RANSAC
Robust Model Estimation
From Data Contaminated By Outliers

Lecturer: Jiří Matas

Authors: Ondřej Chum, Jiří Matas, Ondřej Drbohlav

Czech Technical University in Prague

http://cmp.felk.cvut.cz

Last update: March 2021

Talk outline

- Standard Single Class Single Instance Fitting Problem (SCSI)

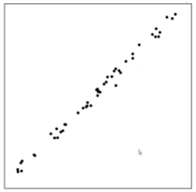 $\Rightarrow$ 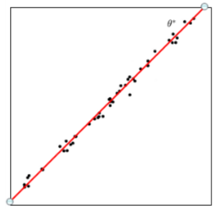

- Robust Single Class Single Instance Fitting Problem (R-SCSI)

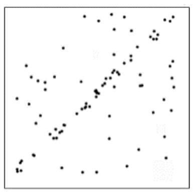 $\Rightarrow$ 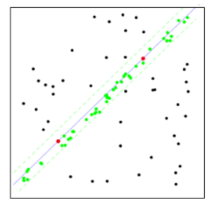

- Single Class Multiple Instance Fitting Problem (SCMI)

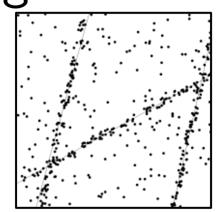 $\Rightarrow$ 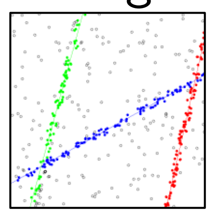

- Multiple Class Multiple Instance Fitting Problem (MCMI)

 $\Rightarrow$ 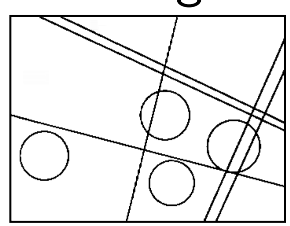

Single/Multi-Class Single/Multi-Instance Fitting Applications

- detection of geometric primitives

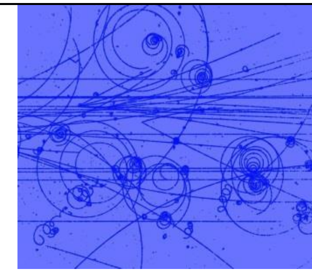

- epipolar geometry estimation
- detection of planar surfaces

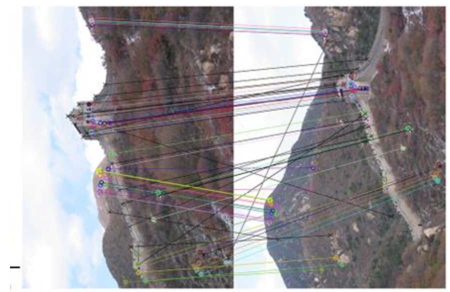

- multiple motion segmentation

- Interpretation of lidar scans

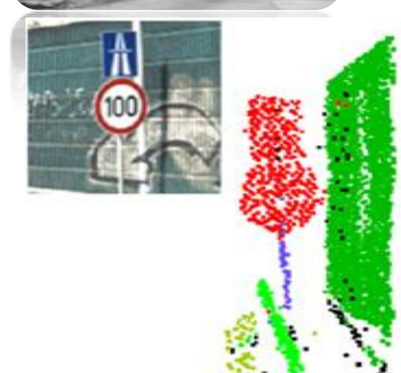

What is RANSAC?

- RANSAC = RANdom SAmple Consensus
- M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. CACM, 24(6):381–395, June 1981.

- **Example**: Finding a line in 2D
  - Not all input points are on the line.
  - Finding a line implicitly divides points to **inliers** (=those on a line) and **outliers** (=those not on a line)
  - Due to noise, "on a line" actually means inside a narrow strip around the line

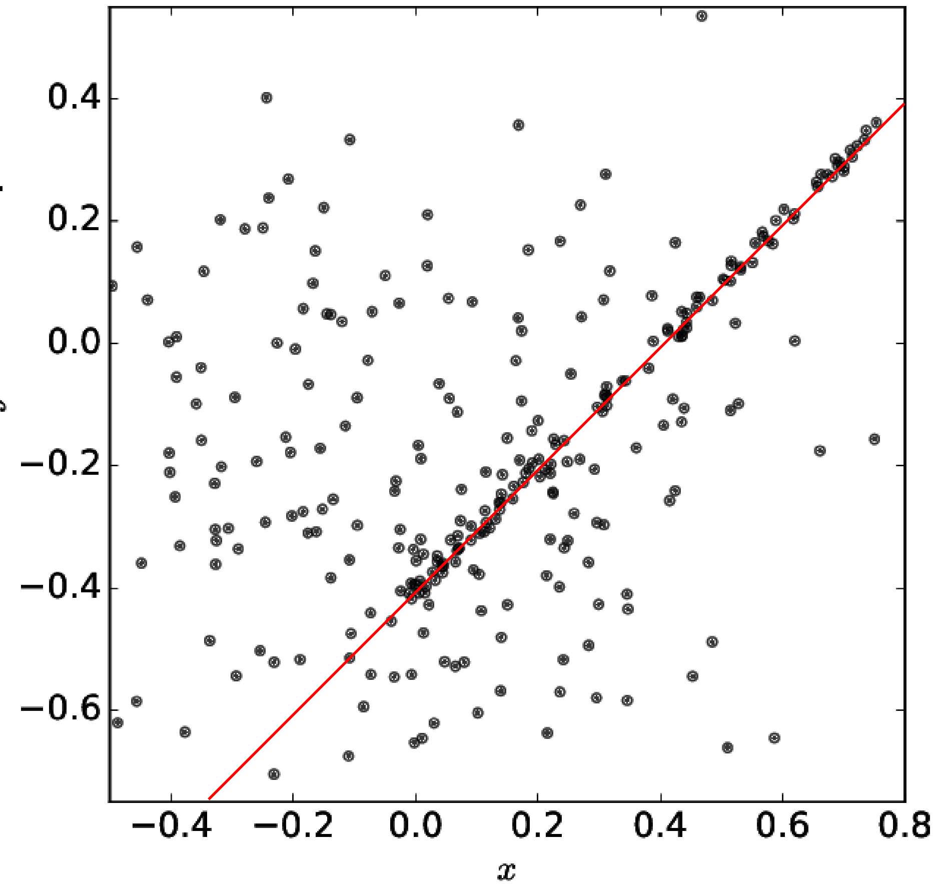

Example: Line Fitting

First, let us introduce a line parametrization and define the "strip around the line" formally:

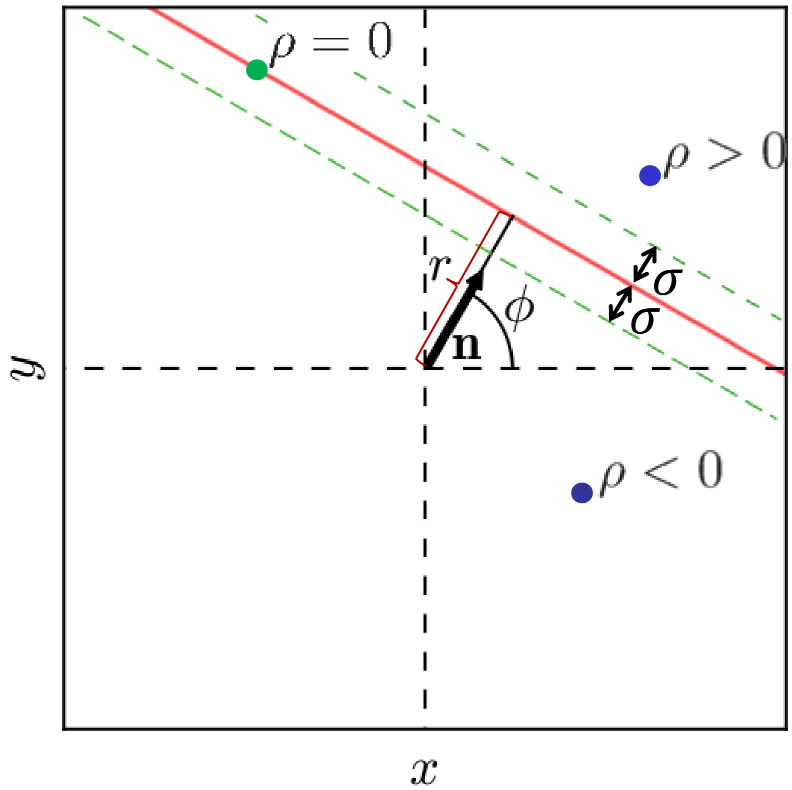

Note: $\mathbf{n} = (\cos\phi, \sin\phi)$
(thus $\|\mathbf{n}\| = 1$)

- Line parameters: $\phi \in [0, \pi[$, $r \in \mathbb{R}$
- Point $\mathbf{x} = (x, y)$ on the line:

$$x \cos\phi + y \sin\phi = r$$
$$\Leftrightarrow \mathbf{x} \cdot (\cos\phi, \sin\phi) = r$$

- Signed distance $\rho(\mathbf{p})$ of point $\mathbf{p}$ from the line:

$$\rho(\mathbf{p}) = \mathbf{p} \cdot (\cos\phi, \sin\phi) - r$$

- Point $\mathbf{p}$ inside a strip of half-width σ:

$$|\rho(\mathbf{p})| \leq \sigma$$

Line Fitting, Inliers Only: Easy!

