Euclidean Reconstruction and Self Calibration

Bill Triggs, INRIA Rhône-Alpes

GIVEN Moving camera with

- Unknown motion
- Fixed, unknown internal calibration

FIND

- Camera calibration
- Camera motion
- Euclidean reconstruction of scene

- Need $m \geq 3$ images in general position

  – OR –

  Constraints on scene/motion/calibration
- Can't recover overall scale of scene
- Many common camera motions inadequate

Previous Work

Maybank & Faugeras

- Minimal data: 3 images, 8 points
- Kruppa equations & algebraic geometry

— complicated & highly nonlinear

— sensitive to noise

Hartley

$\boxed{1}$ $n \geq 8$ points in $m \geq 3$ images

- Levenberg-Marquardt least squares
- Initialization: projective reconstruction

  + guess of ***plane at infinity***

  + algebra gives absolute conic
- Initialization scheme complicated & unreliable

$\boxed{2}$ Self-calibration of ***rotating*** camera

- Image of absolute conic from homographies
- Stable "linear" approach, assuming no translation

Mohr, Boufama & Brand

- Levenberg-Marquardt least squares
- Initialization from scene constraints

Armstrong, Zisserman & Hartley

- 3 images + motion in a plane
- Epipolar constraint gives rotation axis
- Trilinear constraint gives Euclidean structure

Pollefeys, van Gool & Proesmans

- Initial zoom determines principal point

Quadric-Based Approach

1. Initial projective reconstruction

2. From ***quadric projection equations***, recover

— ***dual absolute quadric*** $\mathbf{Q}$

— ***dual absolute image conic*** $\mathbf{q}$

- Stable nonlinear solution from $m \geq 3$ images
- "Linear" solution from $m \geq 4$ images

3. 'Rectify' to get Euclidean structure

Euclidean Structure in Projective Space

- Given by ***dual absolute quadric*** $\mathbf{Q}^{ab}$
- Contravariant, non-negative, symmetric rank 3 matrix
- *Contravariance:* $\mathbf{Q} \to \mathbf{T}\mathbf{Q}\mathbf{T}^\top$ under projective

transformations $\binom{\mathbf{x}}{1} \to \mathbf{T}\binom{\mathbf{x}}{1}$

- *Euclidean coordinates:* $\mathbf{Q} = \begin{pmatrix} \mathbf{I}_{3\times 3} & \mathbf{0} \\ \mathbf{0} & 0 \end{pmatrix}$
- *Euclidean invariance:*

$\left(\begin{smallmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{smallmatrix}\right)^\top = \left(\begin{smallmatrix} \mathbf{R}\mathbf{R}^\top & \mathbf{0} \\ \mathbf{0} & 0 \end{smallmatrix}\right) = \mathbf{Q}$

$\mathbf{Q}$ represents ***dual absolute conic*** in ***plane at infinity***

- *Affine structure*: $\mathbf{p}^\infty \mathbf{Q} = 0$ defines *plane at infinity*
- *Euclidean structure:*

— Angle between planes: $(\mathrm{n}_1^\top \; d_1)\; \mathbf{Q}\; \binom{\mathrm{n}_2}{d_2} = \mathrm{n}_1 \cdot \mathrm{n}_2$

— Sphere centre $\mathbf{c}$ radius r: $\mathbf{x}^\top \left(\mathbf{c}\mathbf{c}^t - r^2\mathbf{Q}\right)^{-1}\mathbf{x} = 0$

Euclidean Structure in Image

- Camera projection: $\mathbf{P} \sim \mathbf{K}(\mathbf{R}\ \mathbf{t})$

— intrinsic calibration matrix $\mathbf{K} = \begin{pmatrix} \alpha_u & s & u_0 \\ 0 & \alpha_v & v_0 \\ 0 & 0 & 1 \end{pmatrix}$

- $\mathbf{Q}$ projects to ***image of dual absolute conic***

$$\boxed{\mathbf{q} \sim \mathbf{P}\mathbf{Q}\mathbf{P}^\top \sim \mathbf{K}\mathbf{K}^\top}$$

