Influence of numerical conditioning on the accuracy of relative orientation

Siniša Šegvić, Gerald Schweighofer and Axel Pinz

Vision-based Measurement Group,

Institute of Electrical Measurement and Measurement Signal Processing,

Graz University of Technology

Purpose

Effects of **numerical conditioning** in the **essential estimation** (calibrated, overconstrained, closed-form)

- analyse the eight-point alg. (`8pt`) forward bias
- discuss the conditioning of five-point alg. (`5pt`)
- validation by comprehensive performance evaluation

Benefits

Why I think this might be of interest to **you**:

- what causes the 8pt alg. forward bias?
- comparison of known conditioning approaches (8pt alg)
- conditioning the 5pt algorithm
- performance evaluation `5pt` vs `8pt` vs hg in the overconstrained case

AGENDA

- Introduction (short)
- Analysis of the `8pt` forward bias
- Review of the `8pt` conditioning (short)
- Conditioning the `5pt` algorithm
- Experimental validation
- Conclusion

INTRODUCTION

Context:

- re-estimating **relative orientation** on the set of inliers
- we can't solve directly for (R,t), use intermediate objects
- $\Rightarrow$ calibrated, overconstrained, closed-form **E**, **H**, ...

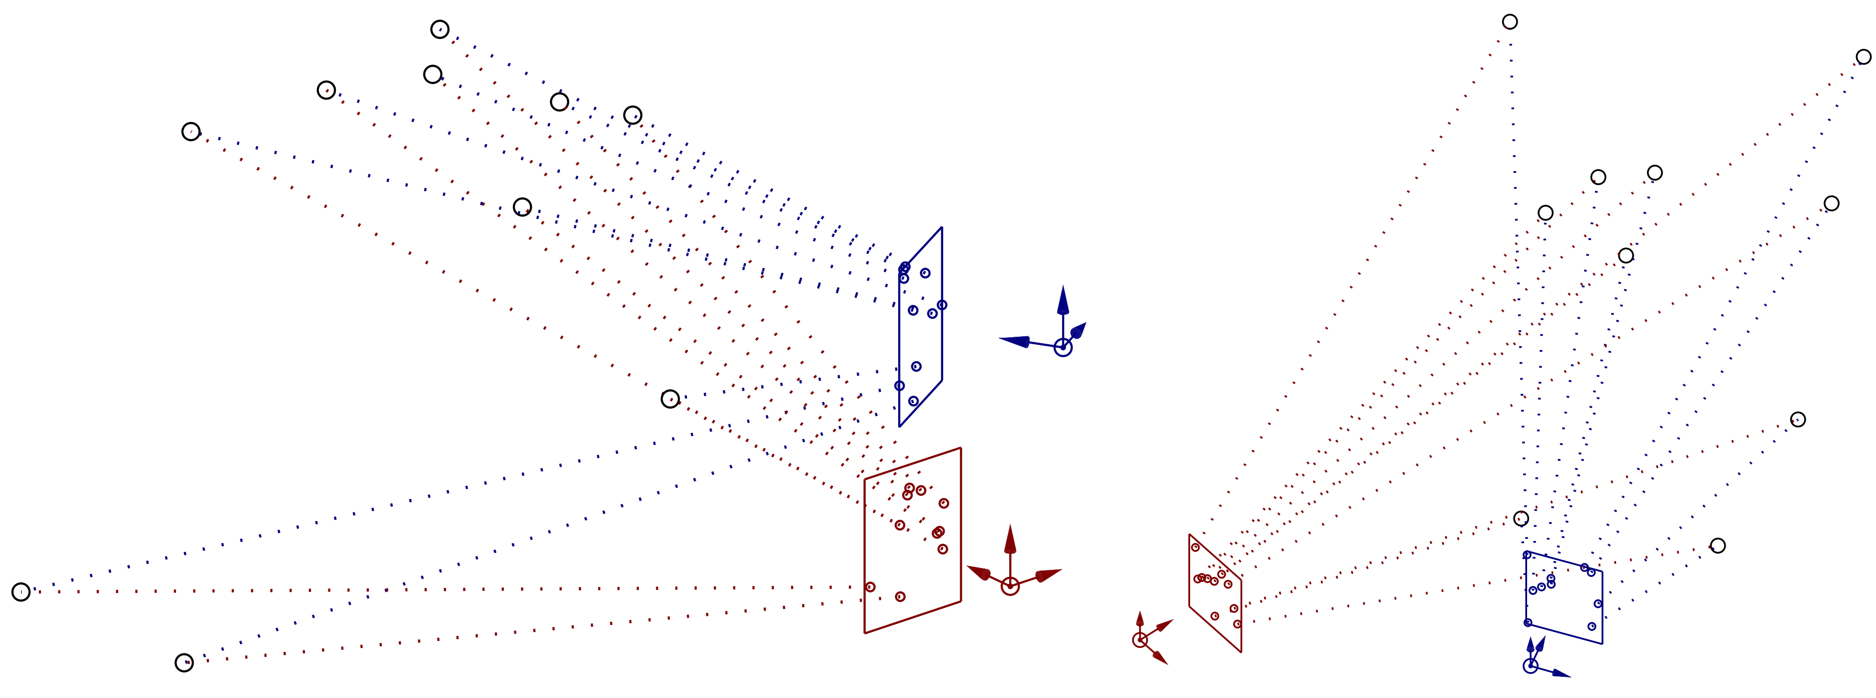

The essential matrix

The recovery approaches rely on **two** constraints:

- the epipolar constraint:

$$\mathbf{q}_{i\mathrm{B}}^{\top} \cdot \mathbf{E} \cdot \mathbf{q}_{i\mathrm{A}} = 0$$

- the calibrated (5DOF) constraint:

$$2 \cdot \mathbf{E}\mathbf{E}^{T}\mathbf{E} - trace(\mathbf{E}\mathbf{E}^{T})\mathbf{E} = 0 \qquad \text{(v1)}$$

THE ESSENTIAL MATRIX

The recovery approaches rely on **two** constraints:

- the epipolar constraint:

$$\mathbf{q}_{i\mathrm{B}}^{\top} \cdot \mathbf{E} \cdot \mathbf{q}_{i\mathrm{A}} = 0$$

- the calibrated (5DOF) constraint:

$$2 \cdot \mathbf{E}\mathbf{E}^{T}\mathbf{E} - trace(\mathbf{E}\mathbf{E}^{T})\mathbf{E} = 0 \qquad \text{(v1)}$$

The eight-point algorithm (`8ptAlg`) employs epipolar constraint:

$\mathbf{A}_{n\times 9} \cdot \mathbf{e} = \mathbf{0}$

The essential matrix

The recovery approaches rely on **two** constraints:

- the epipolar constraint:

$$\mathbf{q}_{i\mathrm{B}}^{\top} \cdot \mathbf{E} \cdot \mathbf{q}_{i\mathrm{A}} = 0$$

- the calibrated (5DOF) constraint:

$$2 \cdot \mathbf{E}\mathbf{E}^T\mathbf{E} - trace(\mathbf{E}\mathbf{E}^T)\mathbf{E} = 0 \quad \text{(v1)}$$

The eight-point algorithm (`8ptAlg`) employs epipolar constraint:
$\mathbf{A}_{n\times 9} \cdot \mathbf{e} = \mathbf{0}$

The five-point algorithm (`5ptAlg`) enforces the 5DOF constraint on the span of lower 4 right-singular vectors of $\mathbf{A}$:

$$\mathbf{E} = a \cdot \mathbf{E_6} + b \cdot \mathbf{E_7} + c \cdot \mathbf{E_8} + d \cdot \mathbf{E_9}$$

THE ESSENTIAL MATRIX

The recovery approaches rely on **two** constraints:

- the epipolar constraint:

$$\mathbf{q}_{i\text{B}}^{\top} \cdot \mathbf{E} \cdot \mathbf{q}_{i\text{A}} = 0$$

- the calibrated (5DOF) constraint:

$$2 \cdot \mathbf{E}\mathbf{E}^T\mathbf{E} - trace(\mathbf{E}\mathbf{E}^T)\mathbf{E} = 0 \quad \text{(v1)}$$

The eight-point algorithm (`8ptAlg`) employs epipolar constraint:
$\mathbf{A}_{n\times 9} \cdot \mathbf{e} = \mathbf{0}$