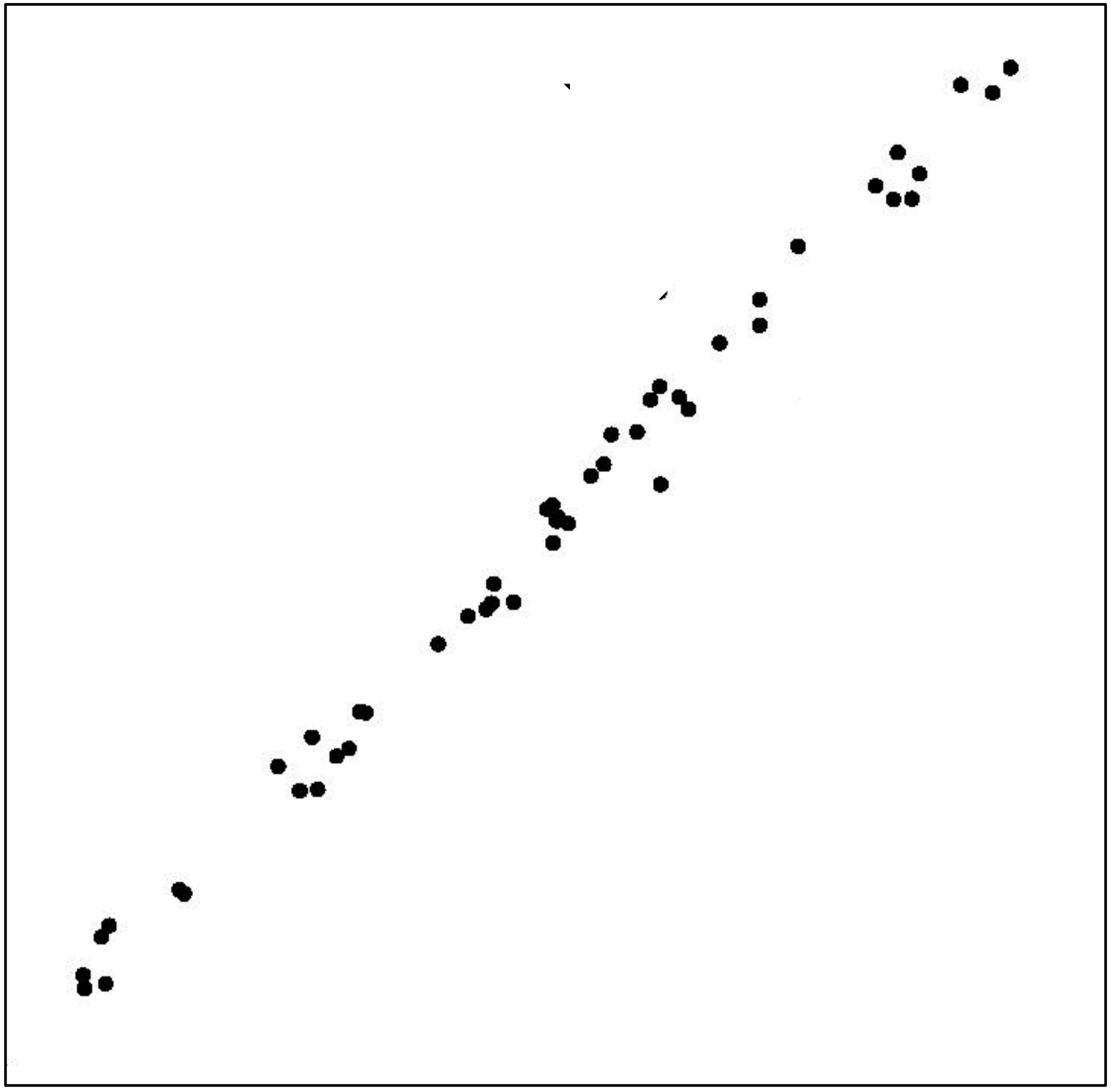

Data points

$$\mathcal{X} = \{\mathbf{x}_j, j = 1, 2, ..., N_p\}$$

$$(\mathbf{x}_j \in \mathbb{R}^2)$$

Find the line which

"best fits" these points.

Line Fitting, Inliers Only: Easy!

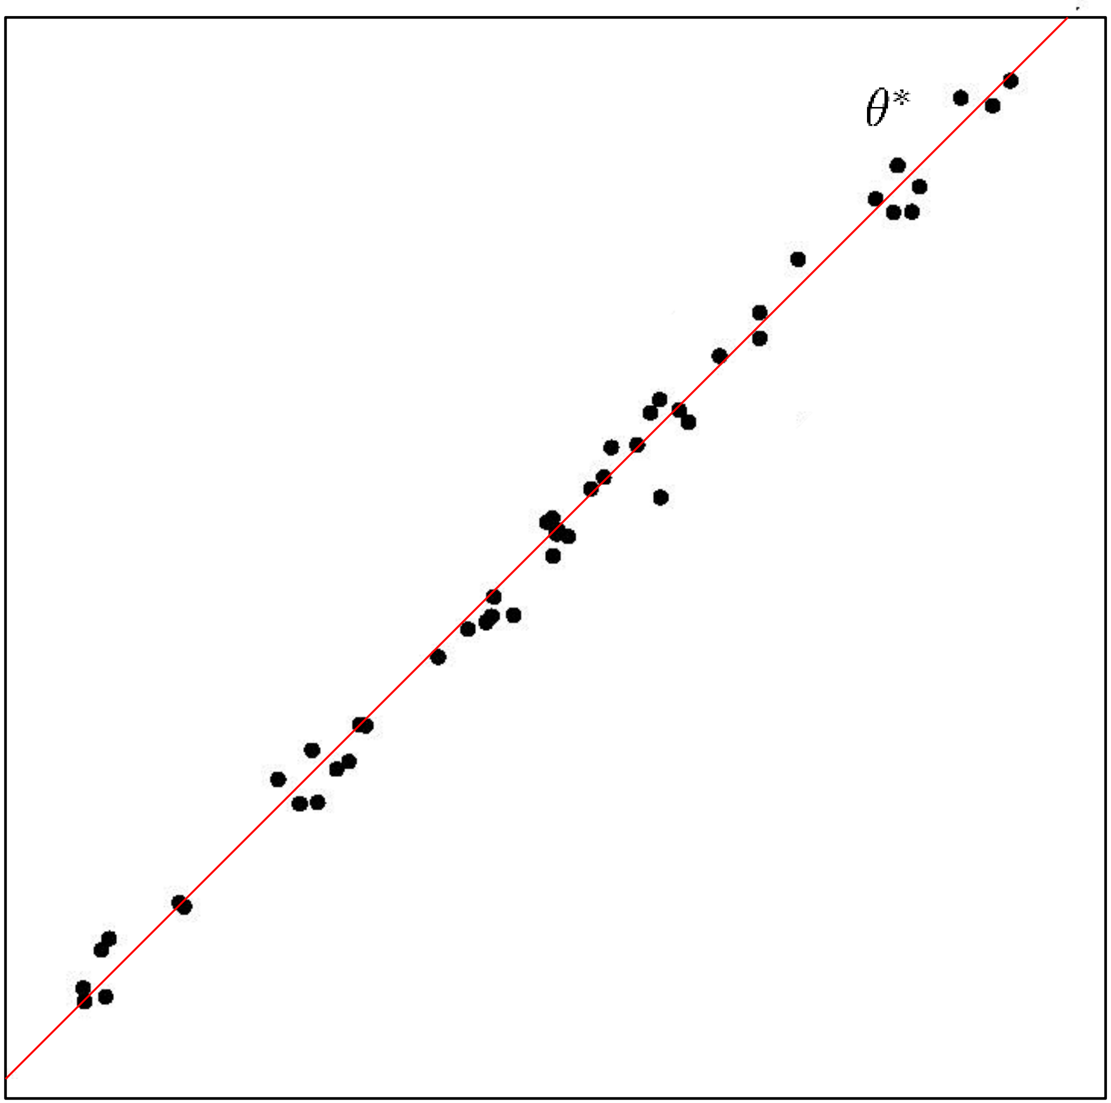

Data points

$$\mathcal{X} = \{\mathbf{x}_j, j = 1, 2, ..., N_p\}$$
$$(\mathbf{x}_j \in \mathbb{R}^2)$$

Find the line which "best fits" these points.

Optimization: Find best line with parameters θ^* :

$$\theta^* = \underset{\theta}{\text{argmin}} \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$$

For $f_{LSQ}(\mathbf{x}, \theta) = [\rho(\mathbf{x})]^2$

This is easily solvable by Singular Value Decomposition (SVD)

