Structure from Motion

Carlo Tomasi

SFM

- moving camera, static environment
- otherwise, motion segmentation
- image motion = f(shape, camera motion)
- shape, camera motion = arg min ||i - f(s, c)||
- two frames and more

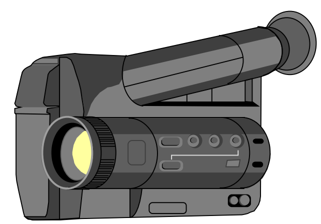

Summary

- measuring image motion
- the geometry of SFM
- SFM is ill-conditioned
- algebra + least squares = trouble
- SFM as MLE
- sensitivity
- promising approaches

Measuring Image Motion

- image motion is a matter of interpretation:
  $I(x, y, t) = I(x_0(x, y, t), y_0(x, y, t), 0)$

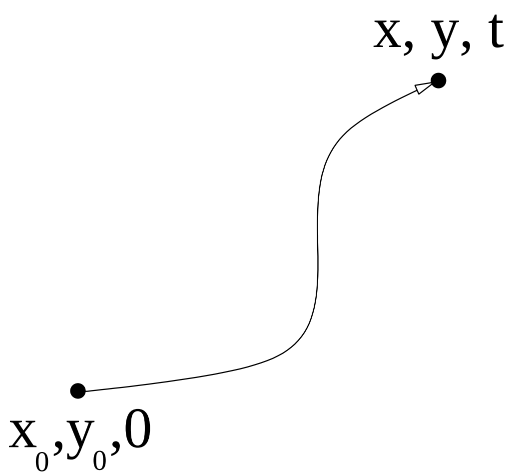

- poor features → "noise"
- false features → outliers
- miscalibration → systematic errors

Epipoles

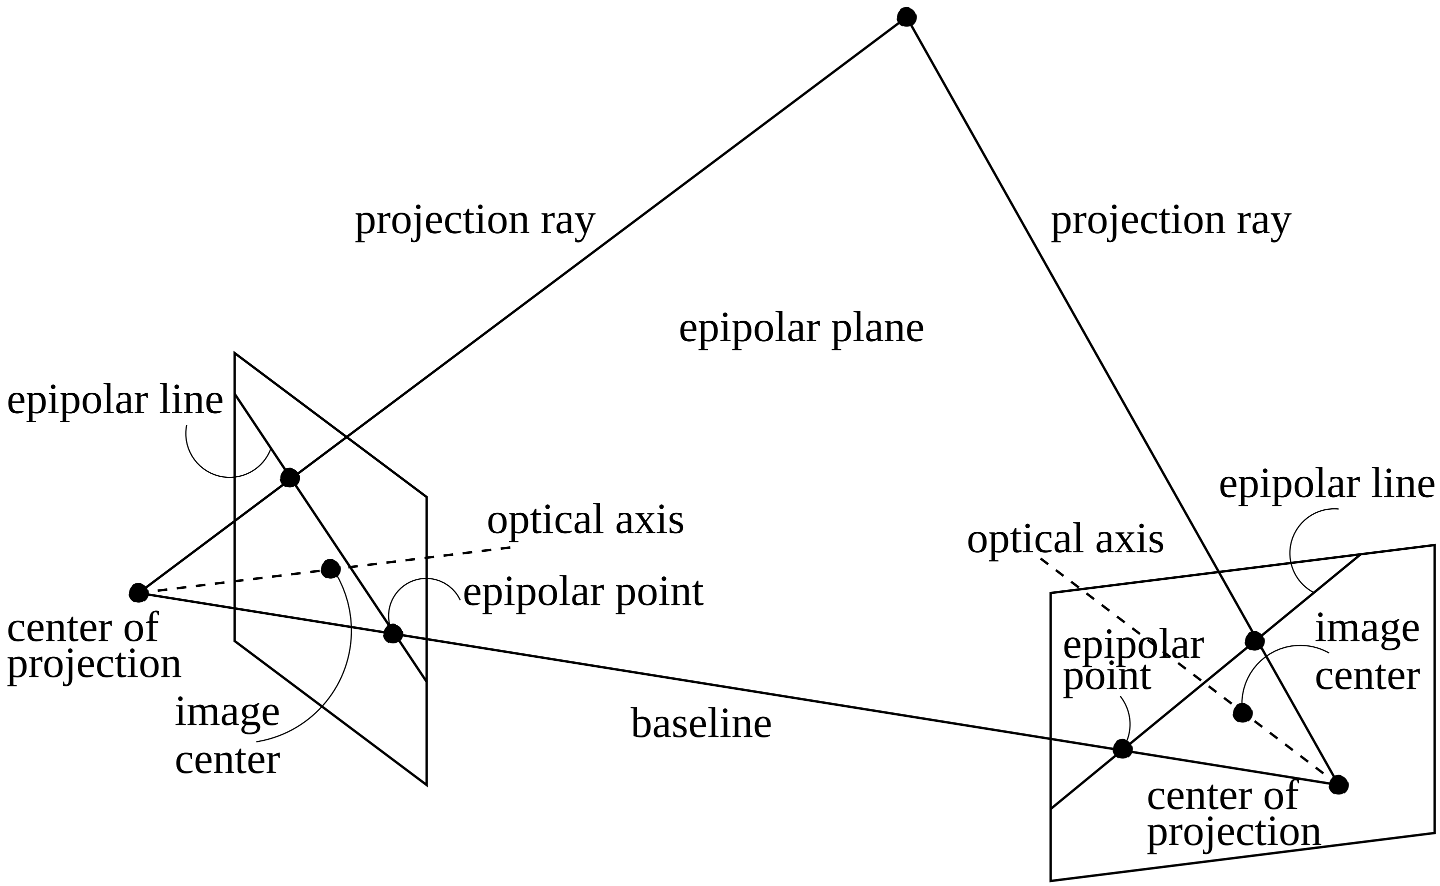

projection ray
projection ray
epipolar plane
epipolar line
optical axis
epipolar point
center of projection
image center
baseline
epipolar line
optical axis
image center
epipolar point
center of projection