The five-point algorithm (`5ptAlg`) enforces the 5DOF constraint on the span of lower 4 right-singular vectors of $\mathbf{A}$:

$$\mathbf{E} = a \cdot \mathbf{E_6} + b \cdot \mathbf{E_7} + c \cdot \mathbf{E_8} + d \cdot \mathbf{E_9}$$

This is equivalent to:

$$\mathbf{e}^{\top} \cdot \begin{bmatrix} \mathbf{e_1} & \mathbf{e_2} & \mathbf{e_3} & \mathbf{e_4} & \mathbf{e_5} \end{bmatrix} = \mathbf{0}^{\top} \quad \text{(v2)}$$

THE 8PT-ALG FORWARD BIAS

The i-th row of the matrix $\mathbf{A}$:

$$\mathbf{A}_i = [x_{i\mathrm{B}} x_{i\mathrm{A}} \quad x_{i\mathrm{B}} y_{i\mathrm{A}} \quad x_{i\mathrm{B}} \quad y_{i\mathrm{B}} x_{i\mathrm{A}} \quad y_{i\mathrm{B}} y_{i\mathrm{A}} \quad y_{i\mathrm{B}} \quad x_{i\mathrm{A}} \quad y_{i\mathrm{A}} \quad 1]$$

THE 8PT-ALG FORWARD BIAS

The i-th row of the matrix $\mathbf{A}$:

$$\mathbf{A}_i = [x_{i\mathrm{B}} x_{i\mathrm{A}} \quad x_{i\mathrm{B}} y_{i\mathrm{A}} \quad x_{i\mathrm{B}} \quad y_{i\mathrm{B}} x_{i\mathrm{A}} \quad y_{i\mathrm{B}} y_{i\mathrm{A}} \quad y_{i\mathrm{B}} \quad x_{i\mathrm{A}} \quad y_{i\mathrm{A}} \quad 1]$$

Compare "quadratic" (1,2,4,5) and "linear" (3,6,7,8) columns:

$$a_{i1} = \hat{a}_{i1} + \hat{x}_{i\mathrm{B}} \Delta x_{i\mathrm{A}} + \Delta x_{i\mathrm{B}} \hat{x}_{i\mathrm{A}} + \Delta x_{i\mathrm{B}} \Delta x_{i\mathrm{A}}$$
$$a_{i3} = \hat{a}_{i3} + \Delta x_{i\mathrm{B}}$$

The 8pt-alg forward bias

The i-th row of the matrix $\mathbf{A}$:

$$\mathbf{A}_i = [x_{i\mathrm{B}}x_{i\mathrm{A}} \quad x_{i\mathrm{B}}y_{i\mathrm{A}} \quad x_{i\mathrm{B}} \quad y_{i\mathrm{B}}x_{i\mathrm{A}} \quad y_{i\mathrm{B}}y_{i\mathrm{A}} \quad y_{i\mathrm{B}} \quad x_{i\mathrm{A}} \quad y_{i\mathrm{A}} \quad 1]$$

Compare "quadratic" (1,2,4,5) and "linear" (3,6,7,8) columns:

$$a_{i1} = \hat{a}_{i1} + \hat{x}_{i\mathrm{B}}\Delta x_{i\mathrm{A}} + \Delta x_{i\mathrm{B}}\hat{x}_{i\mathrm{A}} + \Delta x_{i\mathrm{B}}\Delta x_{i\mathrm{A}}$$
$$a_{i3} = \hat{a}_{i3} + \Delta x_{i\mathrm{B}}$$

Under default conditions ($\alpha = 45°$):

$$|a_{i1} - \hat{a}_{i1}| < |a_{i3} - \hat{a}_{i3}|!$$

The 8pt-alg forward bias

The i-th row of the matrix $\mathbf{A}$:

$$\mathbf{A}_i = [x_{i\mathrm{B}}x_{i\mathrm{A}} \quad x_{i\mathrm{B}}y_{i\mathrm{A}} \quad x_{i\mathrm{B}} \quad y_{i\mathrm{B}}x_{i\mathrm{A}} \quad y_{i\mathrm{B}}y_{i\mathrm{A}} \quad y_{i\mathrm{B}} \quad x_{i\mathrm{A}} \quad y_{i\mathrm{A}} \quad 1]$$

Compare "quadratic" (1,2,4,5) and "linear" (3,6,7,8) columns:

$$a_{i1} = \hat{a}_{i1} + \hat{x}_{i\mathrm{B}}\Delta x_{i\mathrm{A}} + \Delta x_{i\mathrm{B}}\hat{x}_{i\mathrm{A}} + \Delta x_{i\mathrm{B}}\Delta x_{i\mathrm{A}}$$
$$a_{i3} = \hat{a}_{i3} + \Delta x_{i\mathrm{B}}$$

Under default conditions ($\alpha = 45°$):

$$|a_{i1} - \hat{a}_{i1}| < |a_{i3} - \hat{a}_{i3}|!$$

The deviation ratio can be determined:

$$r_{Eql} = \sqrt{E[var(a_{i1})]/E[var(a_{i3})]} = tan(\alpha/2) \cdot \sqrt{2/3}$$
$$r_{Eql}(\alpha = 45°) = 0{,}33$$
$$r_{Eql}(\alpha = 102°) = 1$$

THE 8PT-ALG FORWARD BIAS (2)

Estimation favours solutions $\mathbf{E}$ with large $\mathrm{conv}(\mathbf{E}) = |[\mathbf{E}_{13}, \mathbf{E}_{23}, \mathbf{E}_{31}, \mathbf{E}_{32}]|^{-1}$

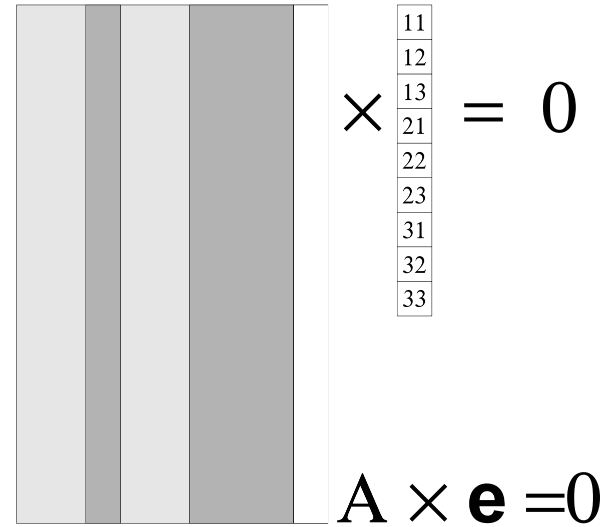

The 8pt-alg forward bias (2)

Estimation favours solutions $\mathbf{E}$ with large $\mathrm{conv}(\mathbf{E}) = ||[\mathbf{E}_{13}, \mathbf{E}_{23}, \mathbf{E}_{31}, \mathbf{E}_{32}]||^{-1}$

For moderate rotations conv attains maximum near the forward direction: $\mathbf{a}_m = \arg\max_a \mathrm{conv}([\mathbf{a}]_\times \mathbf{R}) \approx [\ 0 \quad 0 \quad 1\ ]^\top$

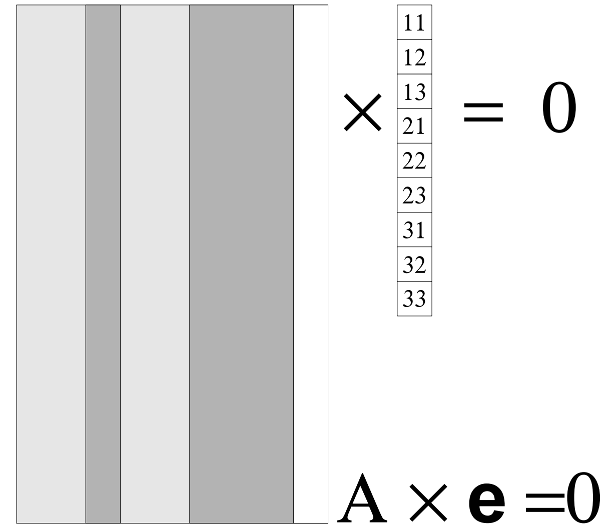

The 8pt-alg forward bias (2)

Estimation favours solutions $\mathbf{E}$ with large $\mathrm{conv}(\mathbf{E}) = ||[\mathbf{E}_{13}, \mathbf{E}_{23}, \mathbf{E}_{31}, \mathbf{E}_{32}]||^{-1}$