Example 1

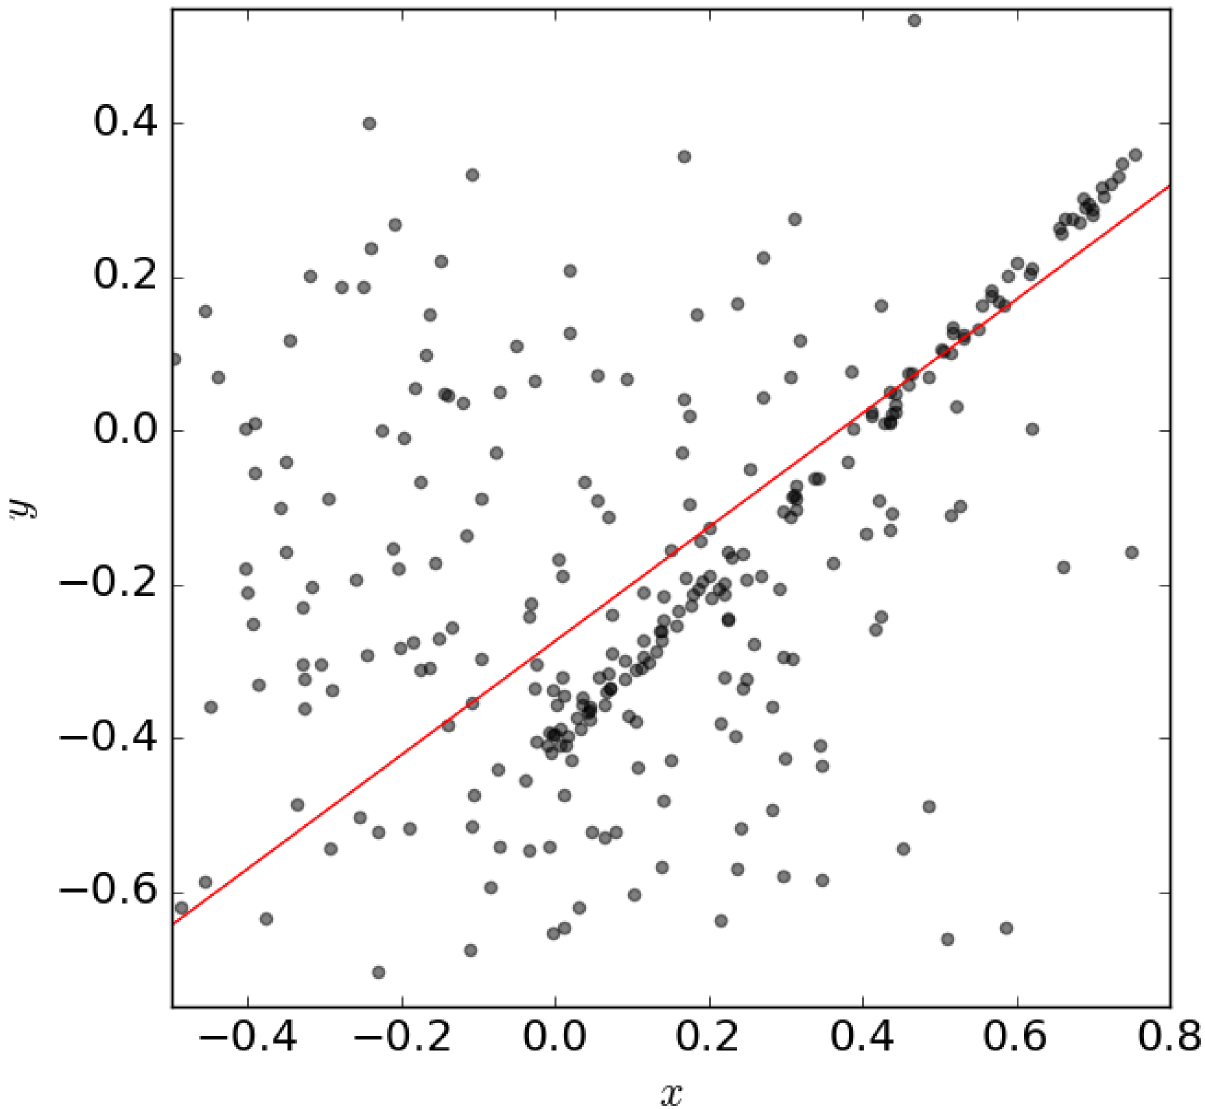

Least squares fit

General Case with Outliers, Example 2

Example 2

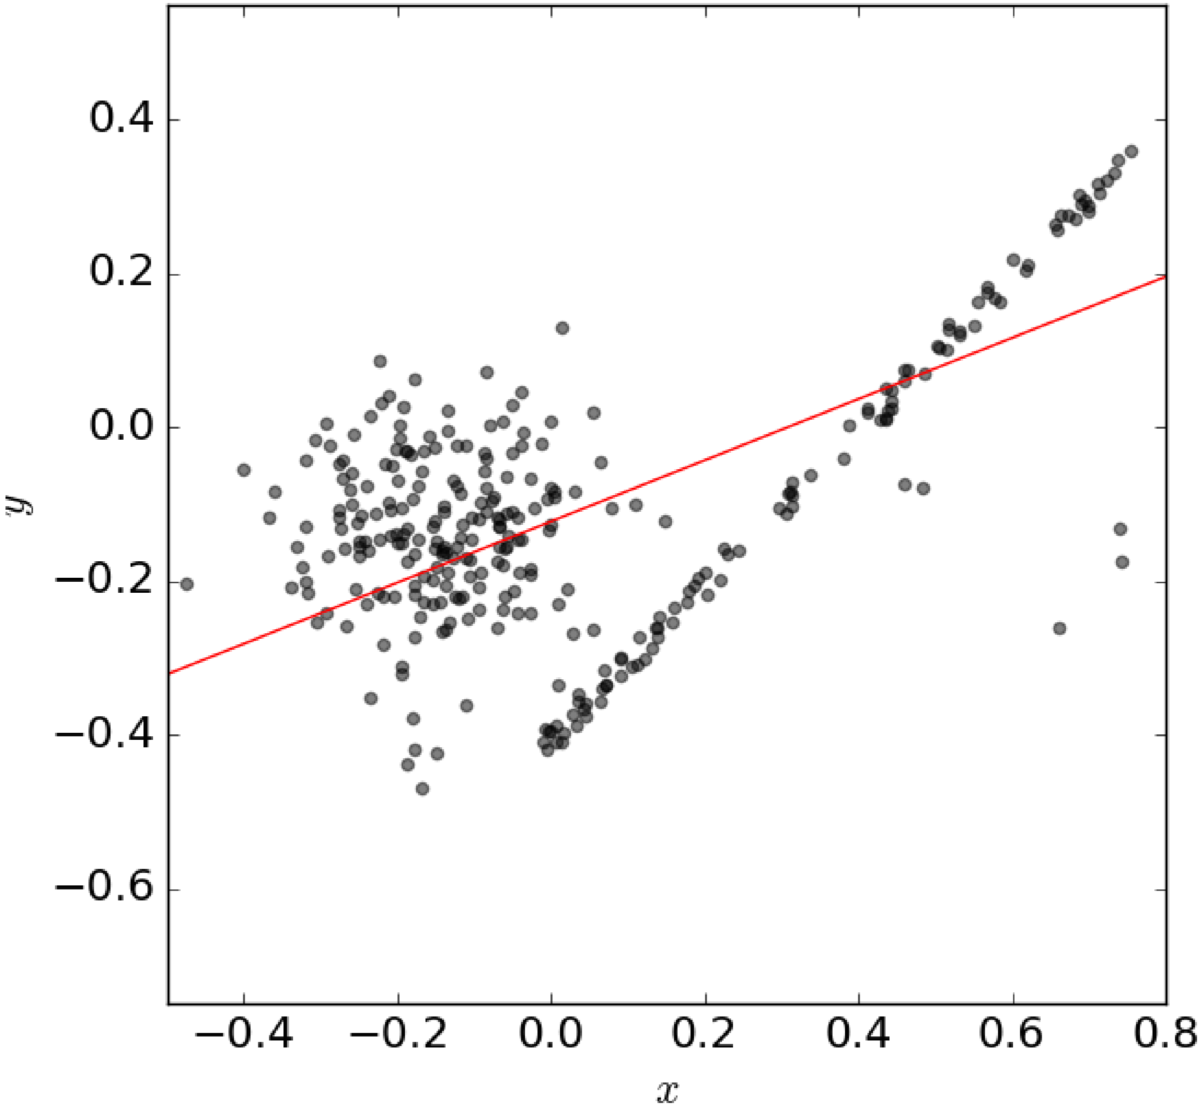

Least squares fit

General Case with Outliers, Robust Cost Function

- $\mathcal{X} = \{\mathbf{x}_j\}_{j=1}^{N_p}$ set of data points

Find:

$$\theta^* = \arg\min_{\theta} \sum_{\mathbf{x}\in\mathcal{X}} f(\mathbf{x}, \theta)$$

$$\theta = (r, \phi)$$

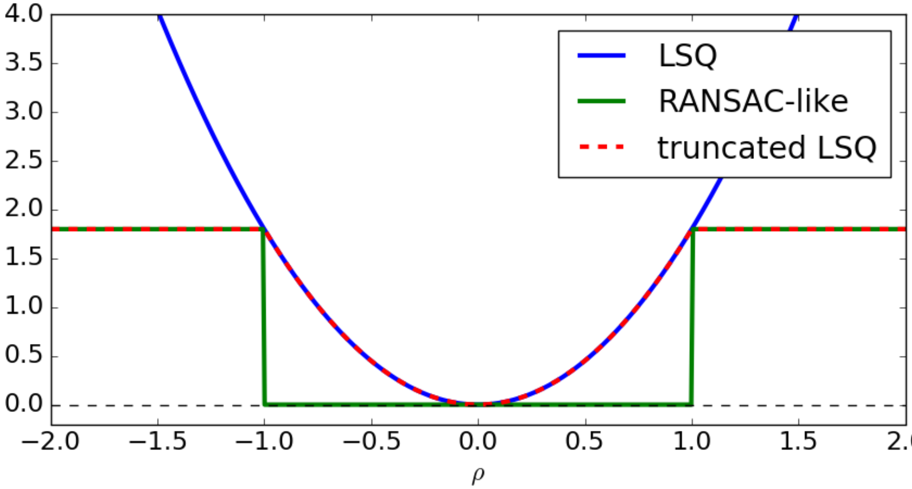

- No outliers: $f_{LSQ}(\mathbf{x}, \theta) = [\rho(\mathbf{x})]^2$

- **For robust fitting, use instead**:

$$f_{\mathsf{RANSAC}}(\mathbf{x}, \theta) = \begin{cases} 0, & \text{if } |\rho(\mathbf{x})| \leq \text{threshold } \sigma \\ \text{const}, & \text{otherwise} \end{cases}$$

- Such cost function is non-convex (and the optimization task is to minimize the number of outliers)

- How to find optimal line parameters?

RANdom SAmple Consensus - RANSAC

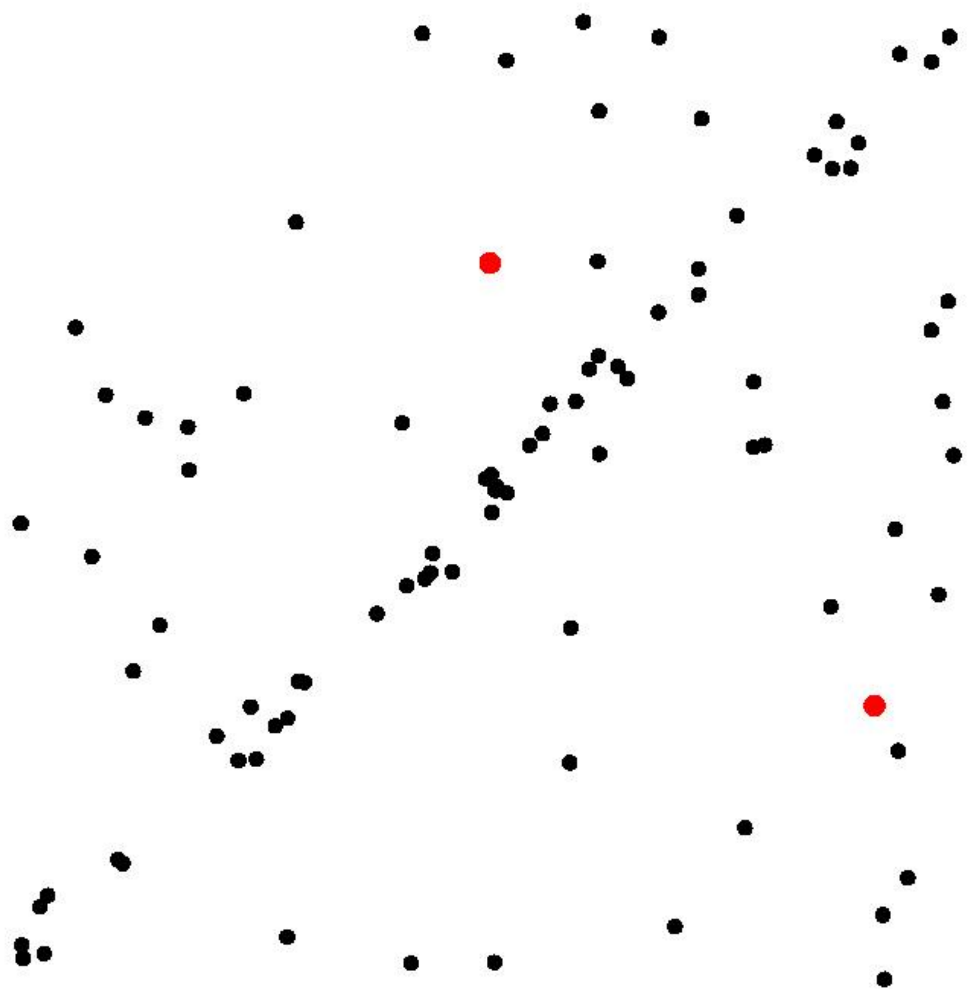

Select sample of m points at random (here m=2)