Focus of Expansion

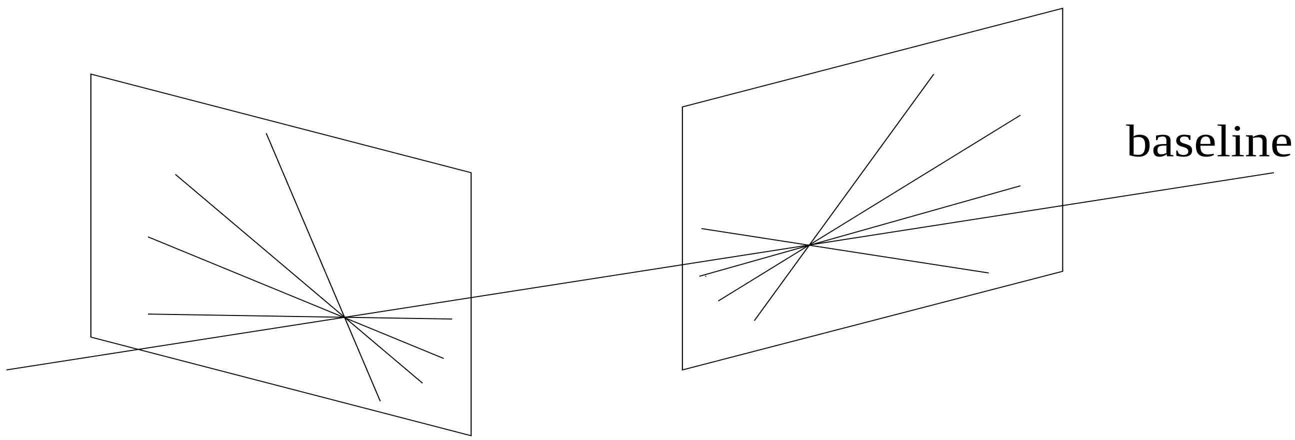

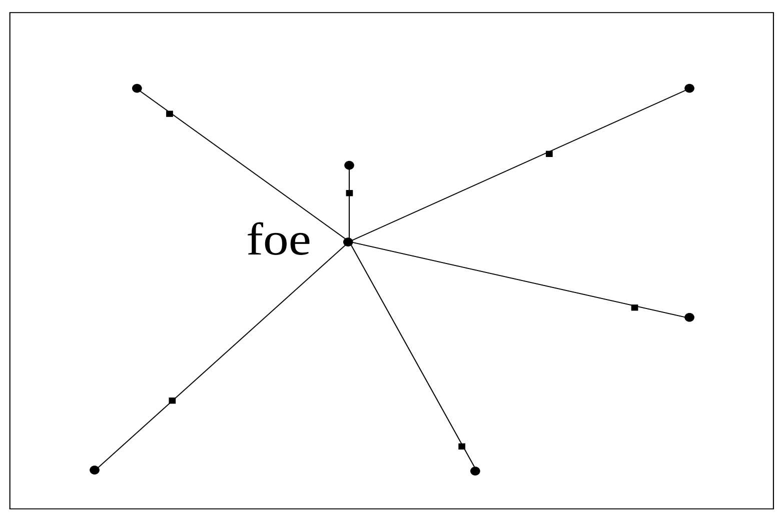

f.o.e. = direction

of translation

2-Frame SFM

- 3D rotate one image until epipolars intersect
- find f.o.e. (translation) as intersection point
- compute depths (shape) from disparities

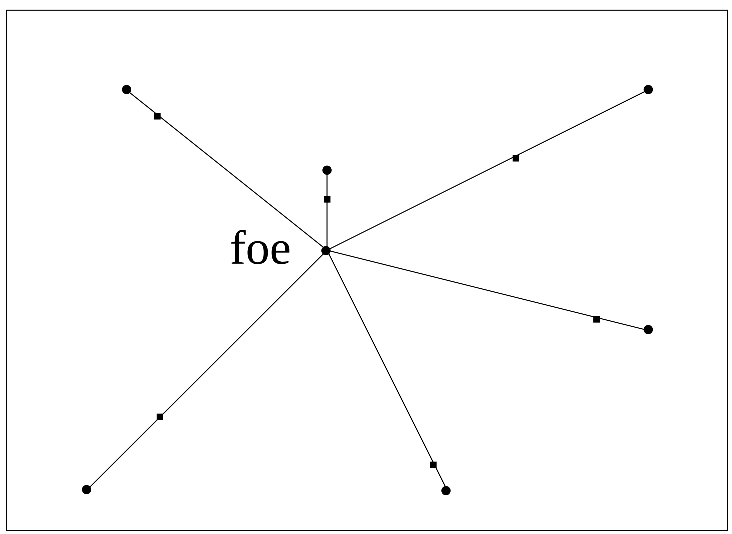

SFM is an Inverse Problem

- vision and graphics
- a hard case: small field of view, lateral motion

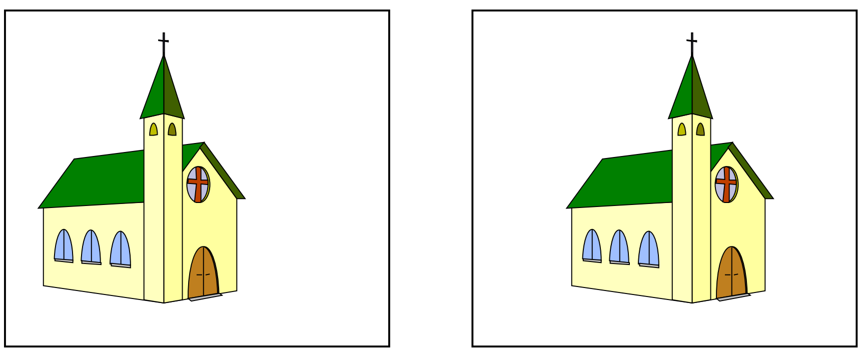

rotation? translation?

