Computer Vision: Lecture 9

Carl Olsson

2020-02-18

Today's Lecture

Reconstruction and Optimization

- Objective Function: Reconstruction Error
- Principles of Local Optimization
- Least Squares Optimization
- Non-Linear Least Squares

Minimizing Reprojection Error

Under the assumption that image points are corrupted by Gaussian noise, minimize the reprojection error.

The reprojection error

In regular coordinates $(\mathbf{x} = (x, y))$ the projection is

$$\left(\frac{P^1\mathbf{X}}{P^3\mathbf{X}}, \frac{P^2\mathbf{X}}{P^3\mathbf{X}}\right),$$

P^1, P^2, P^3 are the rows of P.
The reprojection error is

$$|| \left(x - \frac{P^1\mathbf{X}}{P^3\mathbf{X}}, y - \frac{P^2\mathbf{X}}{P^3\mathbf{X}}\right) ||^2.$$

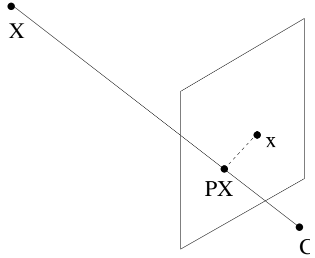

Minimizing Reprojection Error

Calibrated Structure and Motion

Given image projections $\{(x_{ij}, y_{ij})\}$ (i = point nr, j = image nr), find 3D points X_i and cameras $P_j = \begin{bmatrix} R_j & t_j \end{bmatrix}$ such that

$$\sum_{ij} || \left(x_{ij} - \frac{P_j^1 \mathbf{X}_i}{P_j^3 \mathbf{X}_i}, y_{ij} - \frac{P_j^2 \mathbf{X}_i}{P_j^3 \mathbf{X}_i} \right) ||^2,$$

is minimized.

- Complicated non linear expression.
- No closed form solution.

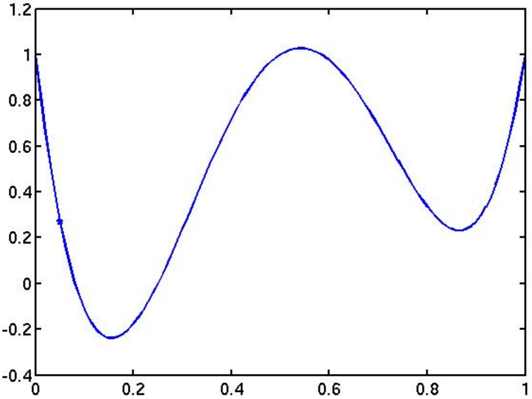

- Pick a starting point.

Minimizing Reprojection Error, Locally

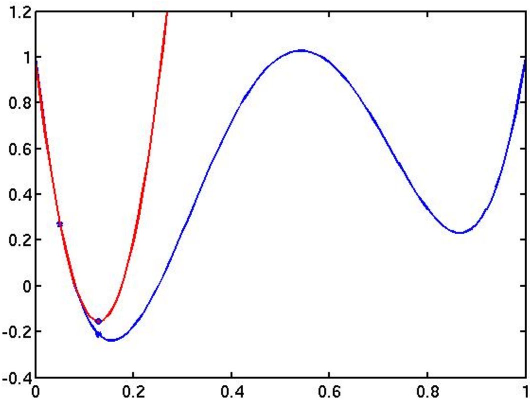

- Approximate the function using 2nd order Taylor expansion and minimize.

Minimizing Reprojection Error, Locally

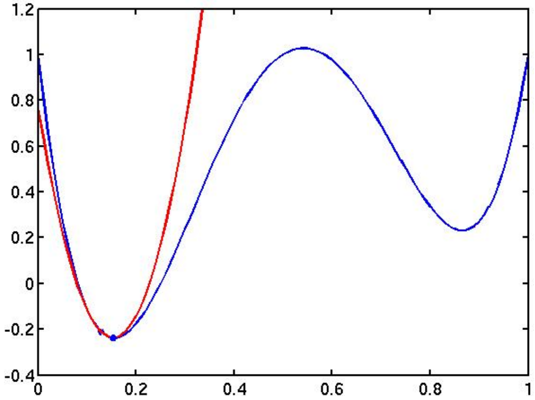

- Repeat.