RANSAC

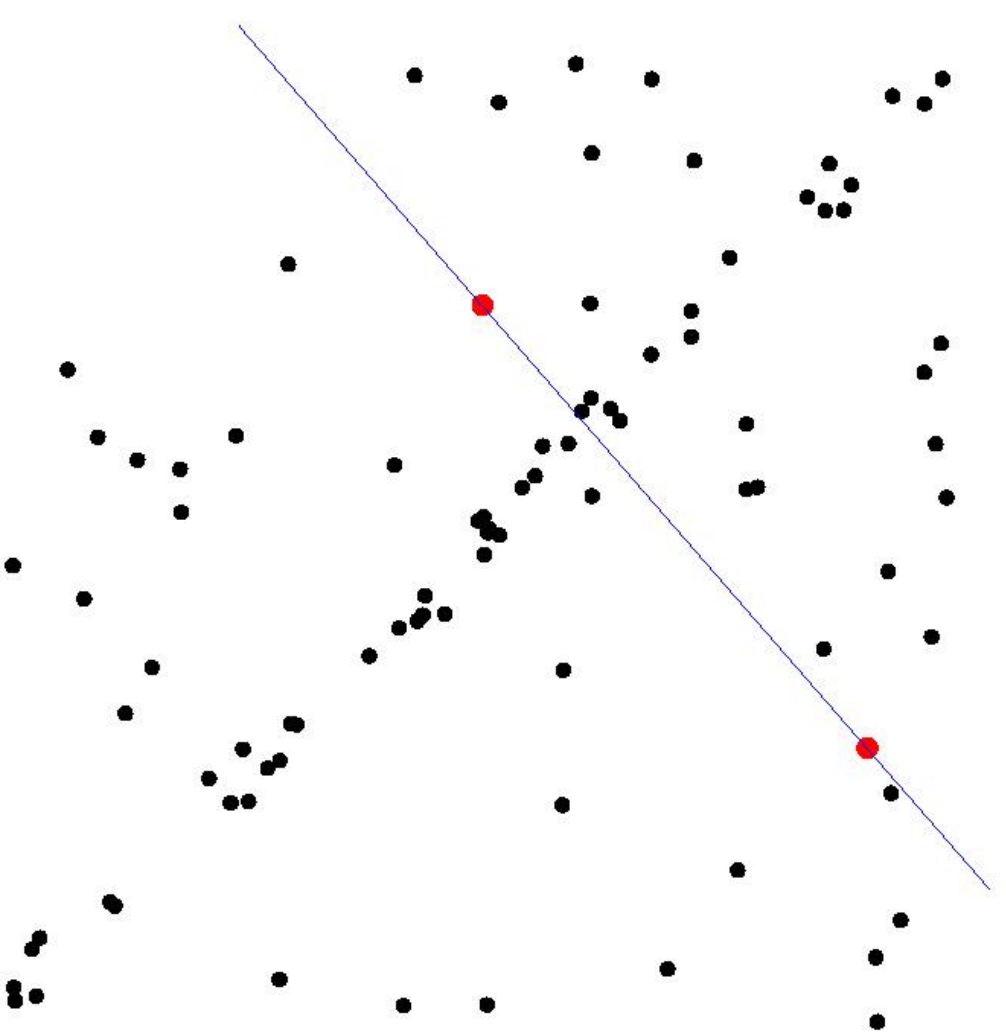

Select sample of m points at random

Estimate model parameters from the data in the sample

RANSAC

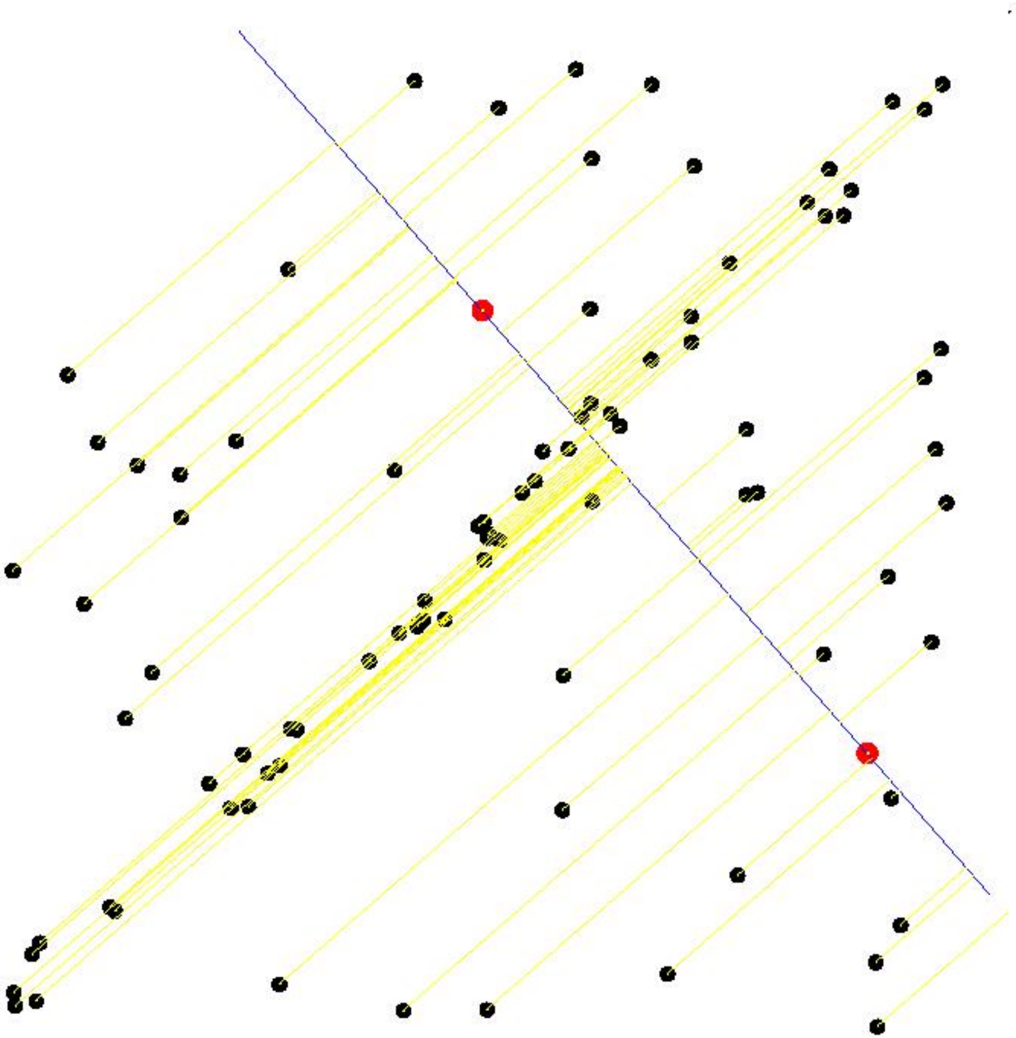

Select sample of m points at random

Estimate model parameters from the data in the sample

Evaluate the distance from model for each data point

RANSAC

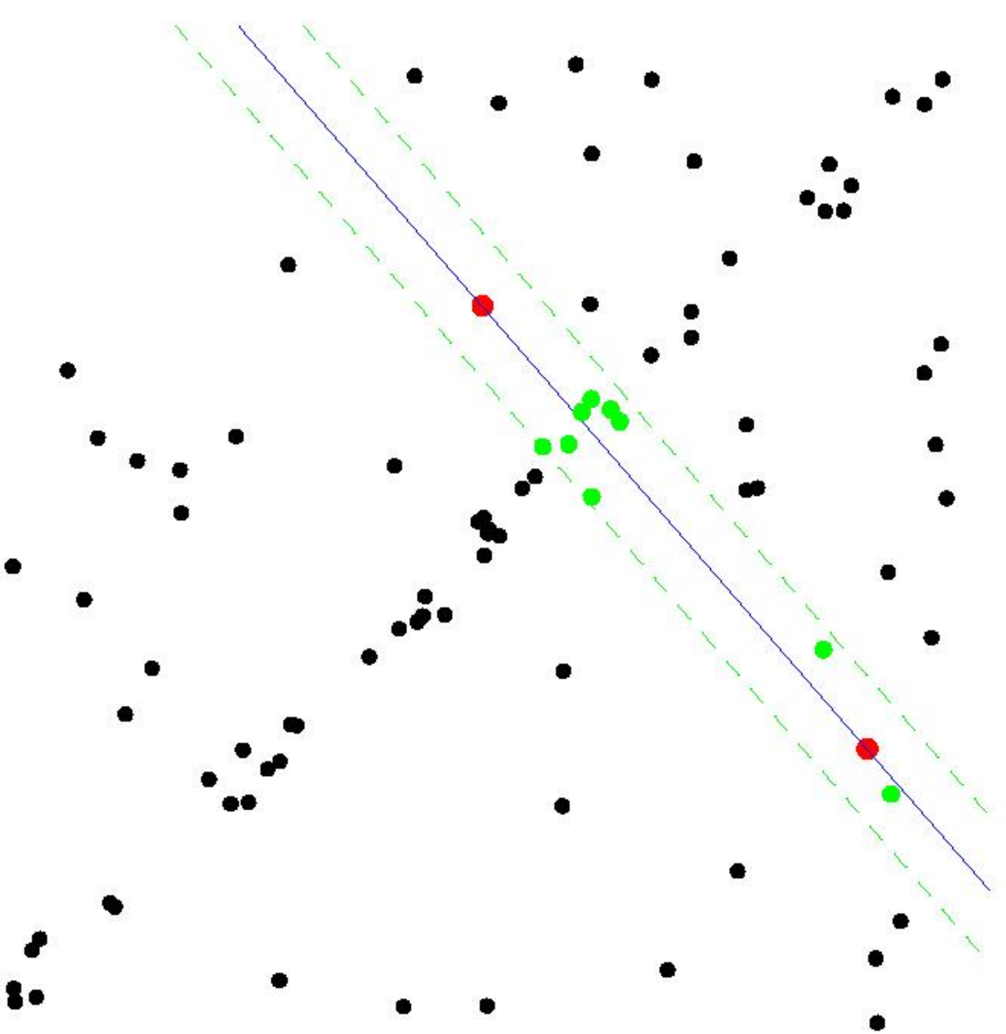

Select sample of m points at random

Estimate model parameters from the data in the sample

Evaluate the distance from model for each data point

Select data that support the current hypothesis

RANSAC

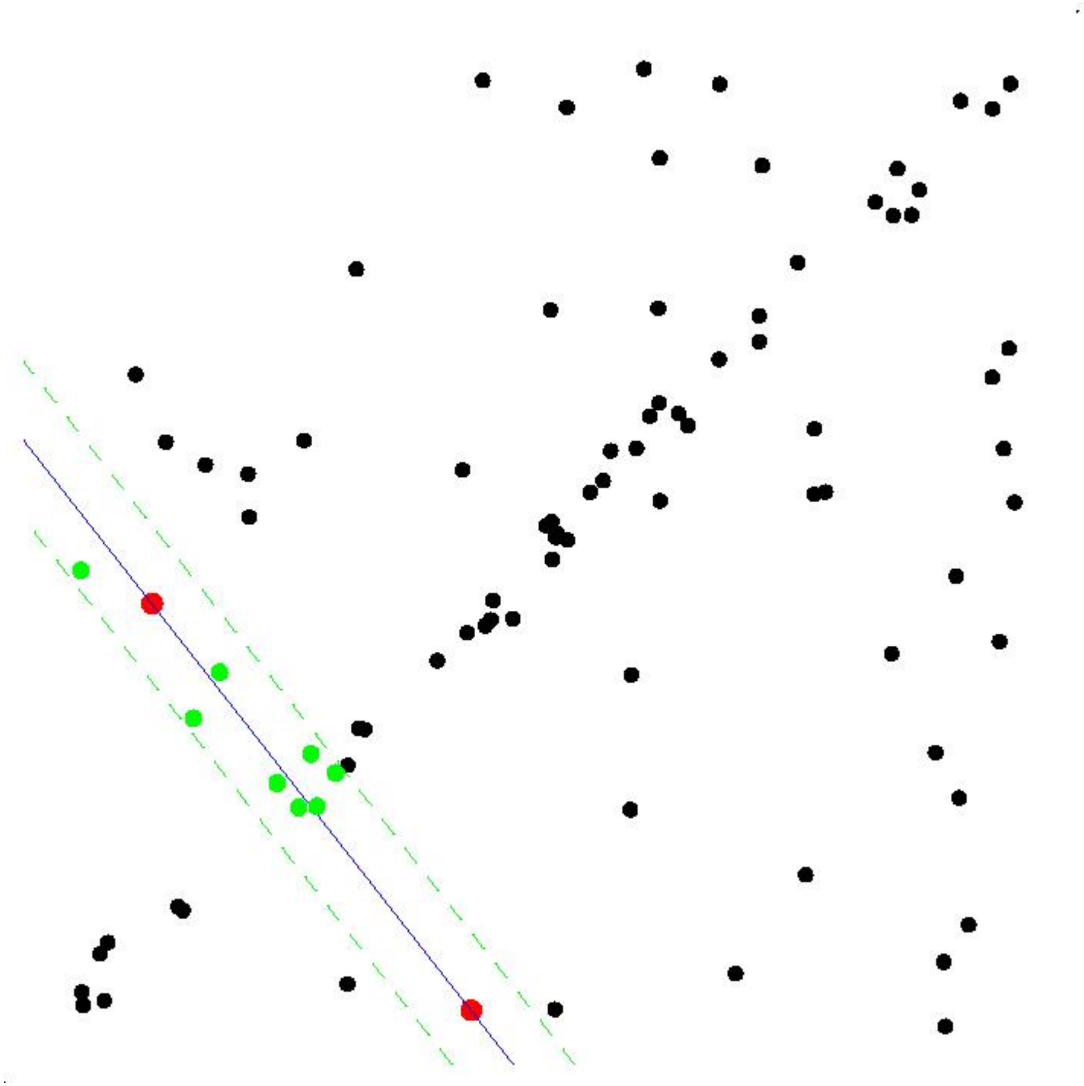

Select sample of m points at random

Estimate model parameters from the data in the sample

Evaluate the distance from model for each data point

Select data that support the current hypothesis

Repeat sampling

RANSAC

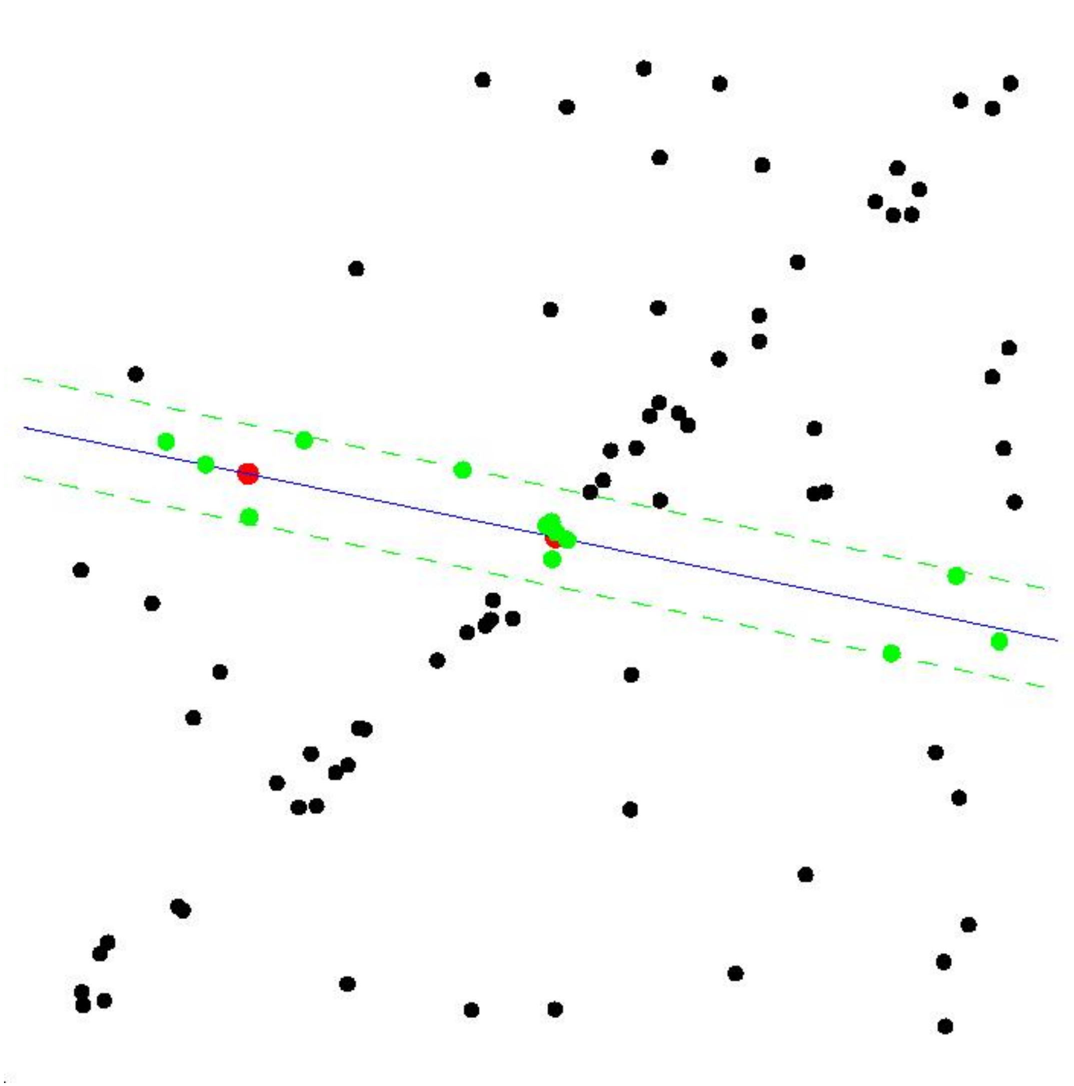

Select sample of m points at random

Estimate model parameters from the data in the sample

Evaluate the distance from model for each data point

Select data that support the current hypothesis

Repeat sampling

RANSAC

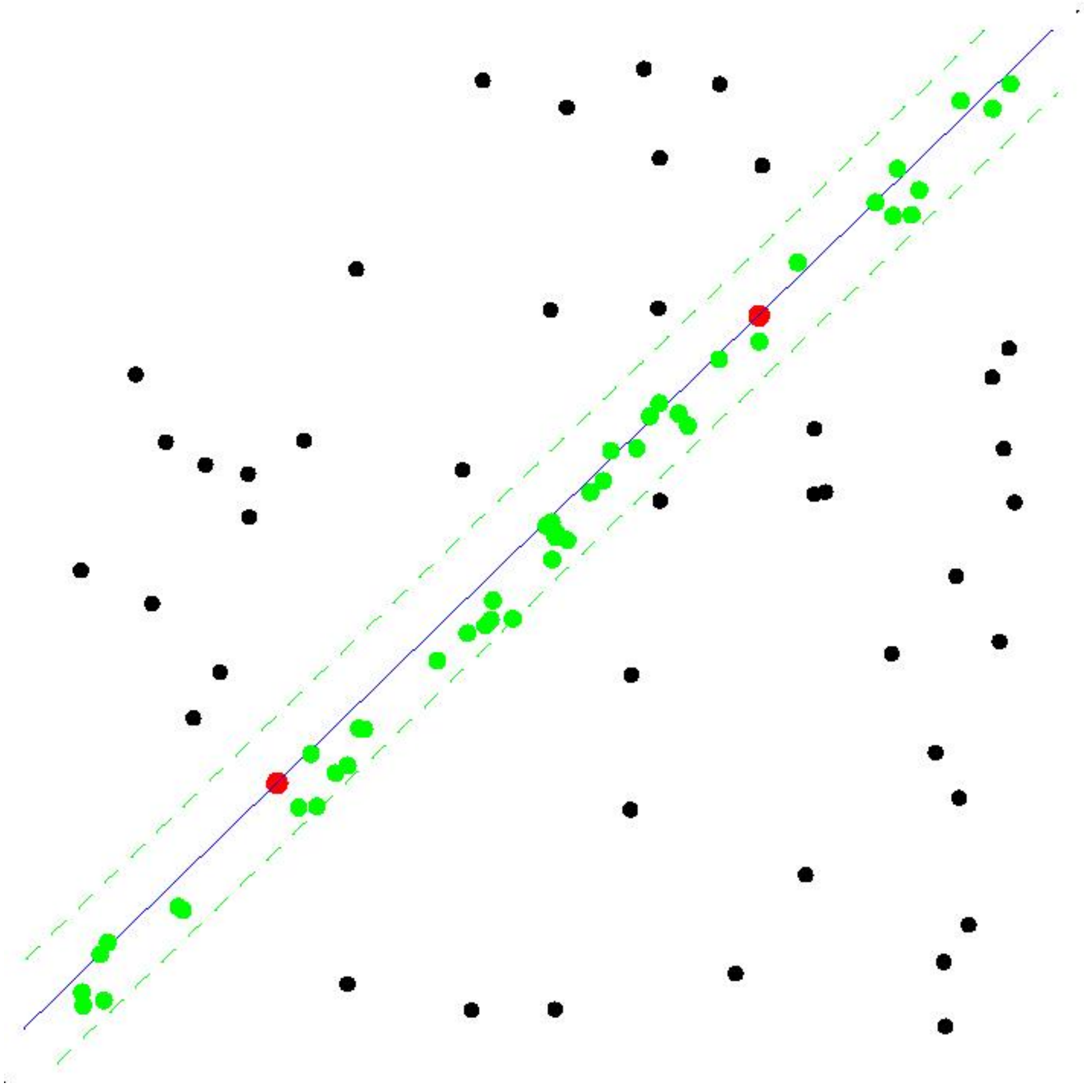

Select sample of m points at random

Estimate model parameters from the data in the sample

Evaluate the distance from model for each data point

Select data that support the current hypothesis

Repeat sampling

RANSAC [Fischler and Bolles 1981]

Input: $\mathcal{X} = \{\mathbf{x}_j\}_{j=1}^{N}$ — data points

