{\times 3}$ is a **point/line image correspondence**.
- ▶ E.g., an element of $\mathbb{P}^2 \times \mathbb{P}^2 \times (\mathbb{P}^2)^\vee$ is called a point-point-line image correspondence, denoted PPL .

Definition

A calibrated camera configuration (A, B, C) is **consistent** with a given point/line image correspondence if there exist a point in \mathbb{P}^3 and a line in \mathbb{P}^3 containing that point such that (A, B, C) respectively map these to the given points and lines in \mathbb{P}^2 .

Example

A configuration (A, B, C) is consistent with a given PPL image correspondence $(x, x', \ell'') \in \mathbb{P}^2 \times \mathbb{P}^2 \times (\mathbb{P}^2)^\vee$ if there exist $(X, L) \in \mathbb{P}^3 \times \text{Gr}(\mathbb{P}^1, \mathbb{P}^3)$ with $X \in L$ such that $AX = x$, $BX = x'$, and $CL = \ell''$. In particular, this implies that $X \notin \ker(A)$, $\ker(B)$ and $\ker(C) \notin L$.

Say configuration (A, B, C) is consistent with a set of point/line correspondences if it is consistent with each correspondence.

Main result

Theorem (K)

The rows of the following table display the algebraic degree for 66 minimal problems across three calibrated views. Given generic point/line image correspondences in the amount specified by the entries in the first five columns, then the number of calibrated camera configurations over \mathbb{C} that are consistent with those correspondences equals the entry in the sixth column.

#PPP	#PPL	#PLP	#LLL	#PLL	#configurations
3	1	0	0	0	272
3	0	0	1	0	216
3	0	0	0	2	448
2	2	0	0	1	424
2	1	1	0	1	528
2	1	0	1	1	424
2	1	0	0	3	736
2	0	0	2	1	304
2	0	0	1	3	648
2	0	0	0	5	1072
1	4	0	0	0	160
1	3	1	0	0	520
1	3	0	1	0	360
1	3	0	0	2	520
1	2	2	0	0	672

Main result (cont'd)

#PPP	#PPL	#PLP	#LLL	#PLL	#configurations
1	2	1	1	0	552
1	2	1	0	2	912
1	2	0	2	0	408
1	2	0	1	2	704
1	2	0	0	4	1040
1	1	1	2	0	496
1	1	1	1	2	896
1	1	1	0	4	1344
1	1	0	3	0	368
1	1	0	2	2	736
1	1	0	1	4	1184
1	1	0	0	6	1672
1	0	0	4	0	360
1	0	0	3	2	696
1	0	0	2	4	1176
1	0	0	1	6	1680
1	0	0	0	8	2272
0	5	0	0	1	160
0	4	1	0	1	616
0	4	0	1	1	456
0	4	0	0	3	616
0	3	2	0	1	1152
0	3	1	1	1	880
0	3	1	0	3	1280
0	3	0	2	1	672
0	3	0	1	3	1008
0	3	0	0	5	1408

Main result (cont'd again)

#PPP	#PPL	#PLP	#LLL	#PLL	#configurations
0	2	2	1	1	1168
0	2	2	0	3	1680
0	2	1	2	1	1032
0	2	1	1	3	1520
0	2	1	0	5	2072
0	2	0	3	1	800
0	2	0	2	3	1296
0	2	0	1	5	1848
0	2	0	0	7	2464
0	1	1	3	1	1016
0	1	1	2	3	1552
0	1	1	1	5	2144
0	1	1	0	7	2800
0	1	0	4	1	912
0	1	0	3	3	1456
0	1	0	2	5	2088
0	1	0	1	7	2808
0	1	0	0	9	3592
0	0	0	5	1	920
0	0	0	4	3	1464
0	0	0	3	5	2176
0	0	0	2	7	3024
0	0	0	1	9	3936
0	0	0	0	11	4912

Some clarifying comments

Remark

A calibrated camera configuration (A, B, C) has 11 degrees of freedom, and the first five columns in the table above represent conditions of codimension 3, 2, 2, 2, 1, respectively.