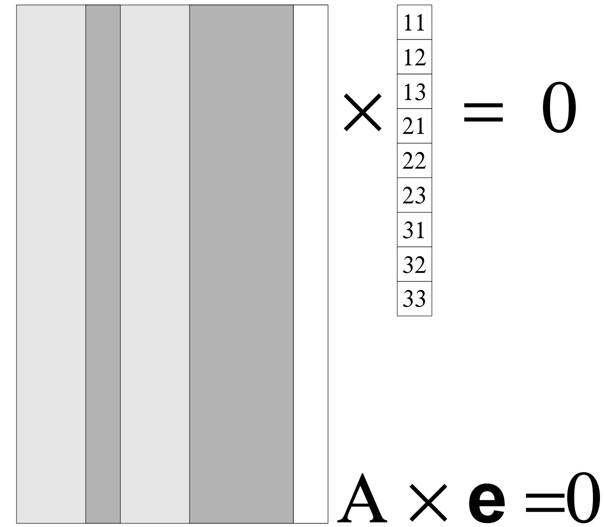

For moderate rotations conv attains maximum near the forward direction: $\mathbf{a}_m = \arg\max_a \mathrm{conv}([\mathbf{a}]_\times \mathbf{R}) \approx [\ 0 \quad 0 \quad 1\ ]^\top$

The bias especially affects the **planar case** when the epipolar constraint is degenerate: $\mathcal{E}(\mathbf{H}) = [\mathbf{a}]_\times \cdot \mathbf{H}, \quad \forall\, \mathbf{a} \in \mathbb{R}^3,$

THE 8PT-ALG FORWARD BIAS (2)

Estimation favours solutions $\mathbf{E}$ with large
$\mathrm{conv}(\mathbf{E}) = ||[\mathbf{E}_{13}, \mathbf{E}_{23}, \mathbf{E}_{31}, \mathbf{E}_{32}]||^{-1}$

For moderate rotations conv attains maximum near the forward direction:
$\mathbf{a}_m = \arg\max_a \mathrm{conv}([\mathbf{a}]_\times \mathbf{R}) \approx [\begin{matrix} 0 & 0 & 1 \end{matrix}]^\top$

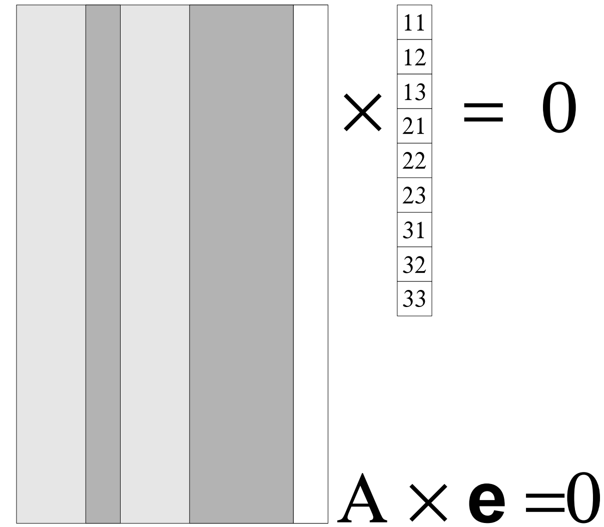

The bias especially affects the **planar case** when the epipolar constraint is degenerate: $\mathcal{E}(\mathbf{H}) = [\mathbf{a}]_\times \cdot \mathbf{H}, \quad \forall\, \mathbf{a} \in \mathbb{R}^3,$

However, the bias also affect the **usual 3D contexts**, where the distance to the target is **much greater** than the baseline

The 8pt-alg forward bias (2)

Estimation favours solutions $\mathbf{E}$ with large
$\text{conv}(\mathbf{E}) = ||[\mathbf{E}_{13}, \mathbf{E}_{23}, \mathbf{E}_{31}, \mathbf{E}_{32}]||^{-1}$

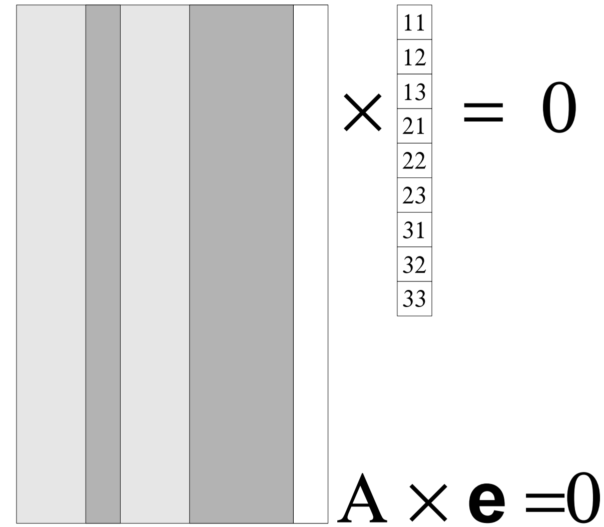

For moderate rotations conv attains maximum near the forward direction:
$\mathbf{a}_m = \arg\max_a \text{conv}([\mathbf{a}]_\times \mathbf{R}) \approx [\begin{array}{ccc} 0 & 0 & 1 \end{array}]^\top$

The bias especially affects the **planar case** when the epipolar constraint is degenerate: $\mathcal{E}(\mathbf{H}) = [\mathbf{a}]_\times \cdot \mathbf{H}, \quad \forall\, \mathbf{a} \in \mathbb{R}^3,$

However, the bias also affect the **usual 3D contexts**, where the distance to the target is **much greater** than the baseline

Here the translation errors **can be approximately compensated** by slight rotation deviations;
small residual changes in the whole translation spectrum!

NUMERICAL CONDITIONING

Review of the 8pt conditioning approaches:

In Hartley's normalization, we recover $\mathbf{E}' = \mathbf{T_2}^{-\top}\mathbf{E}\mathbf{T_1}^{-1}$, relating the transformed points $\mathbf{q}'_{ik} = \mathbf{T_k}\mathbf{q}_{ik}, k = \mathrm{A}, \mathrm{B}$

Numerical conditioning

Review of the 8pt conditioning approaches:

In Hartley's normalization, we recover $\mathbf{E}' = \mathbf{T_2}^{-\top}\mathbf{E}\mathbf{T_1}^{-1}$, relating the transformed points $\mathbf{q'}_{ik} = \mathbf{T_k}\mathbf{q}_{ik}, k = \mathrm{A}, \mathrm{B}$

Mühlich considers an equilibrated matrix $\mathbf{A}_{eq} = \mathbf{W_L} \cdot \mathbf{A} \cdot \mathbf{W_R}$

The new system is $\mathbf{A}_{eq} \cdot \mathrm{e}' = 0$, where $\mathrm{e}' = \mathbf{W_R}^{-1} \cdot \mathbf{e}$

The proposed $\mathbf{W_R}$ ensures a zero-mean expected error in e'

Numerical conditioning

Review of the 8pt conditioning approaches:

In Hartley's normalization, we recover $\mathbf{E}' = \mathbf{T_2}^{-\top}\mathbf{E}\mathbf{T_1}^{-1}$, relating the transformed points $\mathbf{q}'_{ik} = \mathbf{T_k}\mathbf{q}_{ik}, k = \mathrm{A}, \mathrm{B}$

Mühlich considers an equilibrated matrix $\mathbf{A}_{eq} = \mathbf{W_L} \cdot \mathbf{A} \cdot \mathbf{W_R}$

The new system is $\mathbf{A}_{eq} \cdot \mathrm{e}' = 0$, where $\mathrm{e}' = \mathbf{W_R}^{-1} \cdot \mathbf{e}$

The proposed $\mathbf{W_R}$ ensures a zero-mean expected error in e'

Wu et al. have reformulated the linear estimation problem: the new matrix has only linear entries, but is $4n \times (3n + 9)$

Results similar to equilibration

The procedure is much more computationally demanding

CONDITIONING THE 5PT ALGORITHM

Although the **individual** right-singular vectors are very sensitive, their *span* is quite stable!