$e(S) = \theta$ — estimates *model parameters* θ given sample $S \subseteq \mathcal{X}$

$$f(\mathbf{x}, \theta) = \begin{cases} 0, & \text{if distance to model} \leq \text{threshold } \sigma \\ 1, & \text{otherwise} \end{cases}$$

Cost function for single data point $\mathbf{x}$

$\Rightarrow J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ is #outliers

η – required confidence in the solution, σ – outlier threshold

Output: θ^* parameter of the model minimizing the cost function

$iter \leftarrow 0,\ J^* \leftarrow \infty$
repeat
 Select *random* $S \subseteq \mathcal{X}$ (sample size $m = |S|$) — SAMPLING
 Estimate parameters $\theta = e(S)$
 Evaluate $J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ — VERIFICATION
 If $J(\theta) < J^*$ then — SO-FAR-THE-BEST
  $\theta^* \leftarrow \theta,\ J^* \leftarrow J(\theta)$
 $iter \leftarrow iter + 1$
until P(better solution exists) $= f(|\mathcal{X}|, J^*, iter) < 1 - \eta$
Compute θ^* from all inliers $\mathcal{X}_{in}$: $\theta^* \leftarrow$ LocalOptimization$(\mathcal{X}_{in}, \theta^*)$

RANSAC – how many samples?

- N Number of points
- Q Number of inliers, $Q = N - J^*$
- m Size of sample
- $\epsilon = Q/N$ Inlier ratio

Probability of all-inlier (uncontaminated) sample:

$$P(\text{inlier sample}) = \frac{\binom{Q}{m}}{\binom{N}{m}} = \frac{Q(Q-1)...(Q-m+1)}{N(N-1)...(N-m+1)} \approx \epsilon^m$$

Mean time for hitting all-inliers sample is proportional to $1/P$.