Remark

*The proof technique relies on trifocal tensors. These break the symmetry between the three views. The numbers reported above are the **true, intrinsic** degrees for those 66 cases, **based on the underlying camera geometry**. However, using correspondences of type LPP,LPL,LLP, there are other minimal problems with degrees that are not covered by the result and proof technique.*

Remark

*Roughly speaking, these algebraic degrees are **complexity measures**.*

Example: '1PPP + 4PPL' has degree 160

Given the following set of real, random correspondences:

$$\begin{array}{l} \text{PPP} : \begin{bmatrix} 0.6132 \\ 0.8549 \\ 0.5979 \end{bmatrix}, \begin{bmatrix} 0.4599 \\ 0.5713 \\ 0.1812 \end{bmatrix}, \begin{bmatrix} 0.6863 \\ 0.4508 \\ 0.1834 \end{bmatrix} \\ \text{PPL} : \begin{bmatrix} 0.4970 \\ 0.6532 \\ 0.8429 \end{bmatrix}, \begin{bmatrix} 0.5405 \\ 0.8342 \\ 0.6734 \end{bmatrix}, \begin{bmatrix} 0.2692 \\ 0.8861 \\ 0.1333 \end{bmatrix} \\ \text{PPL} : \begin{bmatrix} 0.8933 \\ 0.3375 \\ 0.1054 \end{bmatrix}, \begin{bmatrix} 0.7062 \\ 0.6669 \\ 0.7141 \end{bmatrix}, \begin{bmatrix} 0.3328 \\ 0.8228 \\ 0.6781 \end{bmatrix} . \end{array} \quad \begin{array}{l} \text{PPL} : \begin{bmatrix} 0.6251 \\ 0.9248 \\ 0.9849 \end{bmatrix}, \begin{bmatrix} 0.3232 \\ 0.5453 \\ 0.6941 \end{bmatrix}, \begin{bmatrix} 0.3646 \\ 0.1497 \\ 0.1364 \end{bmatrix} \\ \text{PPL} : \begin{bmatrix} 0.2896 \\ 0.6909 \\ 0.4914 \end{bmatrix}, \begin{bmatrix} 0.6898 \\ 0.9855 \\ 0.6777 \end{bmatrix}, \begin{bmatrix} 0.6519 \\ 0.8469 \\ 0.6855 \end{bmatrix} \end{array}$$

This is a generic instance of the minimal problem '1PPP + 4PPL'. Up to the action of \mathcal{G} , there are only a positive finite number of three calibrated cameras that are exactly consistent with this image data, namely 160 complex configurations. For this instance, it turns out that 18 of those configurations are real. For example, one is:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, B = \begin{bmatrix} -0.22 & 0.95 & -0.18 & 1 \\ 0.96 & 0.24 & 0.08 & 1.44 \\ -0.12 & 0.15 & 0.97 & 0.97 \end{bmatrix}, C = \begin{bmatrix} 0.17 & 0.94 & -0.28 & 1.41 \\ -0.95 & 0.22 & 0.18 & -0.13 \\ -0.24 & -0.23 & -0.94 & -1.16 \end{bmatrix}.$$

Multi-view varieties

These give tight equational formulations for point/line correspondences and cameras to be consistent.

Definition

Fix cameras A, B, C corresponding to linear projections $\alpha, \beta, \gamma : \mathbb{P}^3 \dashrightarrow \mathbb{P}^2$. Set $\mathcal{F}l_{0,1} = \{(X, L) \in \mathbb{P}^3 \times \text{Gr}(\mathbb{P}^1, \mathbb{P}^3) \mid X \in L\}$.