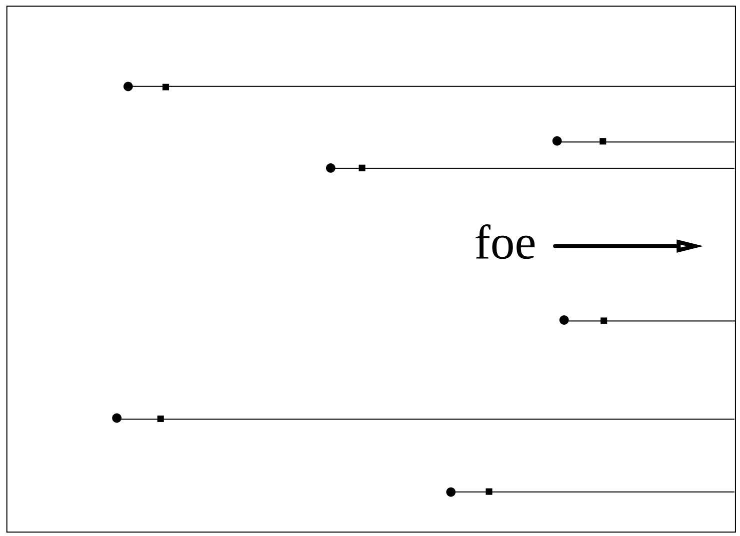

An “Easy” Problem: Translation Only

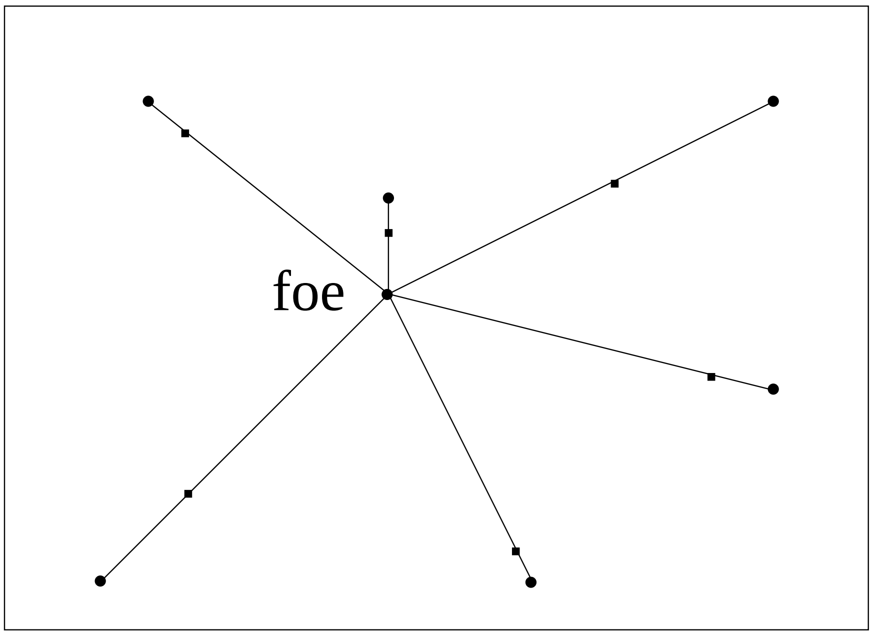

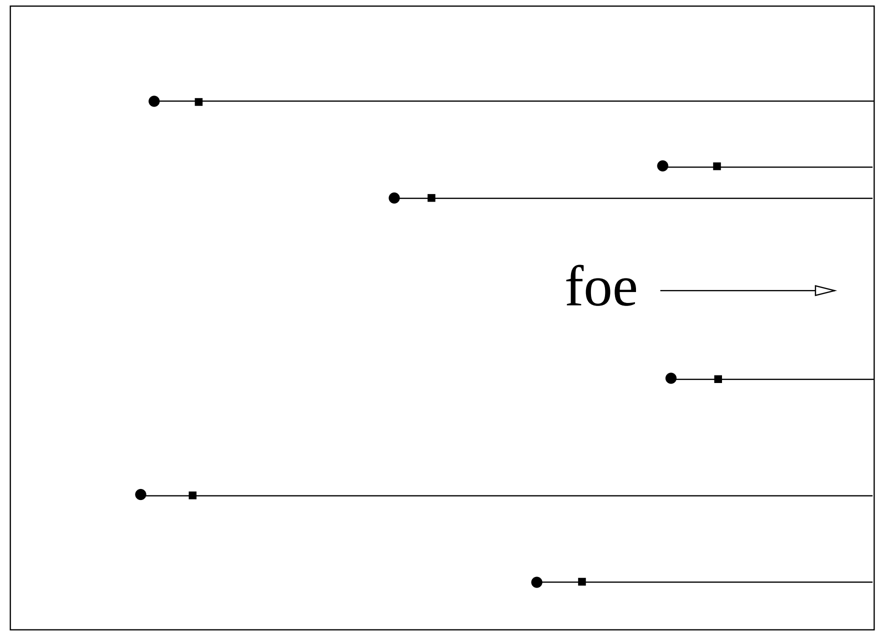

The Standard Solution: Algebra + LSE

- Point pairs → lines
- f.o.e. = point closest to all the lines

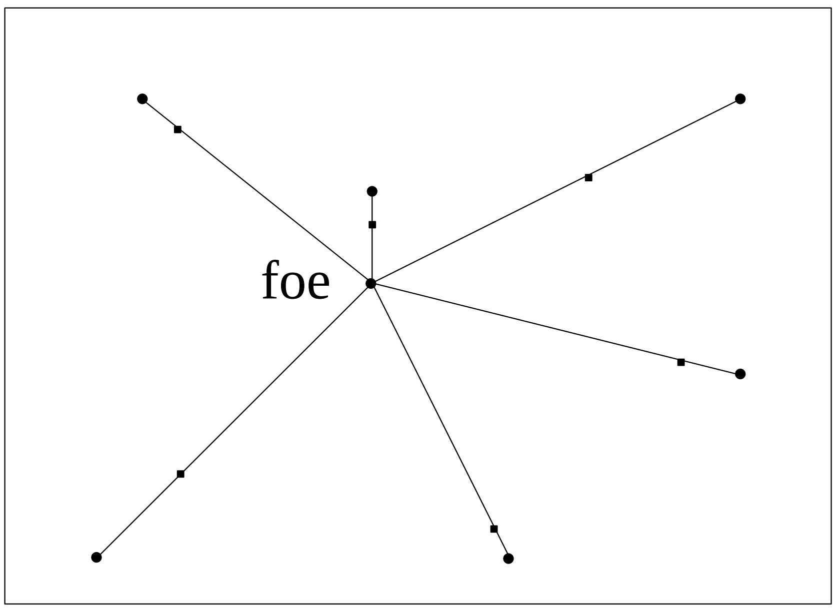

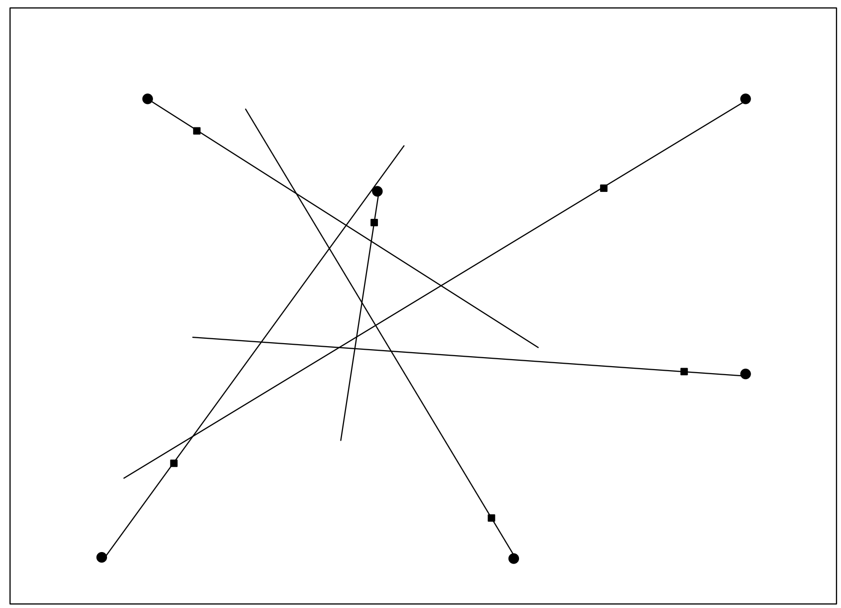

Problems

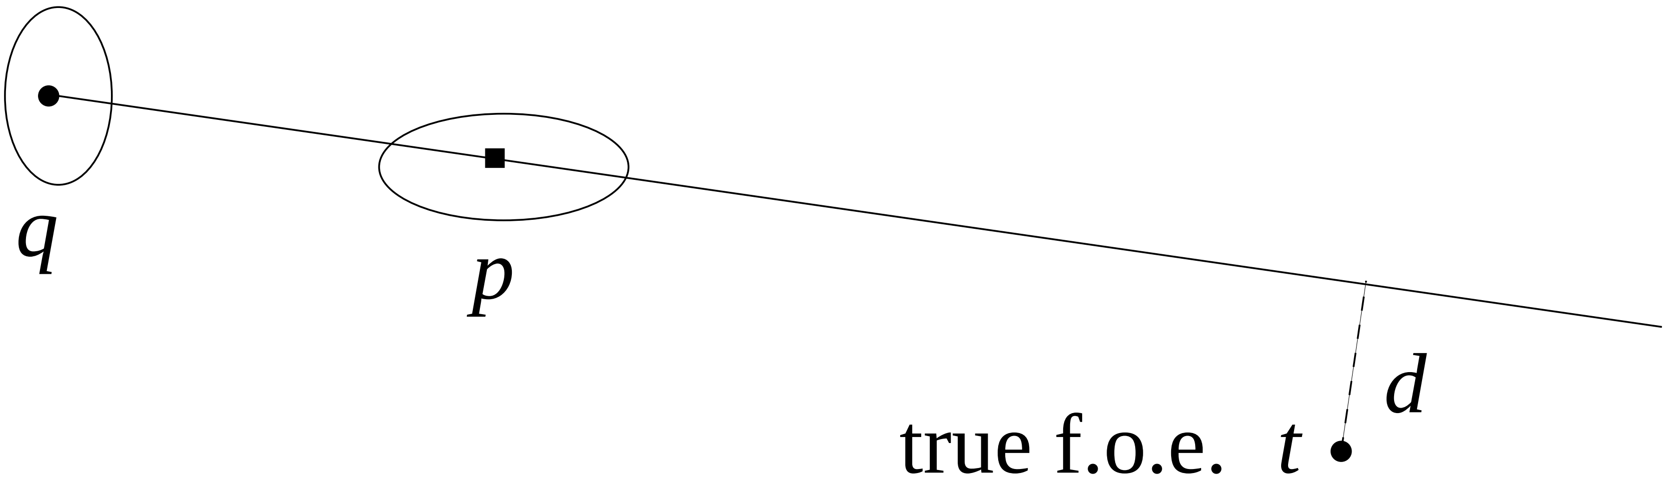

- ellipses, not circles