RANSAC – how many samples?

- How about this formulation:
  - Set the number of samples k such that ***at least one*** pair of points from the line has been hit with probability larger than η
  - Equivalently ... such that ***no*** pair of points from the line has been hit with probability lower than 1 - η
- Q Number of inliers, $Q = N - J^*$
- m Size of sample
- $\epsilon = Q/N$ Inlier ratio

Probability of all-inlier (uncontaminated) sample:

$$P(\text{inlier sample}) = \frac{\binom{Q}{m}}{\binom{N}{m}} = \frac{Q(Q-1)...(Q-m+1)}{N(N-1)...(N-m+1)} \approx \epsilon^m$$

The required confidence in solution:

$$P(\text{bad model } k \text{ times}) = (1 - P(\text{inlier sample}))^k < 1 - \eta$$

Finding the solution with confidence η therefore requires at least k samples:

$$k \geq \log(1-\eta) \;/\; \log\left(1 - \epsilon^m\right)$$

RANSAC termination – how many samples?

- m Size of sample
- $\epsilon = Q/N$ Inlier ratio
- η Confidence
- k required number of samples

Probability of all-inlier (uncontaminated) sample:

$$k \geq \log(1-\eta) \;/\; \log\left(1-\epsilon^m\right)$$

RANSAC termination - How many samples?

Inlier ratio $\epsilon = Q/N$ [%]

Size of the sample m

m	15%	20%	30%	40%	50%	70%
2	130 200 300	73 110 170	32 49 73	17 26 40	10 16 24	4 6 10
3	890 1400 2000	370 570 860	110 170 250	45 70 100	22 34 52	7 11 16
4	5900 9100 $1.4 \cdot 10^4$	1900 2900 4300	370 570 850	120 180 270	46 71 110	11 17 25
8	$1.2 \cdot 10^7$ $1.8 \cdot 10^7$ $2.7 \cdot 10^7$	$1.2 \cdot 10^6$ $1.8 \cdot 10^6$ $2.7 \cdot 10^6$	$4.6 \cdot 10^4$ $7.0 \cdot 10^4$ $1.1 \cdot 10^5$	4600 7000 $1.1 \cdot 10^4$	770 1200 1800	50 78 120
12	$2.3 \cdot 10^{10}$ $3.5 \cdot 10^{10}$ $5.3 \cdot 10^{10}$	$7.3 \cdot 10^8$ $1.1 \cdot 10^9$ $1.7 \cdot 10^9$	$5.6 \cdot 10^6$ $8.7 \cdot 10^6$ $1.3 \cdot 10^7$	$1.8 \cdot 10^5$ $2.7 \cdot 10^5$ $4.1 \cdot 10^5$	$1.2 \cdot 10^4$ $1.9 \cdot 10^4$ $2.8 \cdot 10^4$	210 330 500
18	$2.1 \cdot 10^{15}$ $3.2 \cdot 10^{15}$ $4.8 \cdot 10^{15}$	$1.1 \cdot 10^{13}$ $1.8 \cdot 10^{13}$ $2.6 \cdot 10^{13}$	$7.7 \cdot 10^9$ $1.2 \cdot 10^{10}$ $1.8 \cdot 10^{10}$	$4.4 \cdot 10^7$ $6.7 \cdot 10^7$ $1.0 \cdot 10^8$	$7.9 \cdot 10^5$ $1.2 \cdot 10^6$ $1.8 \cdot 10^6$	1800 2800 4200
30	∞ ∞ ∞	∞ ∞ ∞	$1.3 \cdot 10^{16}$ $2.1 \cdot 10^{16}$ $3.1 \cdot 10^{16}$	$2.6 \cdot 10^{12}$ $3.1 \cdot 10^{12}$ $5.1 \cdot 10^{12}$	$3.2 \cdot 10^9$ $4.9 \cdot 10^9$ $7.4 \cdot 10^9$	$1.3 \cdot 10^5$ $2.0 \cdot 10^5$ $3.1 \cdot 10^5$
50	∞ ∞ ∞	∞ ∞ ∞	∞ ∞ ∞	∞ ∞ ∞	$3.4 \cdot 10^{15}$ $5.2 \cdot 10^{15}$ $7.8 \cdot 10^{15}$	$1.7 \cdot 10^8$ $2.6 \cdot 10^8$ $3.8 \cdot 10^8$

Computed for confidences $\eta = 0.95$ (first row in each cell), $\eta = 0.99$ (second row) and $\eta = 0.999$ (third row)

RANSAC Notes

Pros:

- extremely popular (>17000 citations in Google Scholar)
- used in many applications
- percentage of inliers not needed and not limited
- a probabilistic guarantee for the solution
- mild assumptions: σ known

Cons:

- slow if inlier ratio low
- It was observed experimentally that RANSAC takes several times longer than theoretically expected. This is due to noise –
  not every all-inlier sample generates a good hypothesis:

$$P(\text{inlier sample}) \neq P(\text{good model estimate})$$

RANSAC Variants

- **Cost function:** MLESAC, Huber loss, ...
- **Outlier threshold** σ (how to set it in advance? Or, how to avoid setting it?): Least median of Squares, MINPRAN, MAGSAC, ...
- **Correctness of the results. Degeneracy.**
  Solution: DegenSAC.
- **Accuracy** (parameters are estimated from minimal samples).
  Solution: Locally Optimized RANSAC
- **Speed:** Running time grows with number of data points, number of iterations (polynomial in the inlier ratio)
  Addressing the problem:
  R-RANSAC (Randomized evaluation), RANSAC with SPRT (WaldSAC), PROSAC