— $\mathbf{q}$ is nonsingular, positive, symmetric 3×3 matrix

— $\mathbf{K}$, $\mathbf{q}$ invariant under camera motions

— recover $\mathbf{K}$ from $\mathbf{q}$ by ***Choleski decomposition***

- $\mathbf{q}$ determines ***Euclidean structure of visual sphere***

— $\mathbf{x}^\top \mathbf{q}^{-1} \mathbf{y}$ gives angle between optical rays

Recovering Euclidean Structure

- Given m images from a moving camera
- ***Projective reconstruction*** gives projections up to

— individual rescaling $\mathbf{P}_i \to \lambda_i \mathbf{P}_i$

— overall homography $\mathbf{P}_i \to \mathbf{P}_i \mathbf{T}$

- In this basis, $\mathbf{Q}$ is unknown, 4×4 rank 3 matrix
- In each image, $\mathbf{Q}$ projects to unknown absolute conic

$\mathbf{q} \sim \mathbf{P}_i \mathbf{Q} \mathbf{P}_i^\top$

— If we knew the scales, could recover $\mathbf{Q}$ and $\mathbf{q}$ *linearly*!

- Antisymmetrize to eliminate scales:

$\mathbf{q} \wedge (\mathbf{P} \mathbf{Q} \mathbf{P}^\top) = \mathbf{0}$, *i.e.*

$$\boxed{\mathbf{q}^{AB} \left(\mathbf{P}_i \mathbf{Q} \mathbf{P}_i^\top\right)^{CD} - \left(\mathbf{P}_i \mathbf{Q} \mathbf{P}_i^\top\right)^{AB} \mathbf{q}^{CD} = \mathbf{0}}$$

— *c.f.* $\quad \mathbf{x} \sim \mathbf{P}\mathbf{X} \quad \Leftrightarrow \quad \mathbf{x} \wedge (\mathbf{P}\mathbf{X}) = \mathbf{0}$

— $15m$ ***bilinear equations*** in 10+6 unknowns $\mathbf{Q}, \mathbf{q}$

- Solve by constrained optimization or “linear” method

Solving for Quadrics

Nonlinear method

- Constrained optimization by ***Sequential Quadratic Programming***
- Minimize $\sum_i \|\mathbf{q} \wedge (\mathbf{P}_i \mathbf{Q} \mathbf{P}_i^\top)\|^2$ given $\det(\mathbf{Q}) = 0$
- Rapid convergence from arbitrary initialization
- Further constraints easily imposed

— known skew, aspect ratio, principal point…

"Linear" method

- $m \geq 4$ images, slower & less stable than SQP
- $\mathbf{q} \wedge (\mathbf{P}_i \mathbf{Q} \mathbf{P}_i^\top) = \mathbf{0}$ is linear in 60 unknowns $\mathbf{x}^{ABab} = \mathbf{q}^{AB} \mathbf{Q}^{ab}$

1. solve linearly for $\mathbf{x}$'s
2. write $\mathbf{x}$'s as 6×10 matrix & project to rank 1
3. extract $\mathbf{q}$, $\mathbf{Q}$ & project $\mathbf{Q}$ to rank 3

Sequential Quadratic Programming

- Newton-like iterative ***constrained optimization*** method
- Minimize cost $f(\mathbf{x})$ given constraints $\mathbf{c}(\mathbf{x}) = \mathbf{0}$
- ***Lagrange multipliers*** $\mathbf{z}$ give solution:

$$\nabla f + \mathbf{z} \cdot \nabla \mathbf{c} = \mathbf{0} \quad \text{with} \quad \mathbf{c}(\mathbf{x}) = \mathbf{0}$$

1. Approximate cost to 2nd order, constraint to 1st order
2. Solve ***constrained quadratic programming*** problem

$$\min_{\delta \mathbf{x}} \left(\nabla f \cdot \delta \mathbf{x} + \frac{1}{2} \delta \mathbf{x}^\top \cdot \nabla^2 f \cdot \delta \mathbf{x} \right) \Big|_{\mathbf{c} + \nabla \mathbf{c} \cdot \delta \mathbf{x} = \mathbf{0}}$$

$$\Rightarrow \boxed{\begin{pmatrix} \nabla^2 f & \nabla \mathbf{c}^\top \\ \nabla \mathbf{c} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \delta \mathbf{x} \\ \mathbf{z} \end{pmatrix} = - \begin{pmatrix} \nabla f \\ \mathbf{c} \end{pmatrix}}$$