- line wobble depends on *dist(p,q)*
- effect of line wobble depends on *dist(p,t)*

Standard Fix

- error propagation:

$$\sigma(\Delta p,\, \Delta q) \Rightarrow \sigma(\Delta d)$$

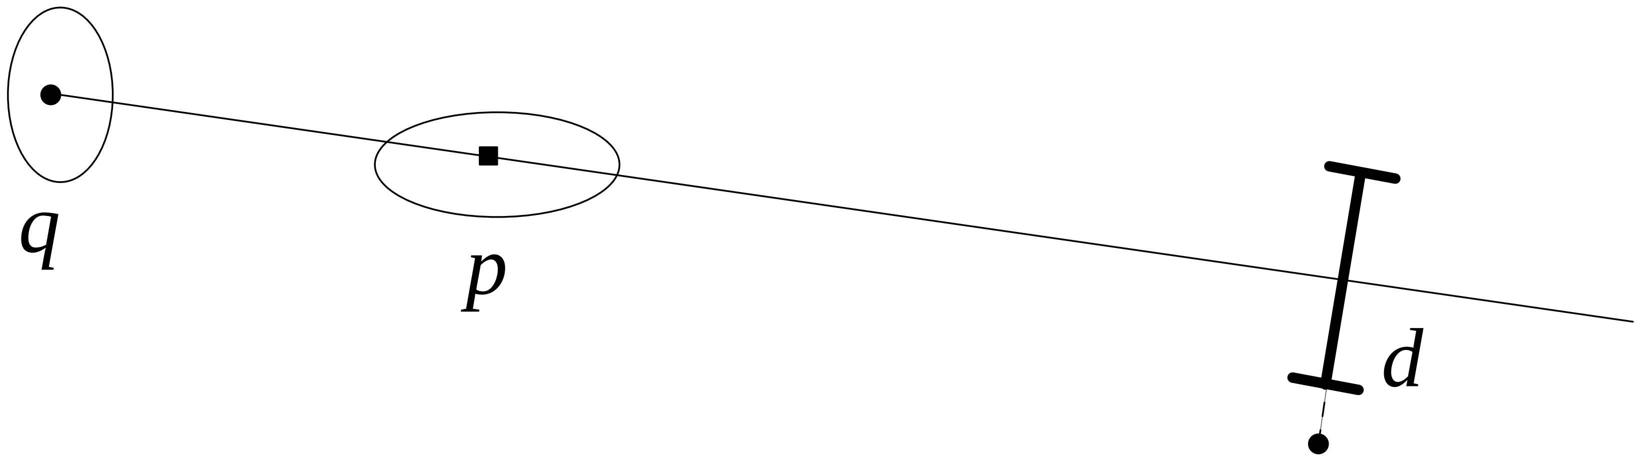

Problems with Standard Solution

- $d(p, q)$ is not linear
- small Δp, Δq, large Δd
- $pdf(\Delta d)$ is not Gaussian
- **bias**
- propagating *pdf* is hard and leads to a hard minimization problem