Locally Optimized RANSAC (LO-RANSAC): Problem Intro

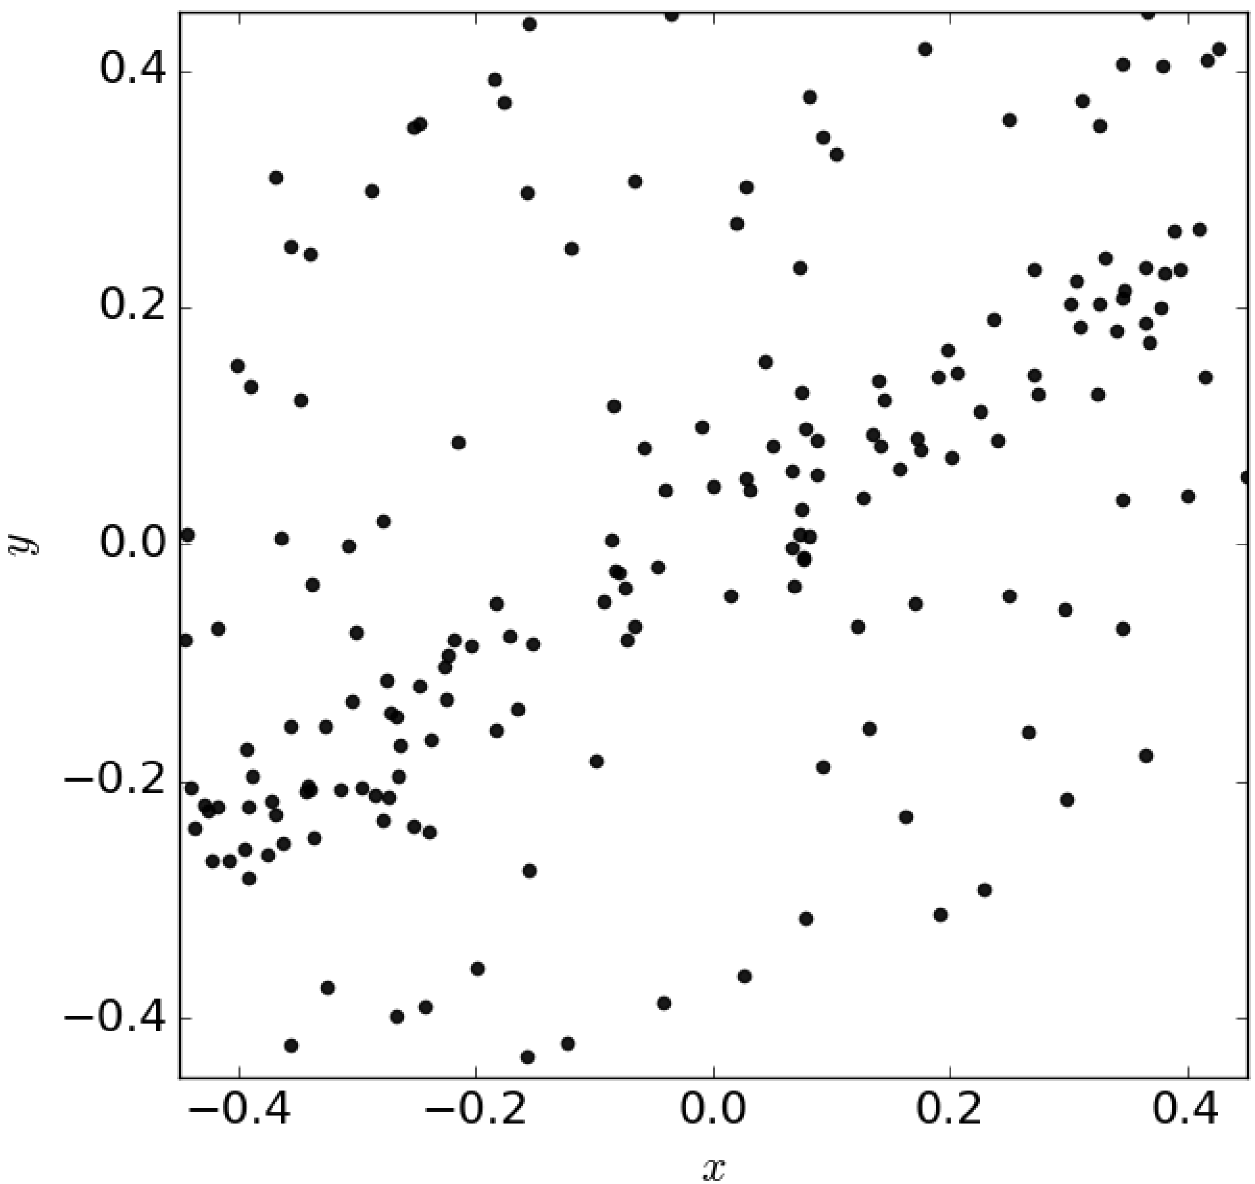

Data: 200 points

LO-RANSAC: Problem Introduction

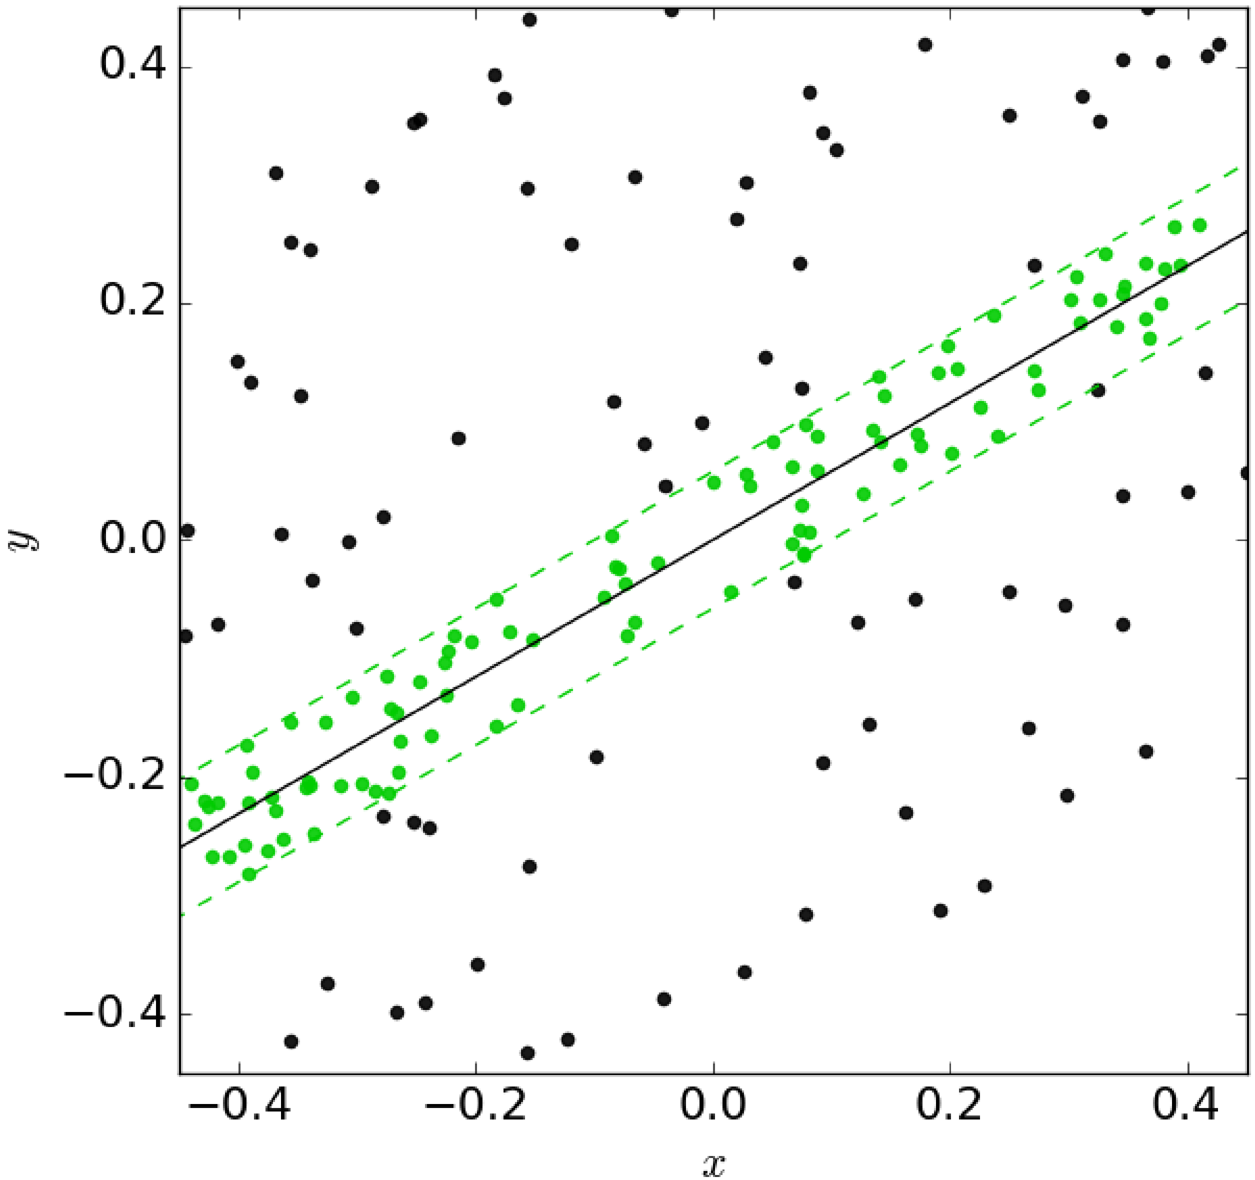

Data: 200 points

Model, 100 inliers

LO-RANSAC: Problem Introduction

For simplicity, consider only points belonging to the model (100 points)

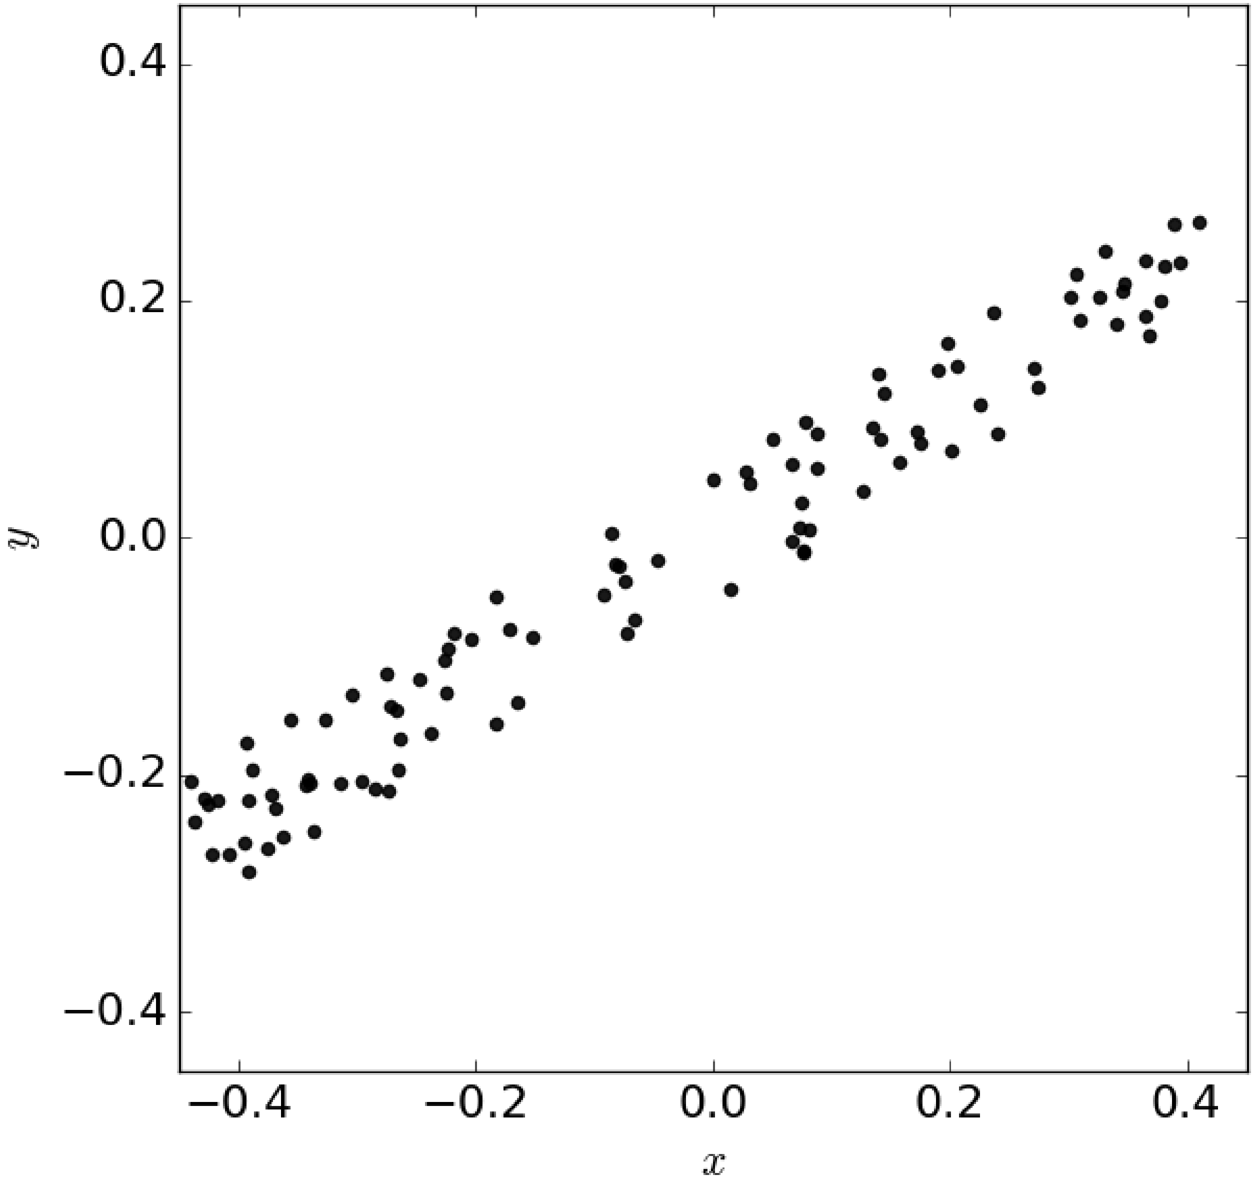

LO-RANSAC: Problem Introduction

For simplicity, consider only points belonging to the model (100 points)