Minimizing Reprojection Error, Locally

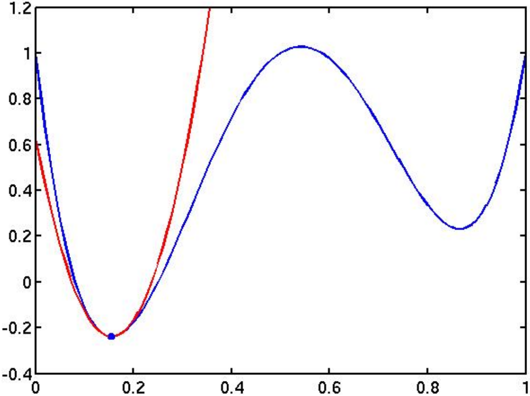

Newtons method.

Minimizing Reprojection Error, Locally

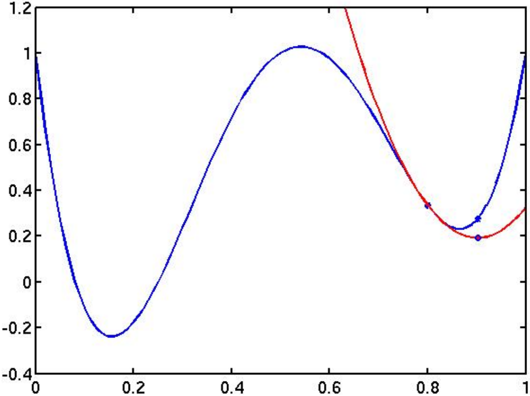

Different starting point.

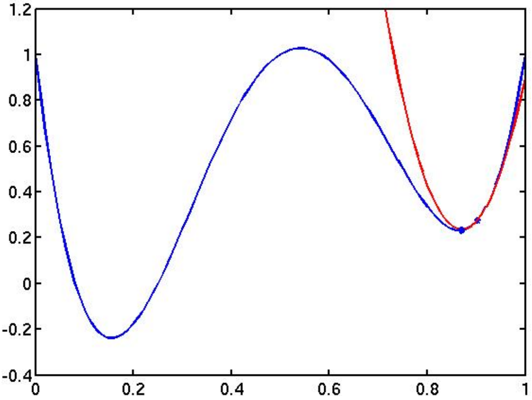

Different starting point.

Minimizing Reprojection Error, Locally

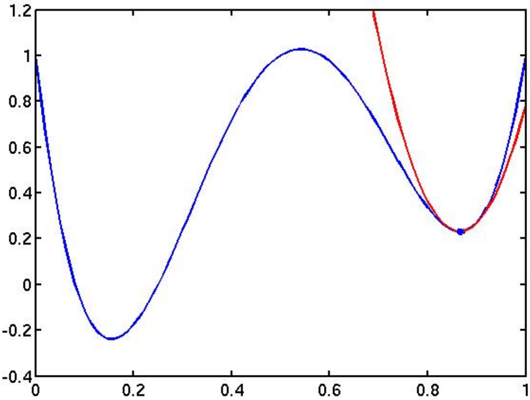

Leads to local minimum.

Minimizing Reprojection Error, Locally

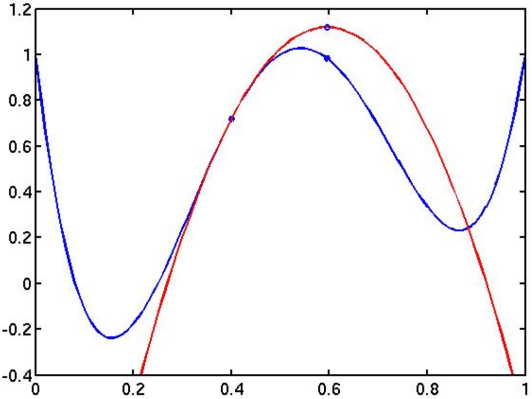

Third starting point, leads to local maximum.