Maximum Likelihood Estimation

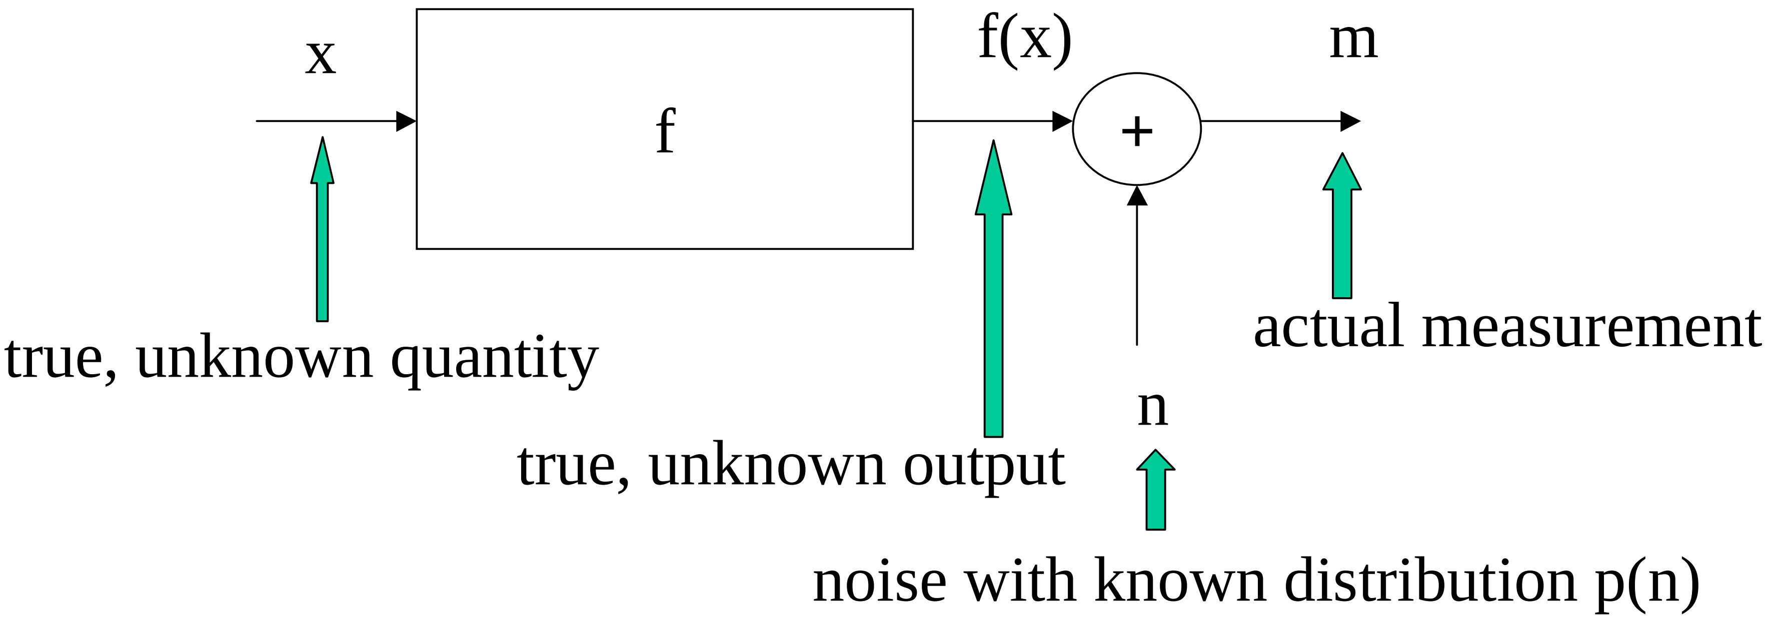

$$\max_{x} \mathrm{pdf}(m|x) = \max_{x} \mathrm{pdf}(m - f(x))$$

Least Squares < MLE < Bayes

LSE vs MLE

LSE

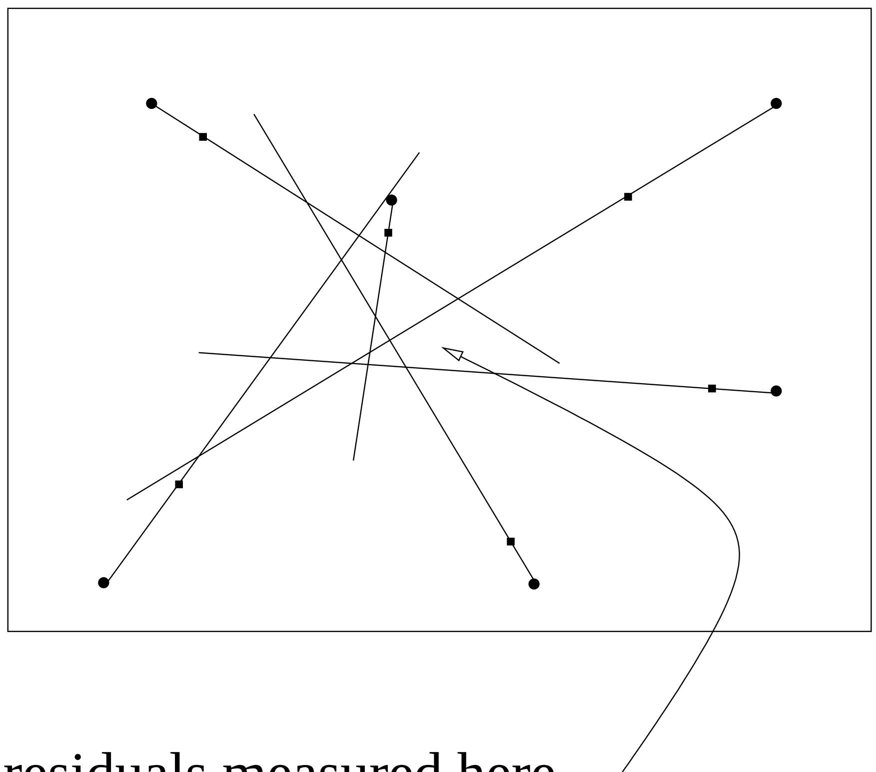

residuals measured here

MLE

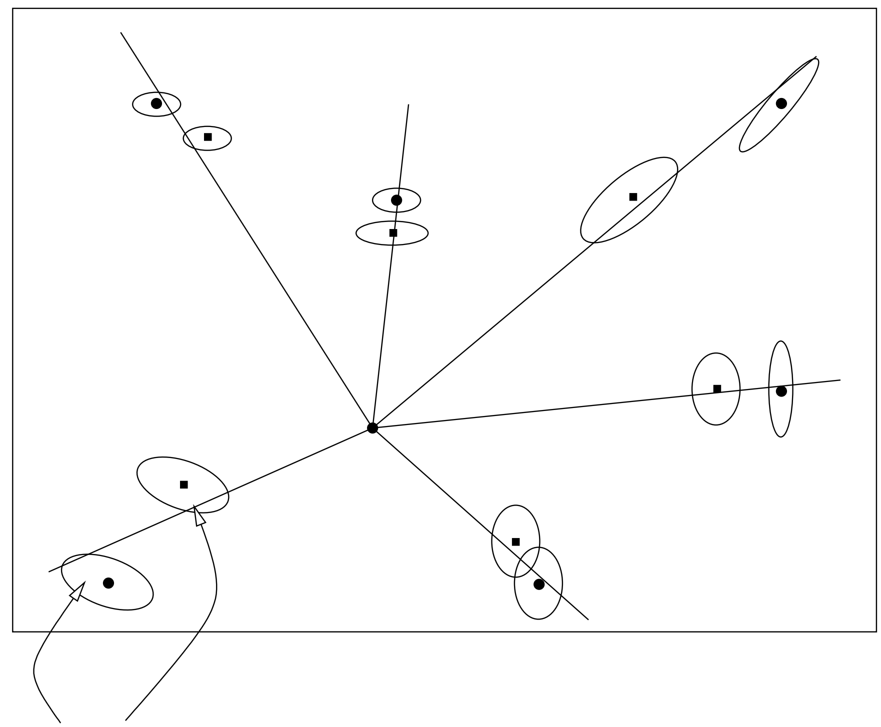

residuals measured here

Cascading MLE Problems