- ▶ **PLL multi-view variety** denoted $X_{A,B,C}^{PLL}$ is the closure of the image of $\mathcal{F}l_{0,1} \dashrightarrow \mathbb{P}_A^2 \times (\mathbb{P}_B^2)^\vee \times (\mathbb{P}_C^2)^\vee$, $(X, L) \mapsto (\alpha(X), \beta(L), \gamma(L))$
- ▶ **LLL multi-view variety** denoted $X_{A,B,C}^{LLL}$ is the closure of the image of $\text{Gr}(\mathbb{P}^1, \mathbb{P}^3) \dashrightarrow (\mathbb{P}_A^2)^\vee \times (\mathbb{P}_B^2)^\vee \times (\mathbb{P}_C^2)^\vee$, $L \mapsto (\alpha(L), \beta(L), \gamma(L))$
- ▶ **PPL multi-view variety** denoted $X_{A,B,C}^{PPL}$ is the closure of the image of $\mathcal{F}l_{0,1} \dashrightarrow \mathbb{P}_A^2 \times \mathbb{P}_B^2 \times (\mathbb{P}_C^2)^\vee$, $(X, L) \mapsto (\alpha(X), \beta(X), \gamma(L))$
- ▶ **PLP multi-view variety** denoted $X_{A,B,C}^{PLP}$ is the closure of the image of $\mathcal{F}l_{0,1} \dashrightarrow \mathbb{P}_A^2 \times (\mathbb{P}_B^2)^\vee \times \mathbb{P}_C^2$, $(X, L) \mapsto (\alpha(X), \beta(L), \gamma(X))$
- ▶ **PPP multi-view variety** denoted $X_{A,B,C}^{PPP}$ is the closure of the image of $\mathbb{P}^3 \dashrightarrow \mathbb{P}_A^2 \times \mathbb{P}_B^2 \times \mathbb{P}_C^2$, $X \mapsto (\alpha(X), \beta(X), \gamma(X))$.

Equations of multi-view varieties

Theorem (Aholt-Sturmfels-Thomas, K)

▶ $\dim(X_{A,B,C}^{PLL}) = 5$ and $I(X_{A,B,C}^{PLL}) = \langle T_{A,B,C}(x, \ell', \ell'') \rangle \subset \mathbb{C}[x_i, \ell'_j, \ell''_k]$

▶ $\dim(X_{A,B,C}^{LLL}) = 4$ and $I(X_{A,B,C}^{LLL}) \subset \mathbb{C}[\ell_i, \ell'_j, \ell''_k]$ is generated by the maximal minors of the matrix $(A^T \ell \quad B^T \ell' \quad C^T \ell'')_{4 \times 3}$

▶ $\dim(X_{A,B,C}^{PPL}) = 4$ and $I(X_{A,B,C}^{PPL}) \subset \mathbb{C}[x_i, x'_j, \ell''_k]$ is generated by the maximal minors of the matrix $\begin{pmatrix} A & x & 0 \\ B & 0 & x' \\ \ell''^T C & 0 & 0 \end{pmatrix}_{7 \times 6}$

▶ $\dim(X_{A,B,C}^{PLP}) = 4$ and $I(X_{A,B,C}^{PLP}) \subset \mathbb{C}[x_i, \ell'_j, x''_k]$ is generated by the maximal minors of the matrix $\begin{pmatrix} A & x & 0 \\ C & 0 & x'' \\ \ell'^T B & 0 & 0 \end{pmatrix}_{7 \times 6}$

▶ $\dim(X_{A,B,C}^{PPP}) = 3$ and $I(X_{A,B,C}^{PPP}) \subset \mathbb{C}[x_i, x'_j, x''_k]$ is generated by the maximal minors of the matrix $\begin{pmatrix} A & x & 0 & 0 \\ B & 0 & x' & 0 \\ C & 0 & 0 & x'' \end{pmatrix}_{9 \times 7}$ together with $\det \begin{pmatrix} A & x & 0 \\ B & 0 & x' \end{pmatrix}_{6 \times 6}$ and $\det \begin{pmatrix} A & x & 0 \\ C & 0 & x'' \end{pmatrix}_{6 \times 6}$ and $\det \begin{pmatrix} B & x' & 0 \\ C & 0 & x'' \end{pmatrix}_{6 \times 6}$