3. Iterate to convergence

✔ Quadratic convergence

— *much* better than penalty function methods

✔ Can handle arbitrary constraints

✘ May get stuck in local minima

✘ May blow up due to singular matrix

Algorithm

0 Normalize all image coordinates

1 Projective reconstruction $\Rightarrow$ projections $\mathbf{P}_i$

2 Solve $15m$ bilinear equations $\mathbf{q} \wedge \left(\mathbf{P}_i \mathbf{Q} \mathbf{P}_i^\top\right) = \mathbf{0}$ for $\mathbf{q}$ & $\mathbf{Q}$ with $\det(\mathbf{Q}) = 0$

3 Choleski decomposition $\mathbf{q} = \mathbf{K}\mathbf{K}^\top$

4 SVD $\mathbf{Q}$ & extract rectifying transform $\mathbf{T}\mathbf{Q}\mathbf{T}^\top = \left(\begin{smallmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & 0 \end{smallmatrix}\right)$

5 $\mathbf{P}_i \to \mathbf{K}^{-1}\mathbf{P}_i\mathbf{T}^{-1} \sim (\mathbf{R}_i \ \mathbf{t}_i)$
— adjust to be exactly Euclidean

6 Re-estimate 3D structure with new $\mathbf{P}_i$'s.

Ambiguous Configurations

- Solution not unique for ***many common configurations***

— arbitrary translation + rotation about fixed axis

e.g. pure translation/axis rotation/planar motion/...

— some point $\mathbf{x}$ on $\mathbf{p}^{\infty}$ fixed in all images:

$\mathbf{Q} \to \mathbf{Q} + \lambda \mathbf{x}\mathbf{x}^{\top}$ with $\det \mathbf{Q} = 0$

$\mathbf{q} \sim \mathbf{P}\mathbf{Q}\mathbf{P}^{\top} \to \mathbf{P}\mathbf{Q}\mathbf{P}^{\top} + \lambda (\mathbf{P}\mathbf{x})(\mathbf{P}\mathbf{x})^{\top}$

- Further ambiguities for some ‘symmetric’ configurations

— *c.f.* Peter Sturm’s talk

- Detect continuous ambiguities via rank of Hessian

— discrete ones harder but less common??

- Apply further constraints to remove ambiguity

— known skew, aspect ratio, motion,...

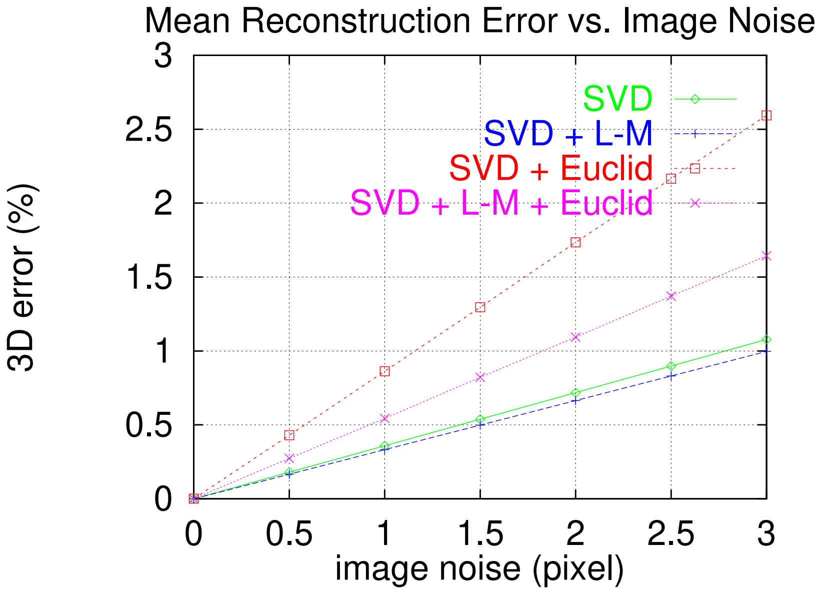
Mean Reconstruction Error vs. Image Noise
3D error (%)
image noise (pixel)
SVD
SVD + L-M
SVD + Euclid
SVD + L-M + Euclid