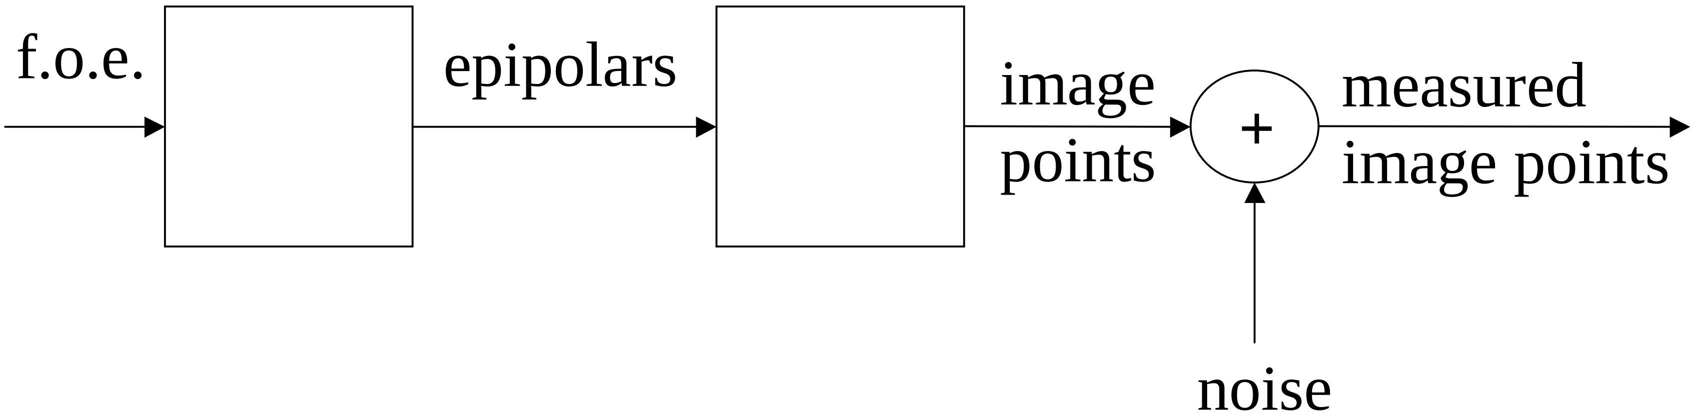

$$\max_{x,y} \mathrm{pdf}(m|x,y) = \max_{x} \max_{y} \mathrm{pdf}(m|y|x)$$

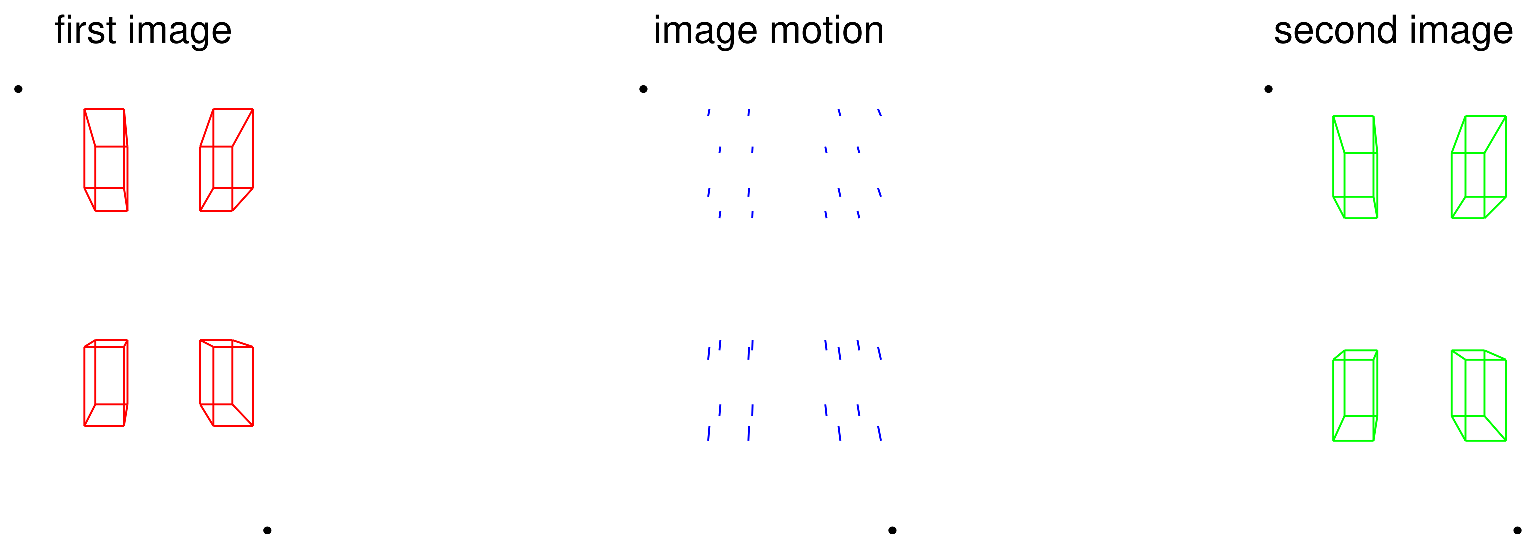
first image
image motion
second image

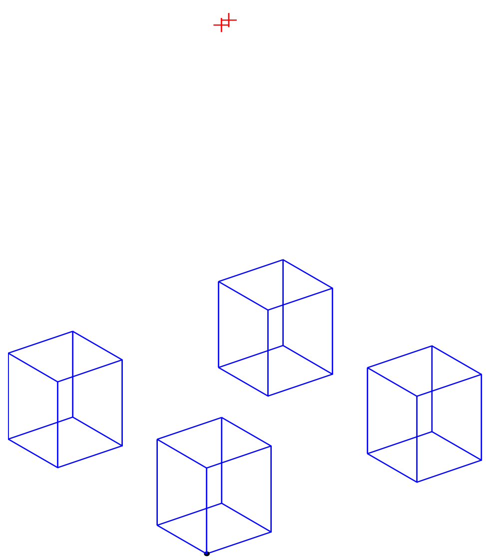

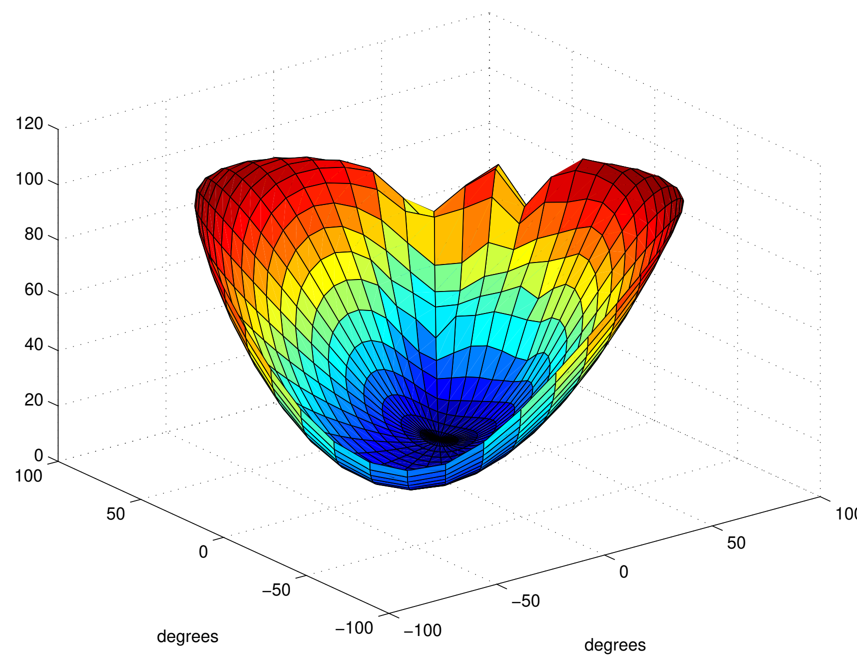
120
100
80
60
40
20
0
100
50
0
−50
−100
degrees
−100
−50
0
50
100
degrees