Example cont'd: '1PPP + 4PPL'

- ▶ Multi-view equations above lead to a parametrized system of polynomial equations for each minimal problem in main result.
- ▶ Minimal problem '1PPP + 4PPL', the **unknowns** are the **36** entries of A, B, C , up to the action of \mathcal{G} . There are $\binom{9}{7} + 3 + 4 \cdot \binom{7}{6} = 67$ **quartic equations**. Coefficients parametrized cubically and quadratically by the image data in $(\mathbb{P}^2)^{11} \times ((\mathbb{P}^2)^\vee)^4$.
- ▶ Since parameter space is irreducible, to find the generic number of solutions to the system, we may specialize to **one** random instance, such as in earlier example.
- ▶ Nonetheless, solving a single instance of this system – 'as is' – is tough, let alone solving systems for the other minimal problems present in main result.
- ▶ Way out: nontrivially **replace** above systems with others, which enlarge the solution sets but amount to accessible computations. This is based on *trifocal tensors* from multi-view geometry.

Trifocal tensors

Definition

Let A, B, C be three calibrated cameras. Their **calibrated trifocal tensor** $T_{A,B,C} \in \mathbb{P}(\mathbb{C}^{3 \times 3 \times 3})$ is computed as follows:

- ▶ Form the 4×9 matrix $(A^T | B^T | C^T)$.
- ▶ Then for $1 \leq i, j, k \leq 3$, the entry $(T_{A,B,C})_{ijk}$ is $(-1)^{i+1}$ times the determinant of the 4×4 submatrix gotten by *omitting* the i^{th} column from A^T , while *keeping* the j^{th} and k^{th} columns from B^T and C^T resp.

Lemma ($T_{A,B,C}$ encodes *PLL* image correspondences)

$T_{A,B,C}(x, \ell', \ell'') := \sum_{1 \leq i, j, k \leq 3} T_{ijk} x_i \ell'_j \ell''_k = 0$ if and only if $\alpha^{-1}(x) \cap \beta^{-1}(\ell') \cap \gamma^{-1}(\ell'') \neq \emptyset$.

Remark (necessary conditions for *PPP, LLL, PLP, PPL*)

Necessary conditions for a *PPP, LLL, PLP, PPL* correspondence to be consistent with (A, B, C) are expressed via polynomials **linear** in $T_{A,B,C}$.

[Hartley-Zisserman: *Multiple View Geometry in Computer Vision*, 2003]

Proposition (Hartley)

Let A, B, C be cameras. Let $x \in \mathbb{P}_A^2, x' \in \mathbb{P}_B^2, x'' \in \mathbb{P}_C^2$ be *points* and $\ell \in (\mathbb{P}_A^2)^\vee, \ell' \in (\mathbb{P}_B^2)^\vee, \ell'' \in (\mathbb{P}_C^2)^\vee$ be *lines*. Putting $T = T_{A,B,C}$, then (A, B, C) is consistent with:

▶ (x, ℓ', ℓ'') only if $T(x, \ell', \ell'') = 0$ **[PLL]**

▶ (ℓ, ℓ', ℓ'') only if $[\ell]_\times T(-, \ell', \ell'') = 0$ **[LLL]**

▶ (x, ℓ', x'') only if $[x'']_\times T(x, \ell', -) = 0$ **[PLP]**

▶ (x, x', ℓ'') only if $[x']_\times T(x, -, \ell'') = 0$ **[PPL]**

▶ (x, x', x'') only if $[x'']_\times T(x, -, -)[x']_\times = 0$. **[PPP]**

In middle bullets, each contraction of T with two vectors gives a column vector in \mathbb{C}^3 . Last bullet: $T(x, -, -) = \sum_{i=1}^3 x_i (T_{ijk})_{1 \leq j, k \leq 3} \in \mathbb{C}^{3 \times 3}$.