Deviations $\delta_i = \min(|\mathbf{e_i} - \hat{\mathbf{e}}_\mathbf{i}|, |\mathbf{e_i} + \hat{\mathbf{e}}_\mathbf{i}|)$, sidewise motion, N=$10^4$, σ=1:

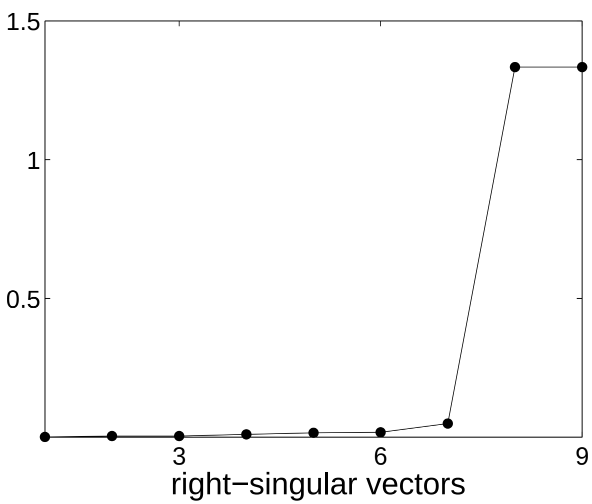

α_H=$45°$, 3D scene

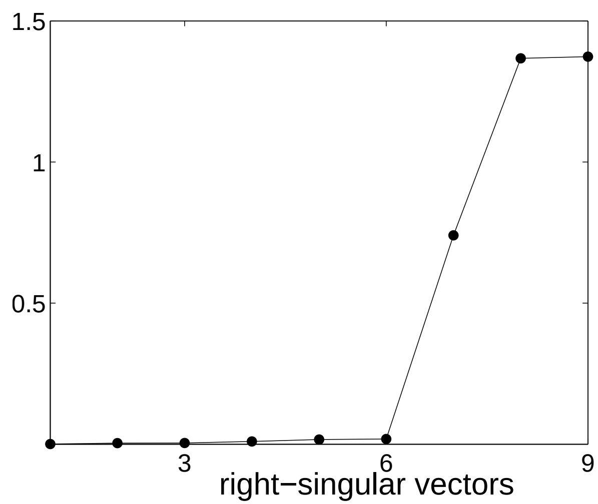

α_H=$45°$, planar scene

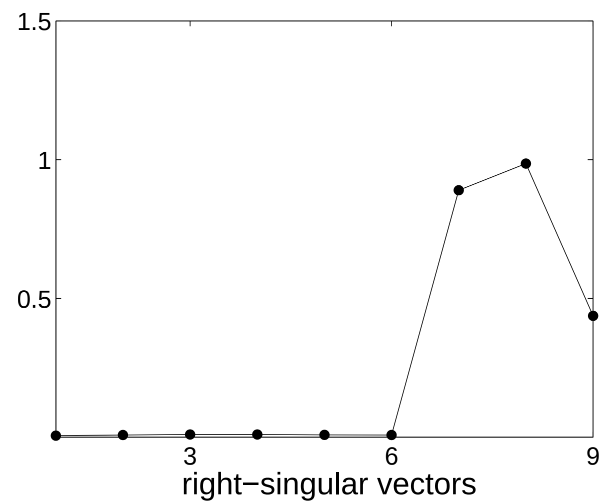

α_H=$120°$, 3D scene

Conditioning the 5pt algorithm

Although the **individual** right-singular vectors are very sensitive, their *span* is quite stable!

Deviations $\delta_i = \min(|\mathbf{e_i} - \hat{\mathbf{e}}_\mathbf{i}|, |\mathbf{e_i} + \hat{\mathbf{e}}_\mathbf{i}|)$, sidewise motion, N=$10^4$, σ=1:

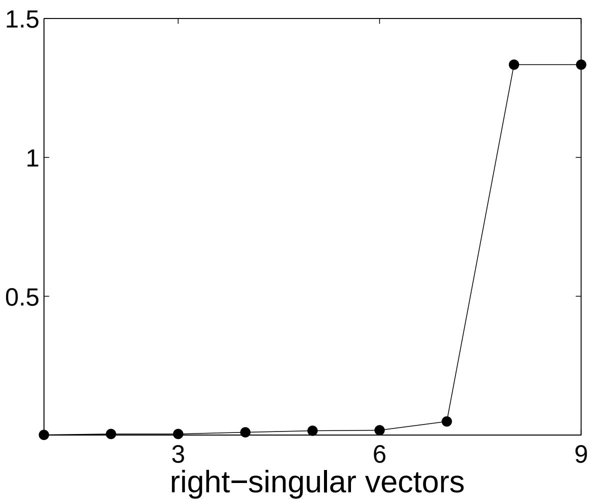

$\alpha_H = 45°$, 3D scene

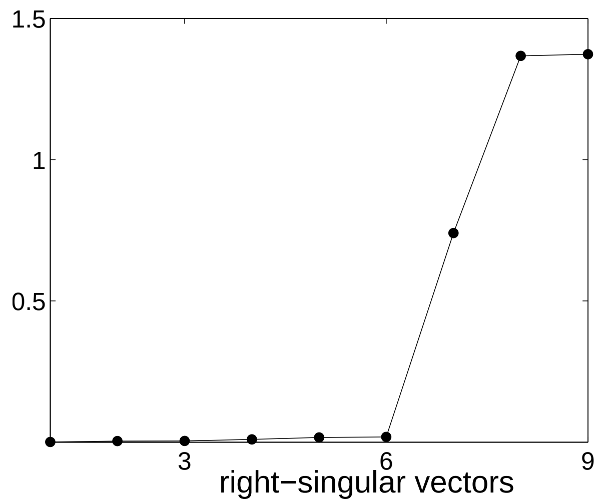

$\alpha_H = 45°$, planar scene

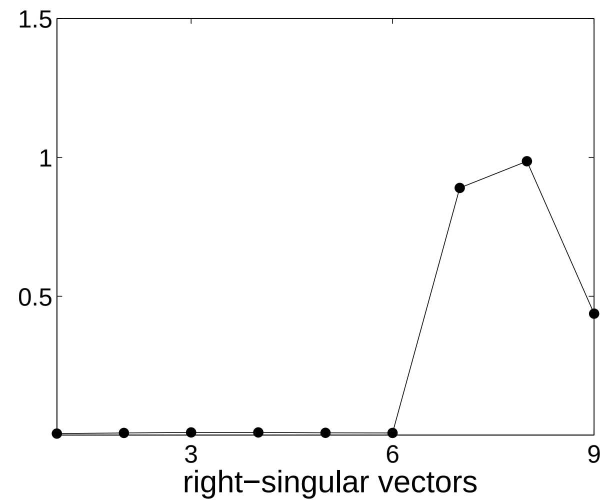

$\alpha_H = 120°$, 3D scene

Hence, the conditioning much less beneficial than with `8ptAlg`.

Conditioning the 5pt algorithm

Although the **individual** right-singular vectors are very sensitive, their *span* is quite stable!

Deviations $\delta_i = \min(|\mathbf{e_i} - \hat{\mathbf{e}}_\mathbf{i}|, |\mathbf{e_i} + \hat{\mathbf{e}}_\mathbf{i}|)$, sidewise motion, N=$10^4$, σ=1:

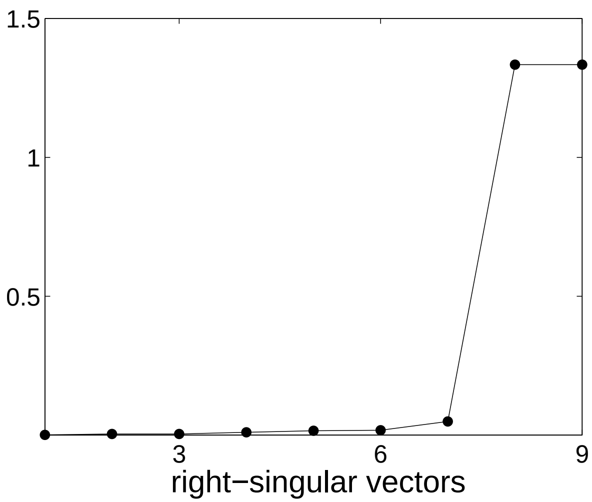

$\alpha_H = 45°$, 3D scene

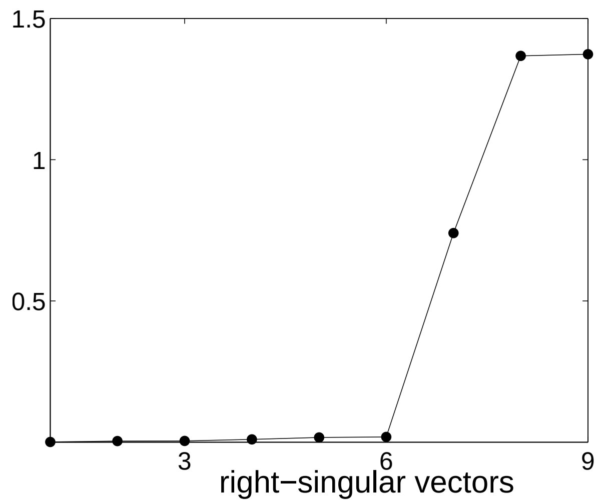

$\alpha_H = 45°$, planar scene

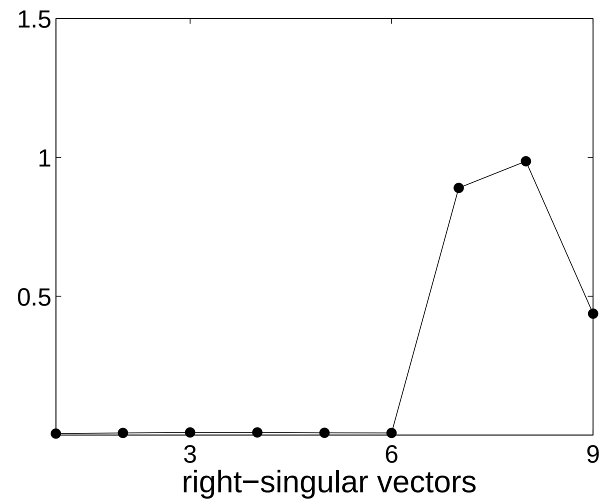

$\alpha_H = 120°$, 3D scene

Hence, the conditioning much less beneficial than with `8ptAlg`.

Conditioning the 5pt algorithm

Although the **individual** right-singular vectors are very sensitive, their *span* is quite stable!

Deviations $\delta_i = \min(|\mathbf{e_i} - \hat{\mathbf{e}}_\mathbf{i}|, |\mathbf{e_i} + \hat{\mathbf{e}}_\mathbf{i}|)$, sidewise motion, N=$10^4$, σ=1:

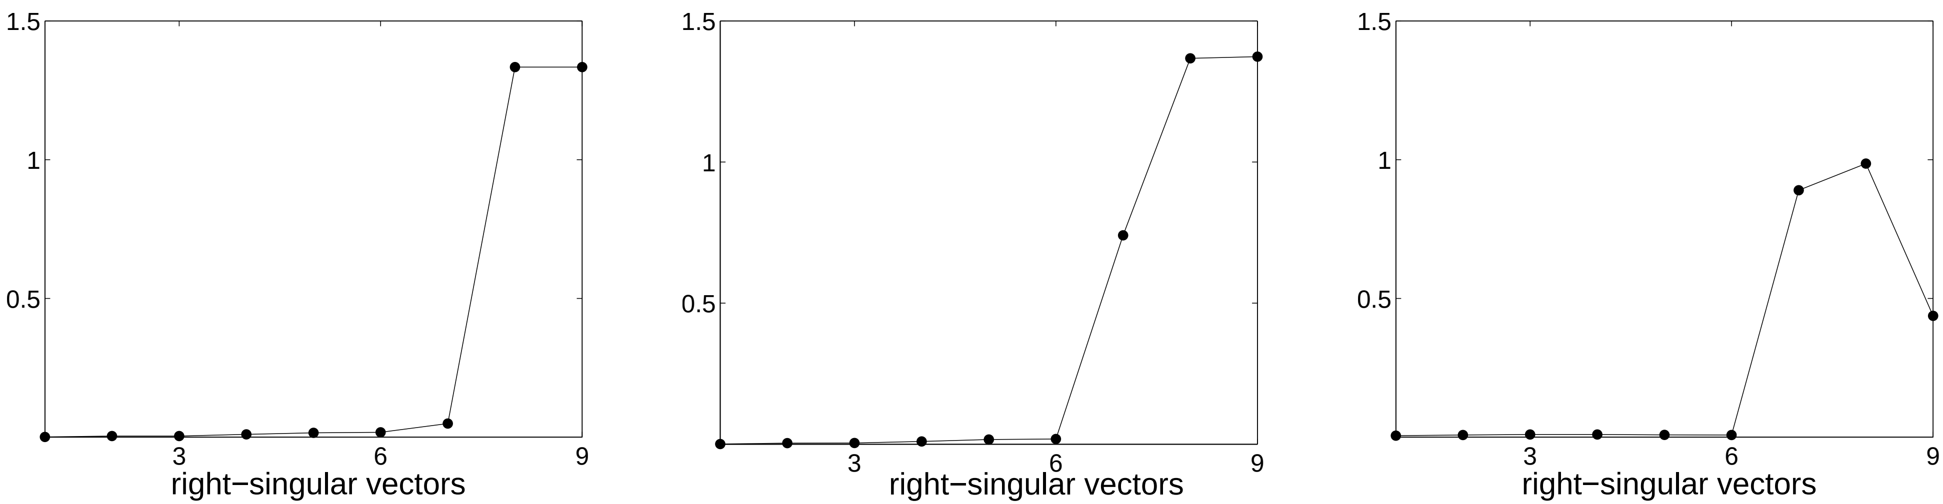

$\alpha_H = 45°$, 3D scene — $\alpha_H = 45°$, planar scene — $\alpha_H = 120°$, 3D scene

Hence, the conditioning much less beneficial than with `8ptAlg`.

Experiments

Parameters of the **artificial** experimental setup:

- geometric: ϕ, θ, `distance`, `depth`, `slant`
- imaging: α_H, σ, resolution[†] for α_H=45° is 384×288

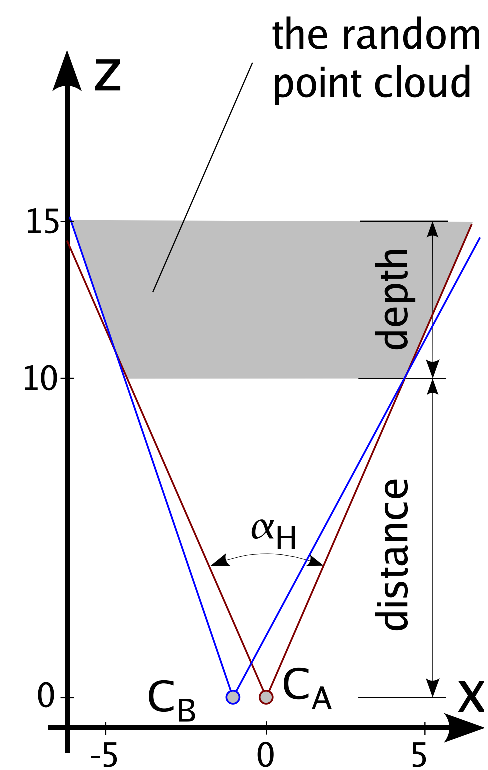

$(-5^\circ, 90^\circ, 10, 5, 0^\circ)$

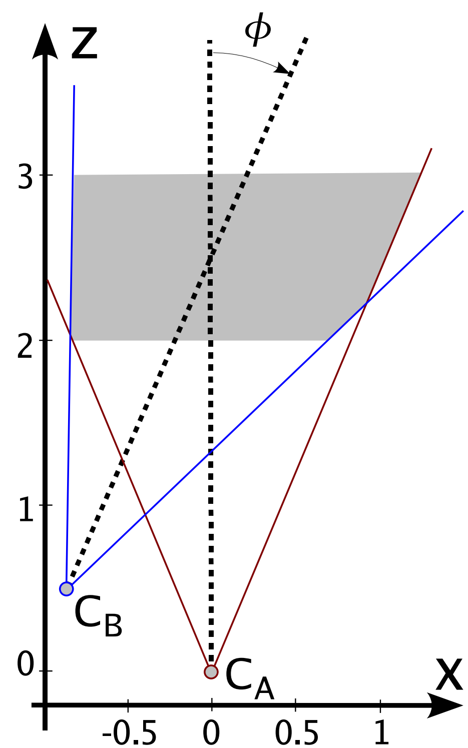

$(-23^\circ, 60^\circ, 2, 1, 0^\circ)$

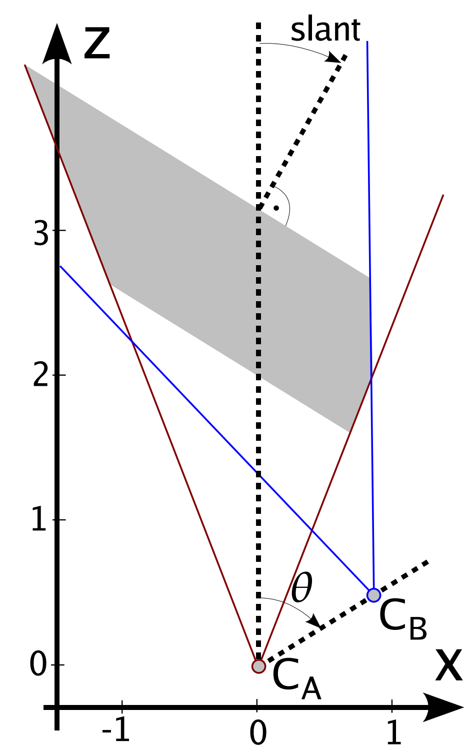

$(23^\circ, -60^\circ, 2, 1, -30^\circ)$

Experiments (2)