LSE vs MLE Errors

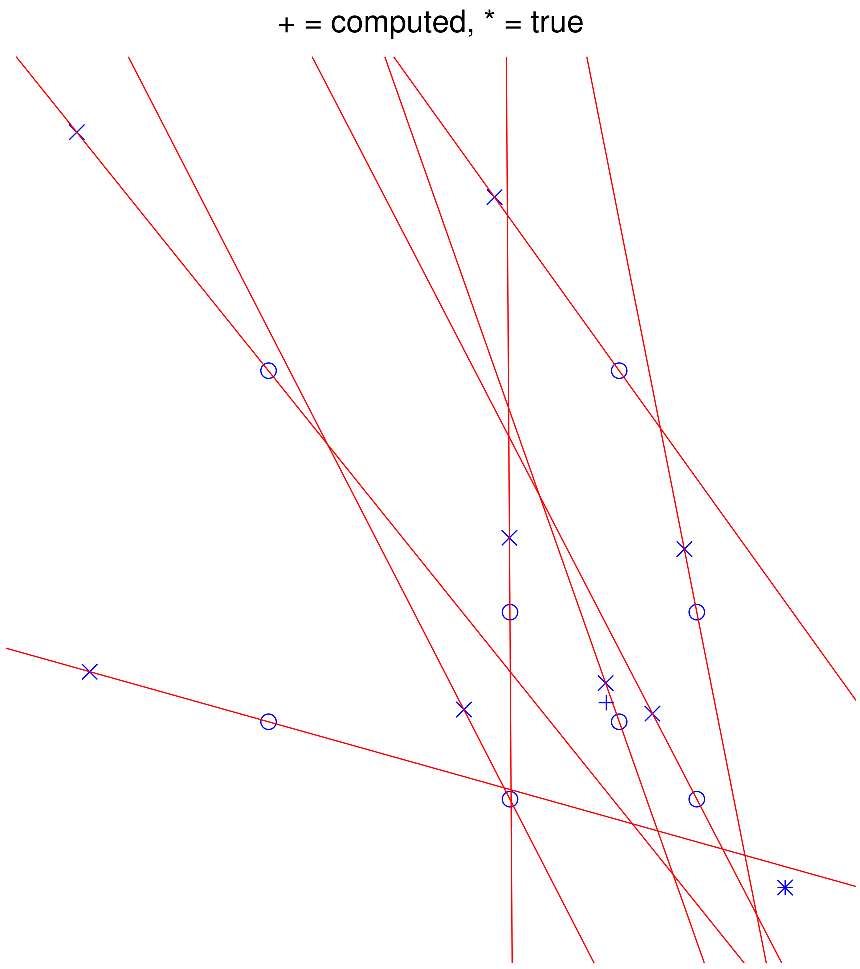

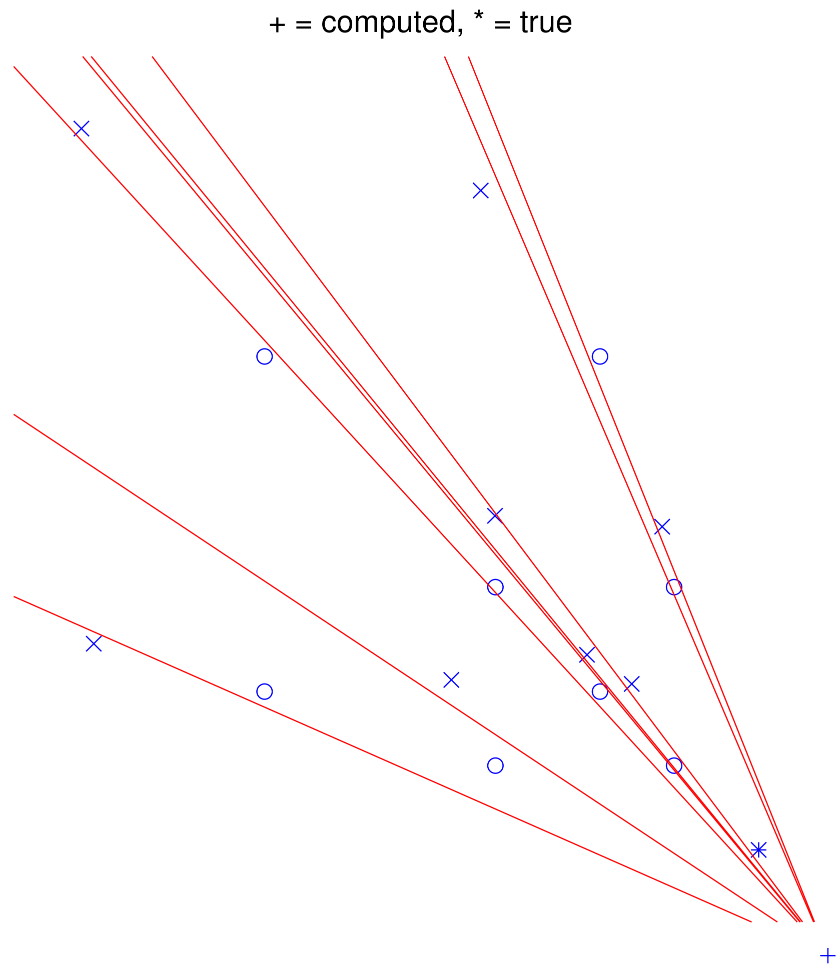

LSE vs MSE Errors, Cont’d

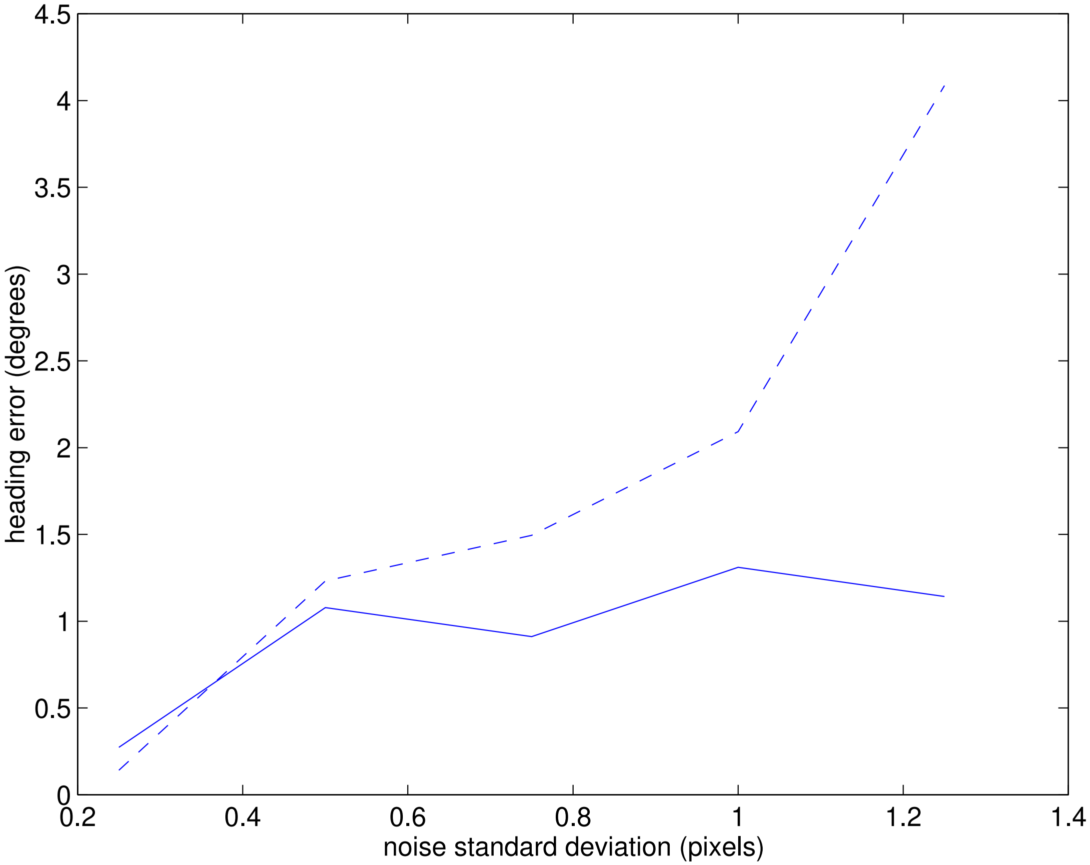

MLE and Error Propagation

- for linear problems,
  MLE + Gaussian = weighted Least Squares
- MLE = correct error propagation (but easier)
- for ill-conditioned problems, the difference can be vital
- the translation-only case is well-conditioned