$$\text{Above } [\ell]_\times = \begin{bmatrix} 0 & \ell_3 & -\ell_2 \\ -\ell_3 & 0 & \ell_1 \\ \ell_2 & -\ell_1 & 0 \end{bmatrix} \text{ etc.}$$

Calibrated Trifocal Variety

Definition

The **calibrated trifocal variety**, denoted $\mathcal{T}_{\text{cal}} \subset \mathbb{P}(\mathbb{C}^{3 \times 3 \times 3})$, is defined to be the Zariski closure of the image of the following rational map:

$$(\text{SO}(3, \mathbb{C}) \times \mathbb{C}^3) \times (\text{SO}(3, \mathbb{C}) \times \mathbb{C}^3) \times (\text{SO}(3, \mathbb{C}) \times \mathbb{C}^3) \dashrightarrow \mathbb{P}(\mathbb{C}^{3 \times 3 \times 3}),$$

$$\left((R_1, t_1), (R_2, t_2), (R_3, t_3) \right) \mapsto T_{[R_1|t_1], [R_2|t_2], [R_3|t_3]}$$

where the formula for T is from the previous slide. So, \mathcal{T}_{cal} is the closure of the set of all calibrated trifocal tensors. *Analogue of Nister's essential variety \mathcal{E}*

Lemma

The calibrated configuration (A, B, C) is equivalent to the tensor $T_{A,B,C}$.

Theorem (K)

*The calibrated trifocal variety $\mathcal{T}_{\text{cal}} \subset \mathbb{P}(\mathbb{C}^{3 \times 3 \times 3})$ is irreducible, **dimension 11** and **degree 4912**. It equals the $\text{SO}(3, \mathbb{C})^{\times 3}$ -orbit closure generated by the following projective plane, parametrized by $[\lambda_1 \ \lambda_2 \ \lambda_3]^T \in \mathbb{P}^2$:*

$$T_{1^{**}} = \begin{bmatrix} 0 & \lambda_1 & \lambda_2 \\ 0 & 0 & 0 \\ \lambda_1 & 0 & 0 \end{bmatrix}, \quad T_{2^{**}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \lambda_1 & \lambda_2 \\ 0 & \lambda_3 & 0 \end{bmatrix}, \quad T_{3^{**}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & \lambda_1 & \lambda_2 + \lambda_3 \end{bmatrix}.$$

Numerical algebraic geometry

- ▶ I obtain the table of degrees by a computational proof, using **homotopy continuation**.
- ▶ **General methodology:** solutions of a **start system** are **tracked** to solutions of a **target system** (RK4 and Newton's method)
- ▶ Already applied to: kinematics, biology, statistics, Schubert calculus . . .
- ▶ Here, it suffices to count number of configurations for **one** random instance per minimal problem.

[Sommese-Wampler: *The numerical solution of systems of polynomials arising in science and engineering*, 2005]

Numerical algebraic geometry

- ▶ I obtain the table of degrees by a computational proof, using **homotopy continuation**.
- ▶ **General methodology:** solutions of a **start system** are **tracked** to solutions of a **target system** (RK4 and Newton's method)
- ▶ Already applied to: kinematics, biology, statistics, Schubert calculus . . .
- ▶ Here, it suffices to count number of configurations for **one** random instance per minimal problem.

[Sommese-Wampler: *The numerical solution of systems of polynomials arising in science and engineering*, 2005]

Proof Sketch:

1. Key maneuver: use the necessary conditions for consistency that are linear in $T_{A,B,C}$. This is a relaxation of each minimal problem. Defines a linear section L_{special} of calibrated trifocal variety \mathcal{T}_{cal} . Move a generic linear section L_{general} of \mathcal{T}_{cal} to L_{special} , by tracking **4912** paths via homotopy continuation.
2. Decide which of the endpoints are indeed consistent, by comparing with multi-view equations. (Use SVD for robustness)

Future project: implement a fast solver to recover (A, B, C) , tracking optimal number of paths.

Open problem: can all of the solutions (A, B, C) be **REAL**?

Thank you!