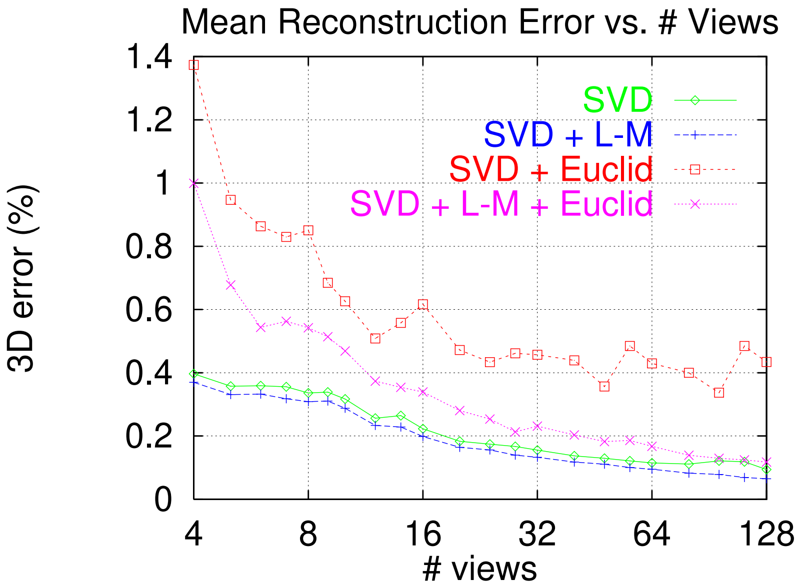
Mean Reconstruction Error vs. # Views
3D error (%)
views
SVD
SVD + L-M
SVD + Euclid
SVD + L-M + Euclid

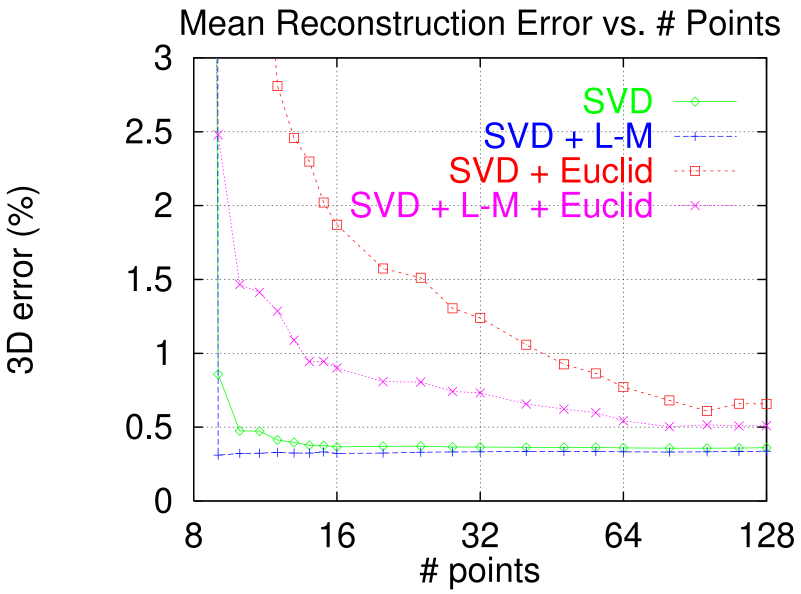
Mean Reconstruction Error vs. # Points
3D error (%)
points
SVD
SVD + L-M
SVD + Euclid
SVD + L-M + Euclid