Cannot Do the Impossible

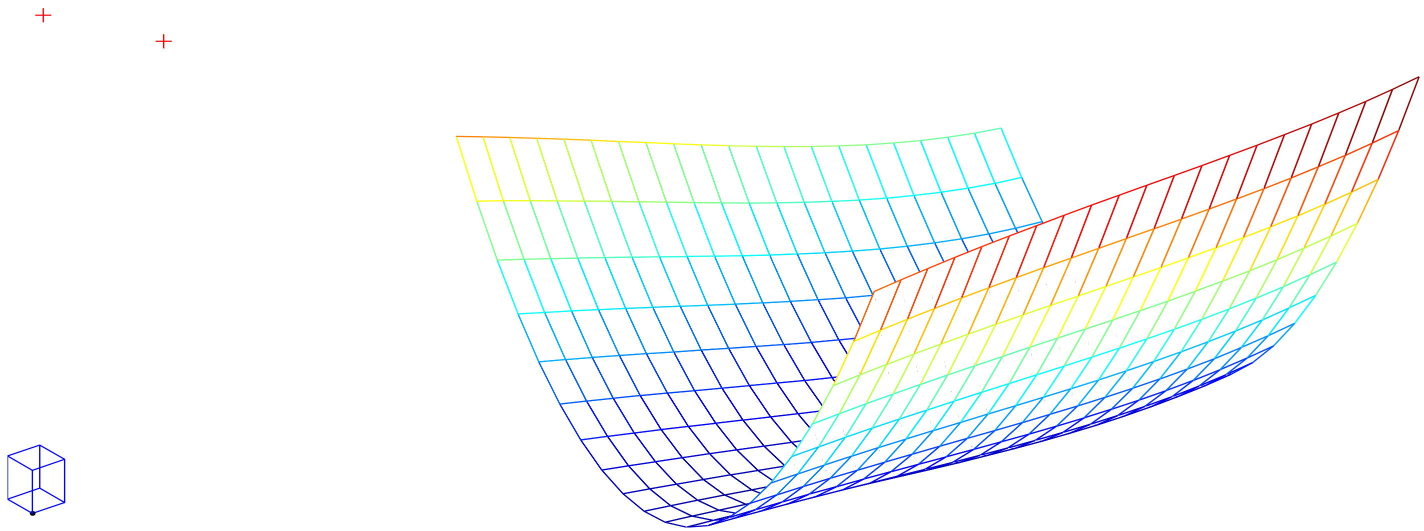

regularization?

Sensitivity

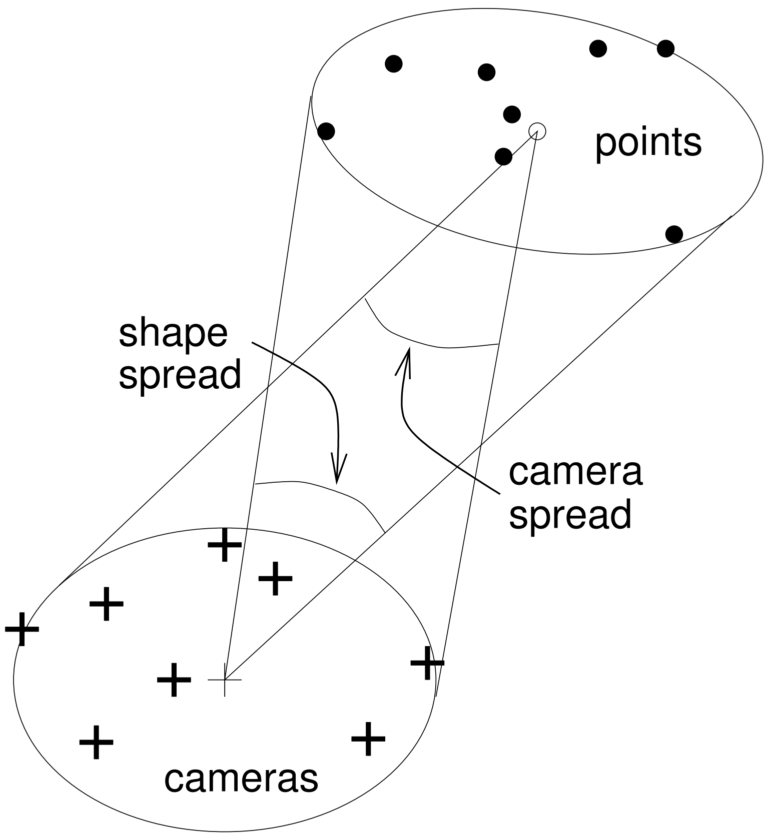

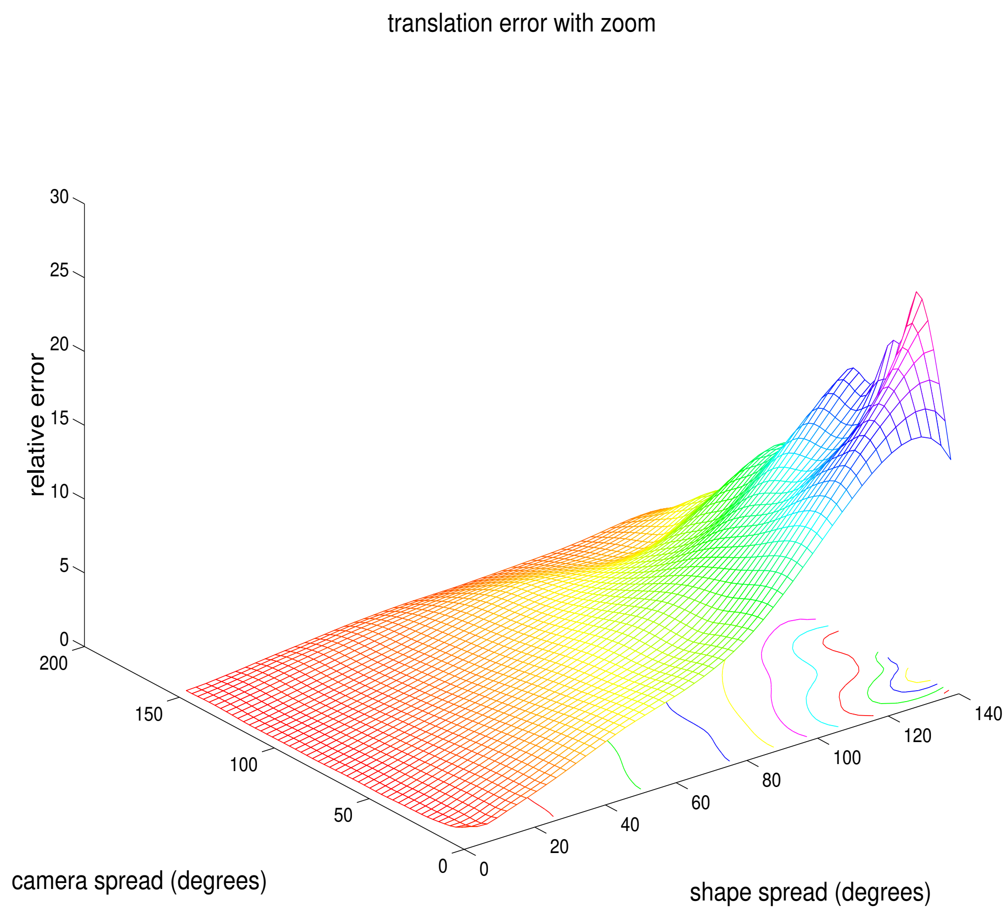

Conclusions

- no algorithm can solve impossible cases
- calibration and attention to noise statistics can make the difference between success and failure for ill-conditioned problems
- regularization, as a last resort, may be necessary
- robust statistics may be necessary for outliers