We consider the **accuracy** of the recovered epipole $\mathbf{t}$ in variants `standard`, `hartley` and `muehlich`

We perform 10^4 experiments with 50 random points and observe:

Experiments (2)

We consider the **accuracy** of the recovered epipole t
in variants `standard`, `hartley` and `muehlich`

We perform 10^4 experiments with 50 random points and observe:

- Spherical distribution of the **epipole** t (the arrow denotes $\hat{\mathrm{t}}$)

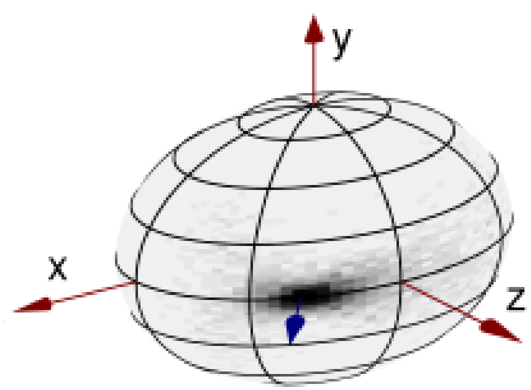

Experiments (2)

We consider the **accuracy** of the recovered epipole $\mathbf{t}$ in variants `standard`, `hartley` and `muehlich`

We perform 10^4 experiments with 50 random points and observe:

- Spherical distribution of the **epipole** $\mathbf{t}$ (the arrow denotes $\hat{\mathbf{t}}$)
- Distribution of the **angular epipole error** $\Delta\mathbf{t} := \measuredangle(\mathbf{t}, \hat{\mathbf{t}})$

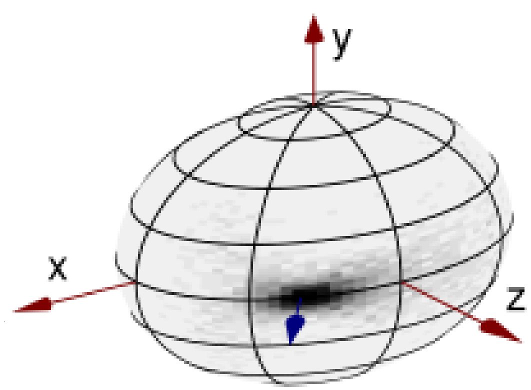

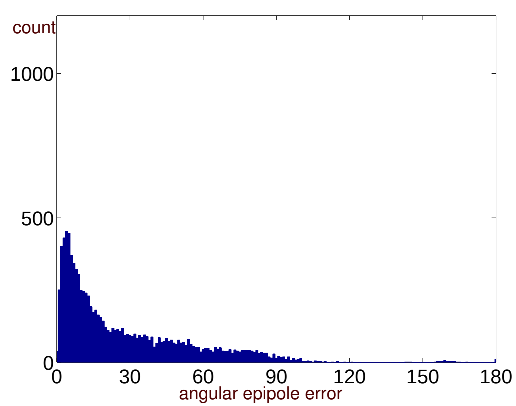

Experiments (2)

We consider the **accuracy** of the recovered epipole t in variants `standard`, `hartley` and `muehlich`

We perform 10^4 experiments with 50 random points and observe:

- Spherical distribution of the **epipole** t (the arrow denotes $\hat{\mathrm{t}}$)
- Distribution of the **angular epipole error** $\Delta\mathrm{t} := \measuredangle(\mathrm{t}, \hat{\mathrm{t}})$
- Dependence of $med\{\Delta\mathrm{t}\}$ on different parameters of the setup

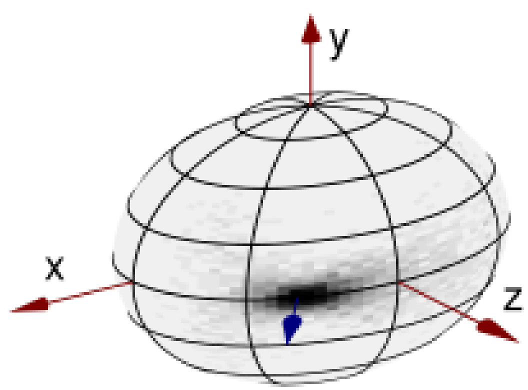

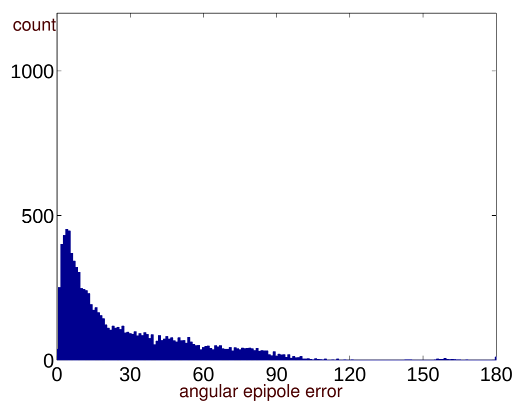

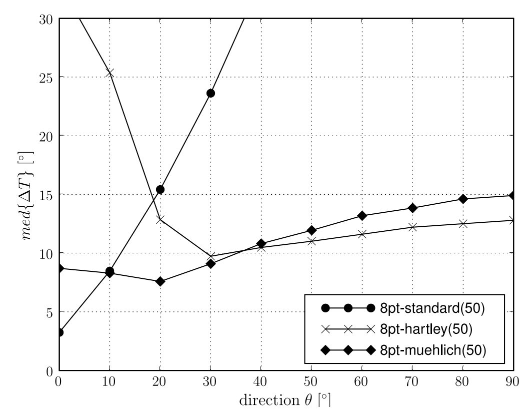

EXPERIMENTS (3)

`8pt-standard` epipoles in **degenerate** and **noisy** datasets:

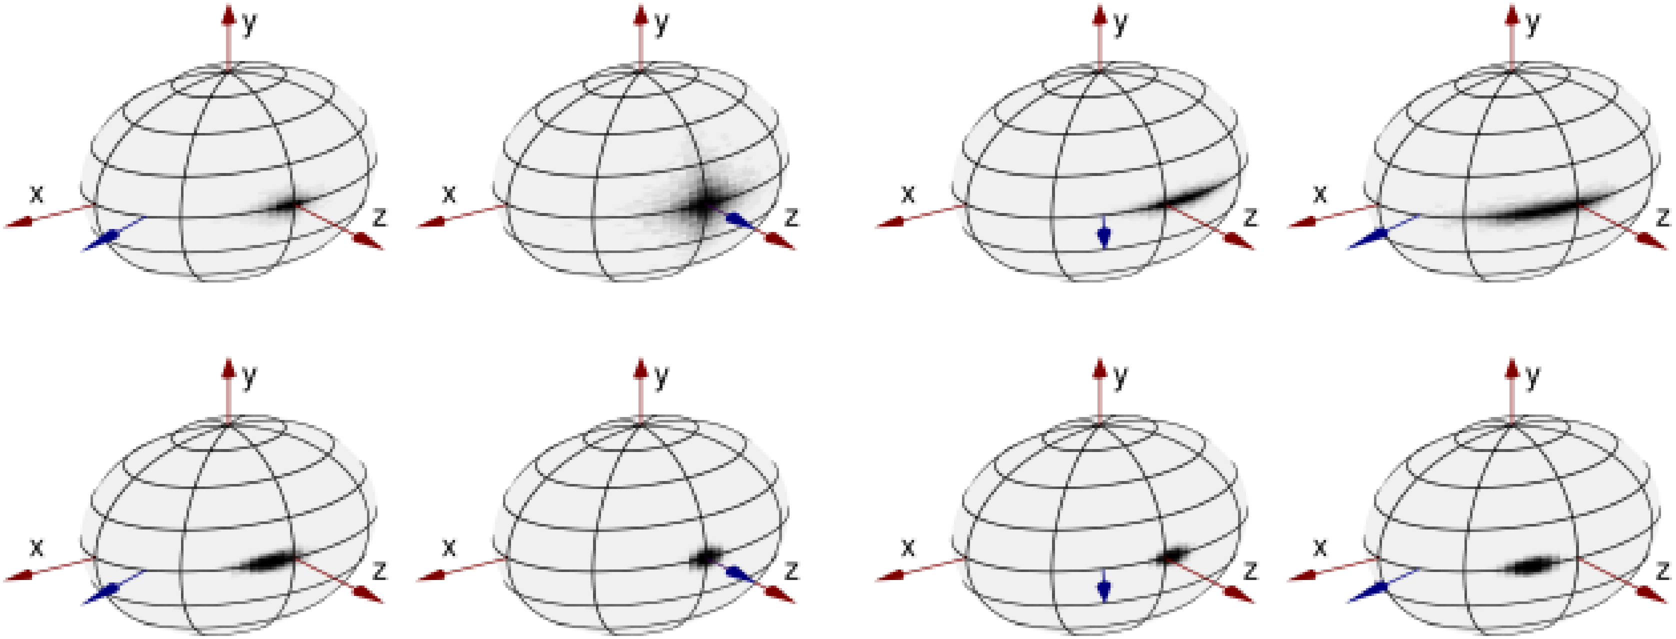

Common: `distance`=10, α_H=45°
Top: `depth`=0, σ=0. Bottom: `depth`=5, σ=1.
Left: θ=(120°,180°), ϕ=0°. Right: θ=135°, ϕ=(-20°,20°).

The shifted modes clearly reflect the forward bias

Backward motion ($|\theta| >$90°) produces t with positive z

EXPERIMENTS (4)

The bias goes away for large α_H, low σ, low `distance` or conditioned data:

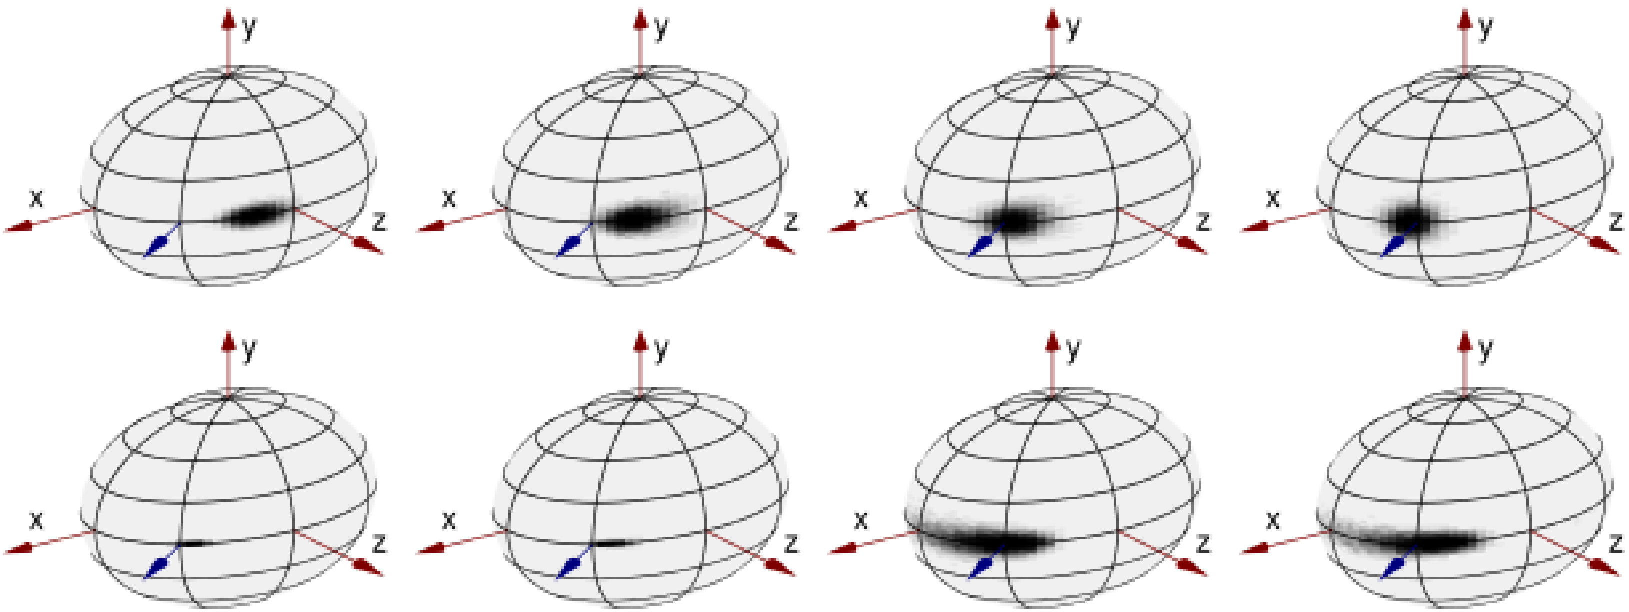

Common: `distance`=10, `depth`=5, θ=135°, ϕ=0°, α_H=45°, $\sigma = 1$
Top: α_H=60°,90°,100°,120°
Bottom: σ=0,2, `distance`=3, normalization, equilibration.

EXPERIMENTS (5)

Normalization and equilibration perform similarly, except for forward motion:

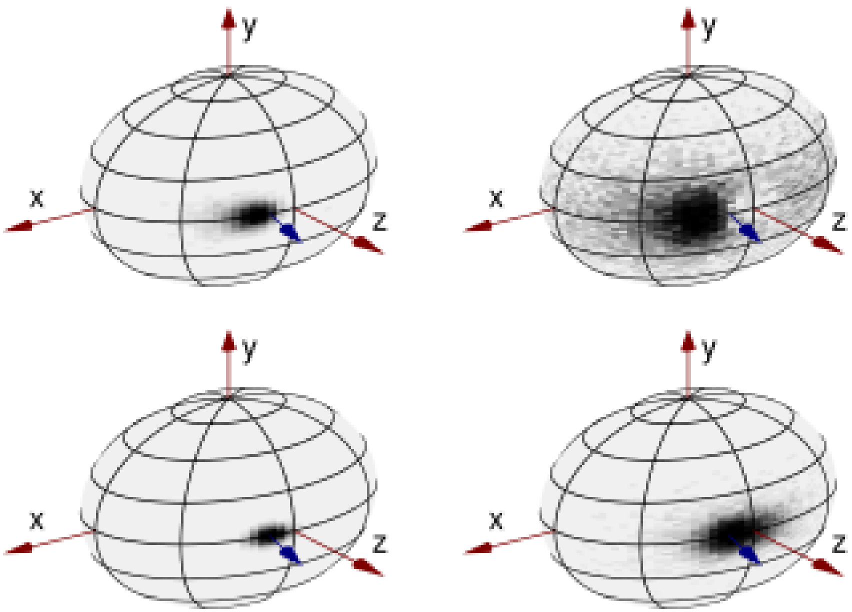

Common: `distance`=10, `depth`=5, θ=170°, ϕ=0°, α_H=45°
Left: σ=$0{,}5$, Right: σ=$1{,}0$
Top: normalization, Bottom: equilibration

Experiments (6)

`5pt` vs `8pt` for 3D scenes ($med\{\Delta\mathbf{t}\}$, `distance`=10, `depth`=5)