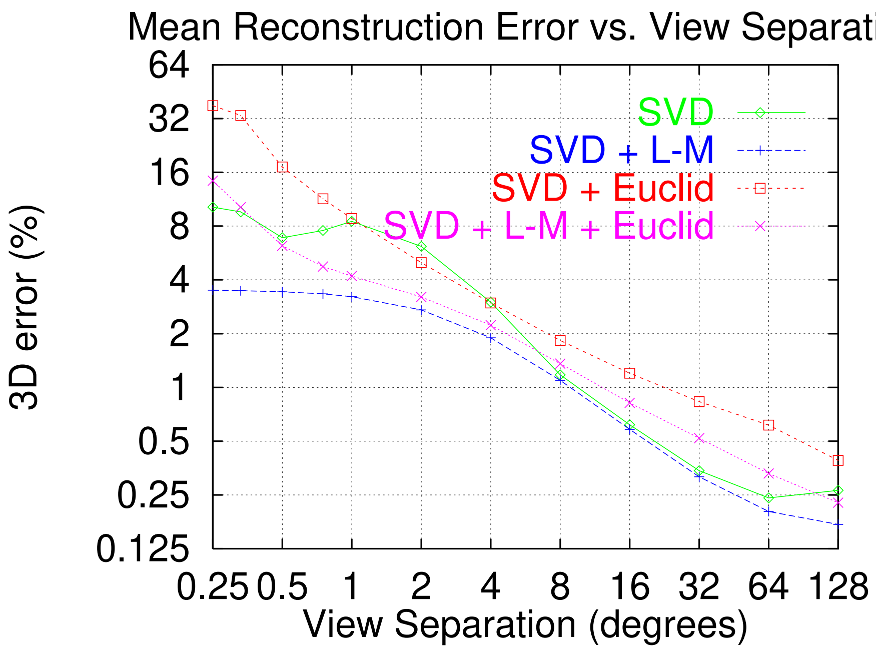
Mean Reconstruction Error vs. View Separation
3D error (%)
64
32
16
8
4
2
1
0.5
0.25
0.125
0.25
0.5
1
2
4
8
16
32
64
128
View Separation (degrees)
SVD
SVD + L-M
SVD + Euclid
SVD + L-M + Euclid

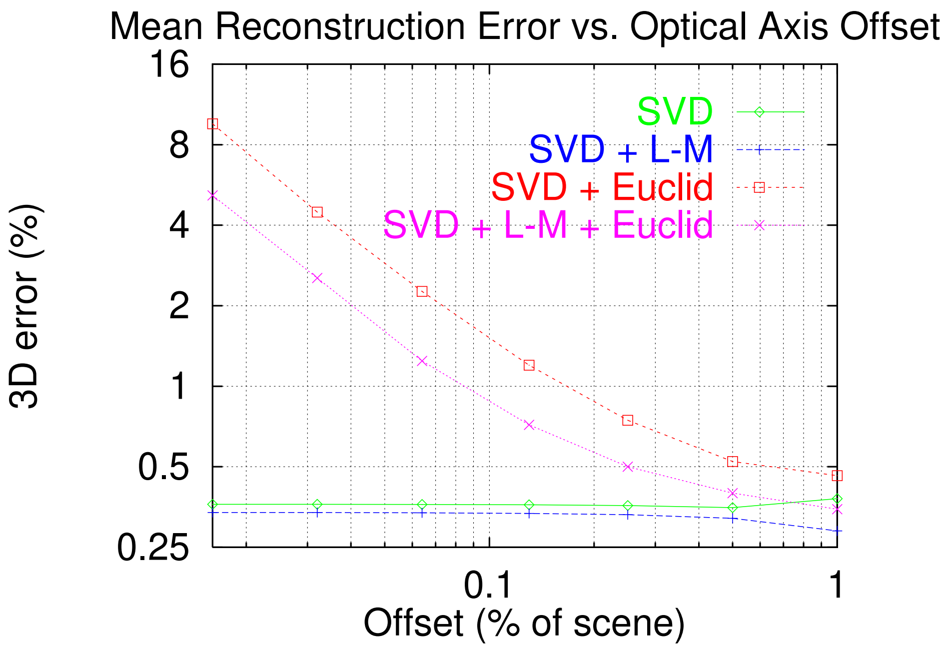
Mean Reconstruction Error vs. Optical Axis Offset
3D error (%)
16
8
4
2
1
0.5
0.25
0.1
1
Offset (% of scene)
SVD
SVD + L-M
SVD + Euclid
SVD + L-M + Euclid

Conclusions & Future Work

Euclidean structure & self-calibration

- Based on projective reconstruction + quadric projection
- Linear & nonlinear methods — nonlinear better!
- Works well for $m \geq 3$ cameras in ***general position***
- **BUT** fails for many common "degenerate" configurations

— extra constraints required

Future Work

- Handle degeneracy, extra constraints & zoom
- Several independent cameras & stereo head
- Recursive formulation & bundle adjustment