Minimizing Reprojection Error, Locally

Why not just sample the function?
One dimensional function:

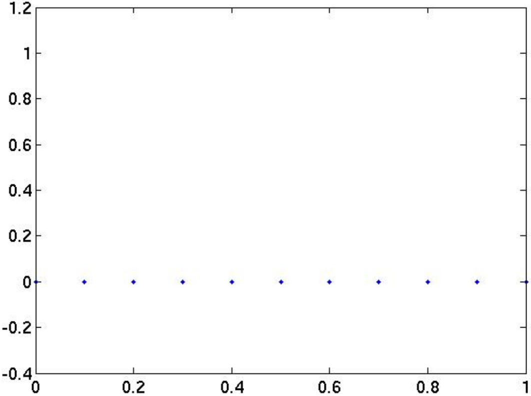

Minimizing Reprojection Error, Locally

Why not just sample the function?
One dimensional function:

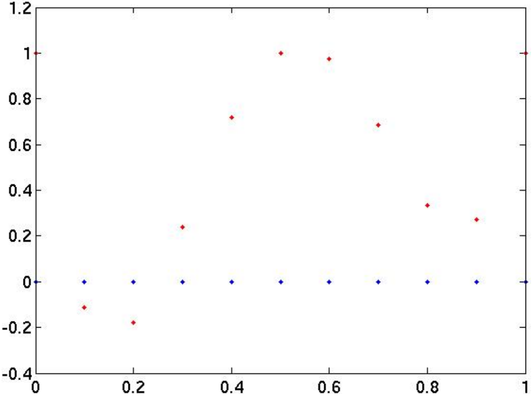

Sample 10 points, pick lowest value. Probably works.

Why not just sample the function?
Two dimensional function:

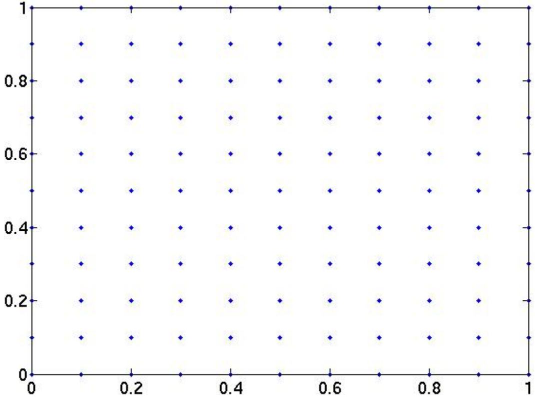

10^2 samples.

Minimizing Reprojection Error, Locally

Why not just sample the function?
Three dimensional function:

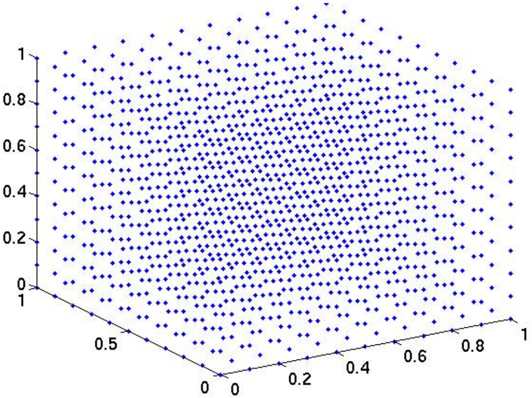

10^3 samples.

Minimizing Reprojection Error, Locally

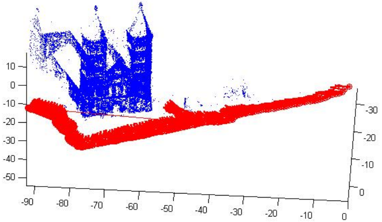

How many variables do we have?

The cathedral dataset:

- 480 camera matrices $[R_i \quad t_i]$.
  Rotation part 3 dof, translation part 3 dof.
  Totally: $480(3+3) = 2880$.
- 91178 3D points.
  3 dof each.
  Totally: $91178 \cdot 3 = 273534$

Local Optimization

See lecture notes.

Steepest Descent

Non-linear least squares:

$$\min_{v} \|r(v)\|^2$$

Linear approximation of residuals:

$$r(v) \approx r(v_0) + J(v_0)(v - v_0).$$

Line-search along the direction

$$d = -\frac{J(v_0)^T r(v_0)}{|J(v_0)^T r(v_0)|}$$

Demonstration...

Gauss-Newton

Non-linear least squares:

$$\min_v \|r(v)\|^2$$

Linear approximation of residuals:

$$r(v) \approx r(v_0) + J(v_0)(v - v_0).$$

Line-search along the direction

$$\min \|r(v_0) + J(v_0)d\|^2.$$

Demonstration...

Levenberg-Marquard

Non-linear least squares:

$$\min_{v} \|r(v)\|^2$$

Linear approximation of residuals:

$$r(v) \approx r(v_0) + J(v_0)(v - v_0).$$

Line-search along the direction

$$\min \|r(v_0) + J(v_0)d\|^2 + \lambda \|d\|^2.$$

Demonstration...