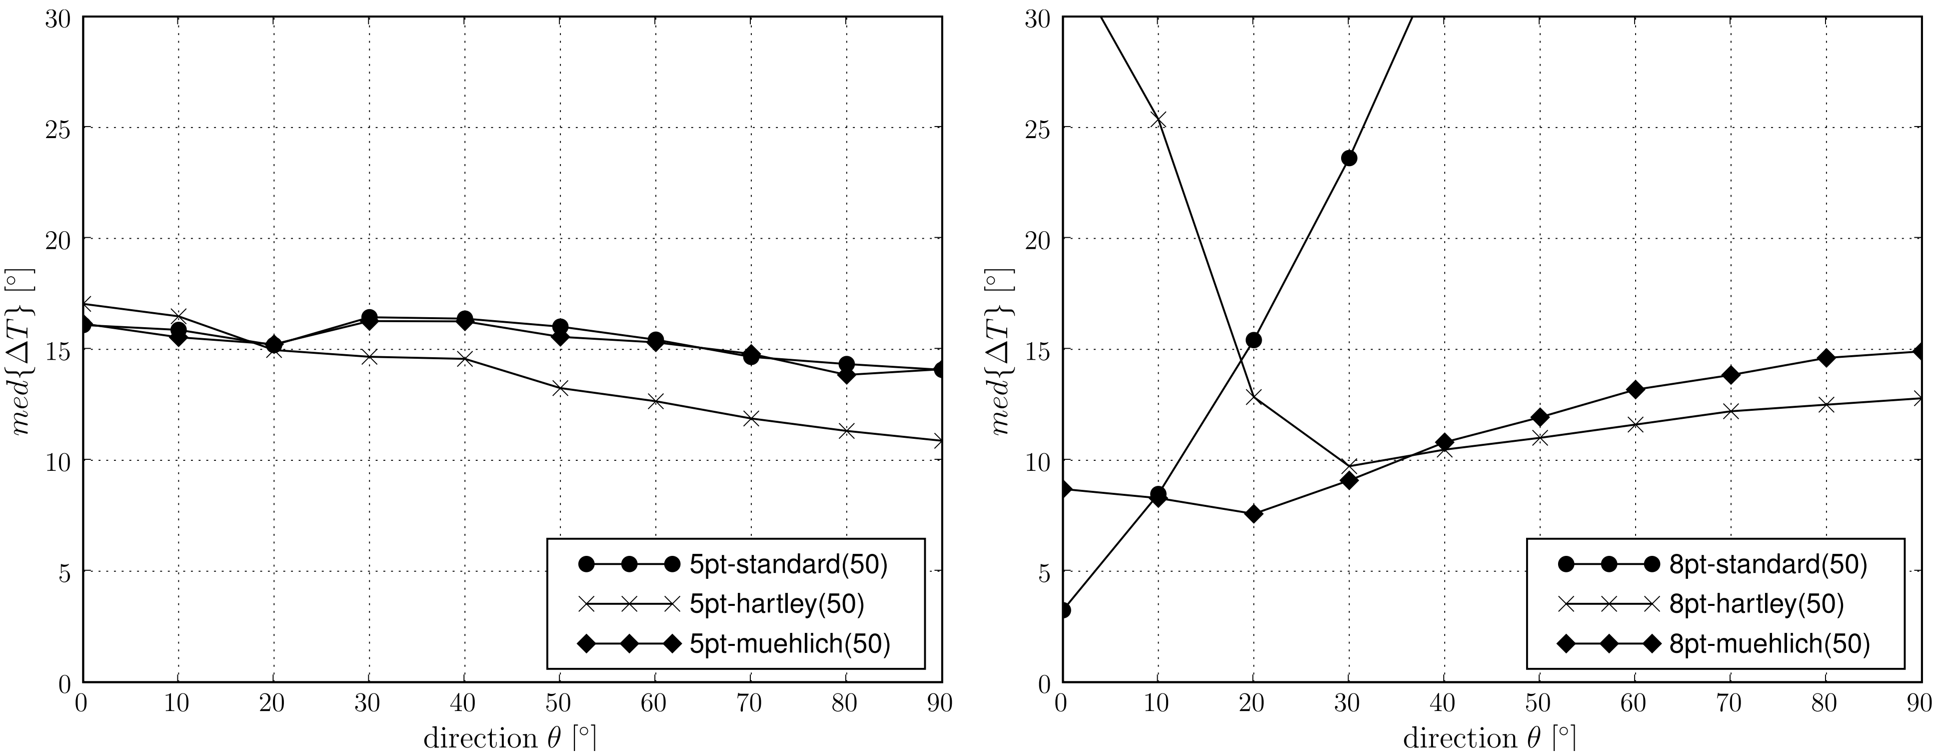

σ=1,0; α_H=45°.

`5pt` disambiguation relies on the total reprojection error

Conditioning helps more `8pt` than `5pt`

EXPERIMENTS (7)

5pt vs hg for planar scenes ($med\{\Delta \mathtt{t}\}$, distance=10, depth=0)

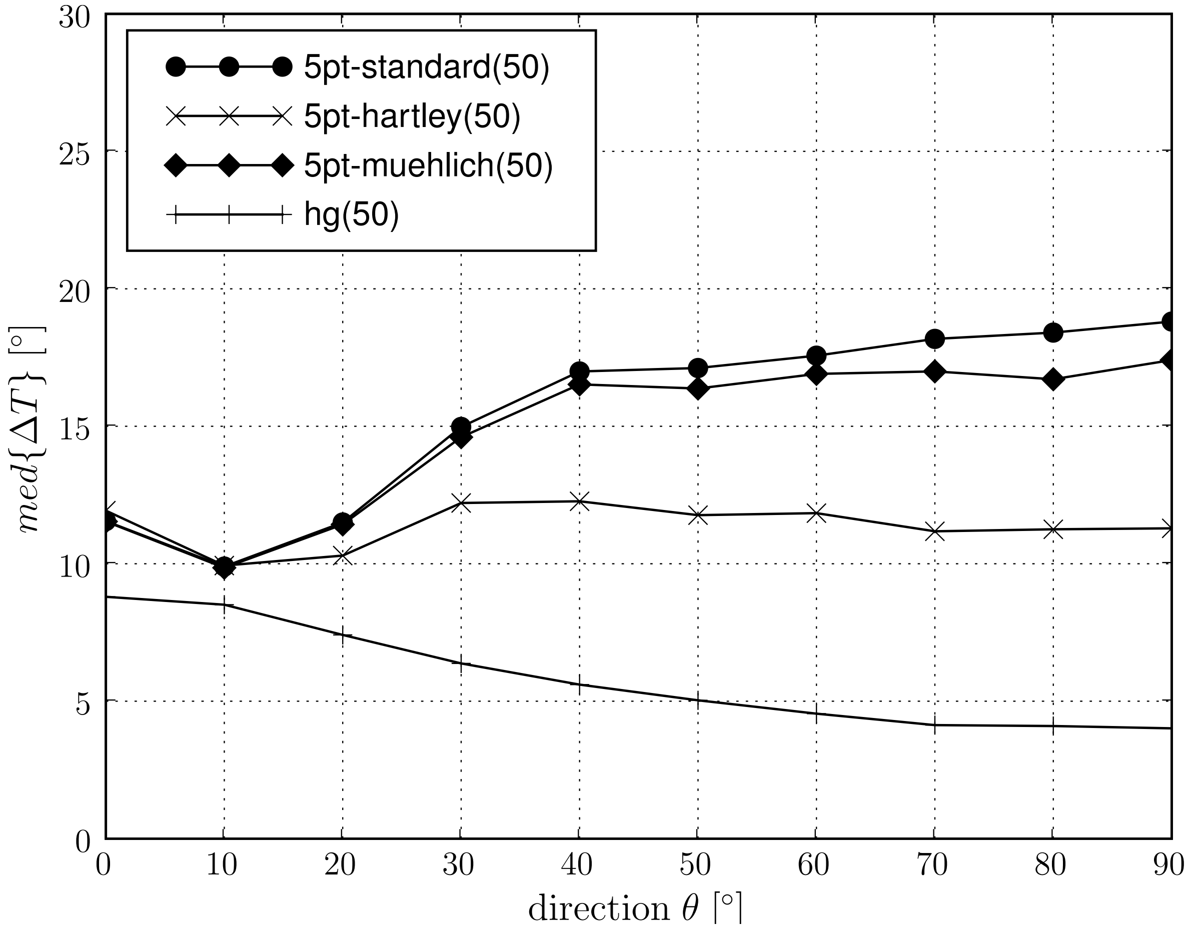

σ=1,0; α_H=45°

5pt and hg disambiguation uses groundtruth!

5pt conditioning always improves the results

hg always better than 5pt

DISCUSSION

The addressed **issues**:

- `8pt` forward bias
- `5pt` numerical conditioning
- experimental validation

DISCUSSION

The addressed **issues**:

- `8pt` forward bias
- `5pt` numerical conditioning
- experimental validation

Conclusions:

- `8pt-standard` performance strongly depends on α_H
- `5pt` conditioning less beneficial than `8pt` conditioning
- `5pt` better than `8pt` for:
  - shallow scenes
  - small number of points
    break-even point: 20 (45°),50 (90°)
- Model selection required for best results