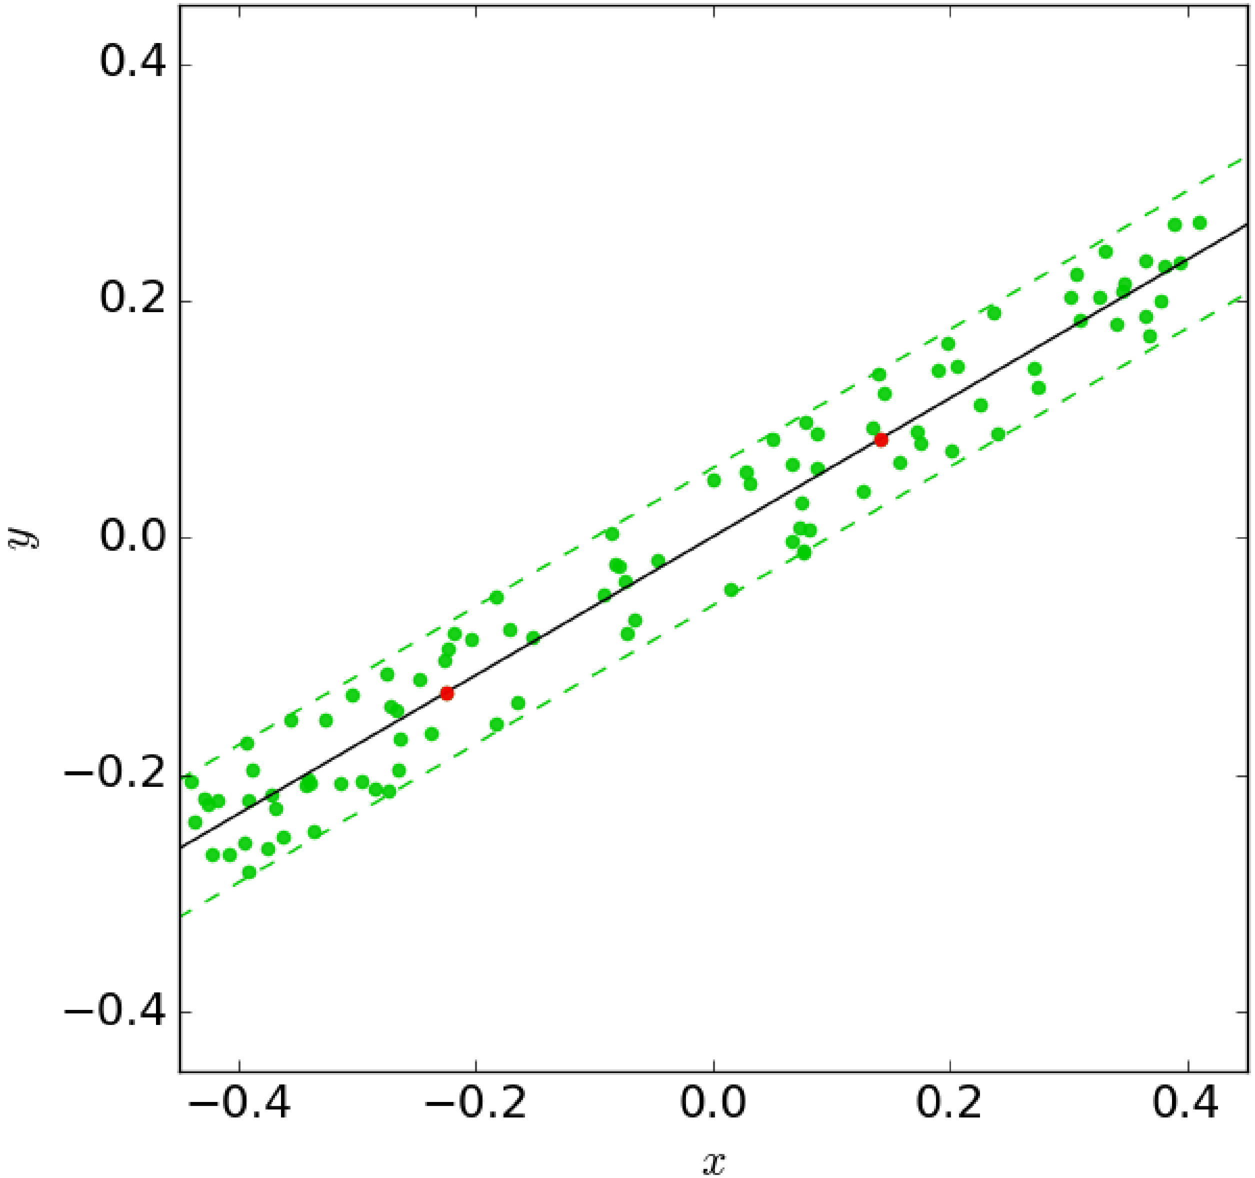

RANSAC

Hypothesis generation from 2 points

Will every two points generate the whole inlier set?

This sample:
YES. 100 inliers.

LO-RANSAC: Problem Introduction

For simplicity, consider only points belonging to the model (100 points)

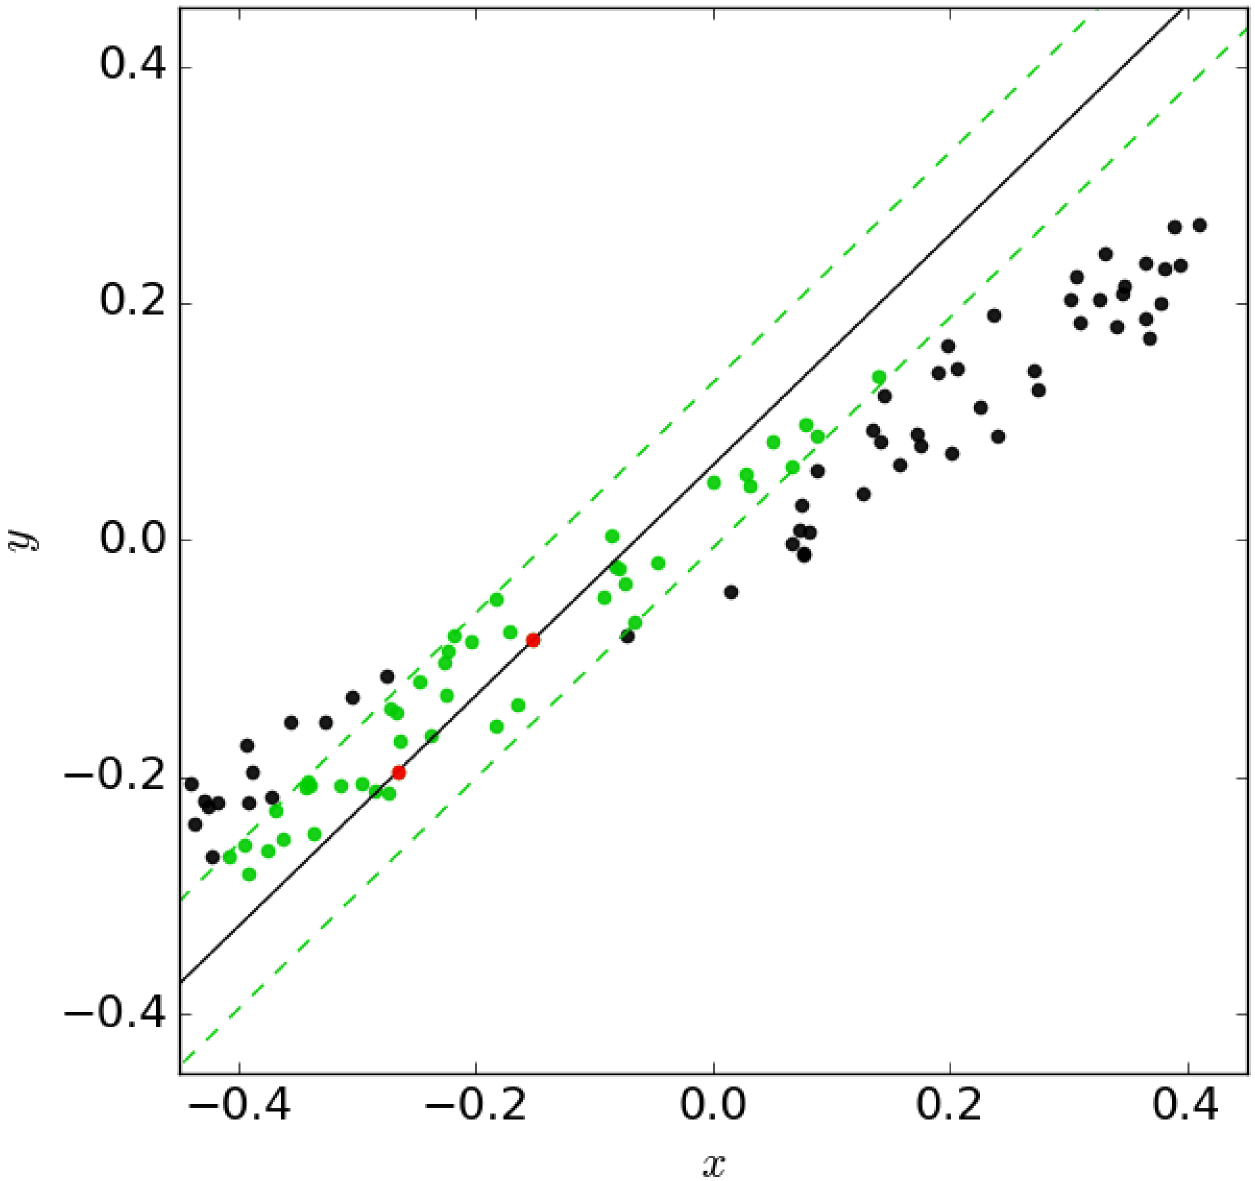

RANSAC

Hypothesis generation from 2 points

Will every two points generate the whole inlier set?

This sample:
NO. 45 inliers.

LO-RANSAC: Problem Introduction

For simplicity, consider only points belonging to model (100 points)

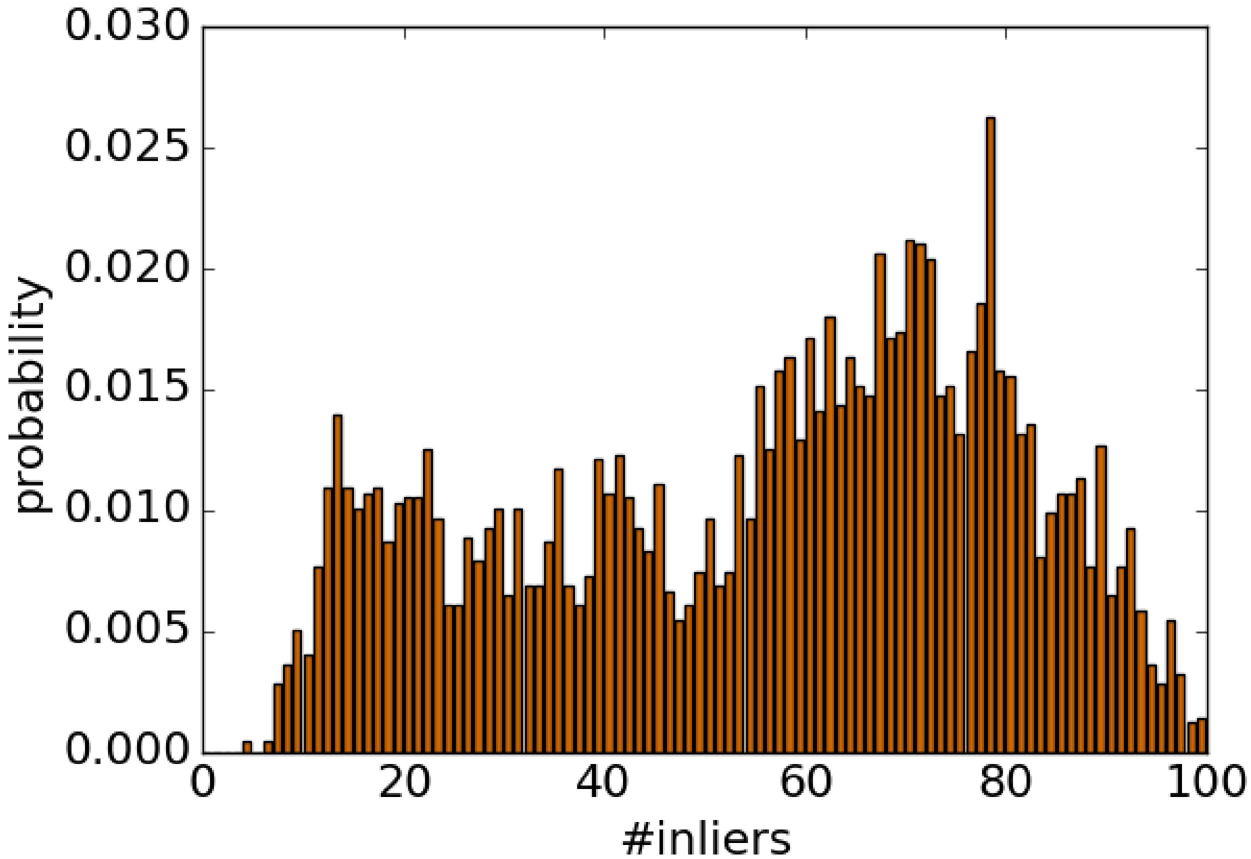

The distribution of the number of inliers obtained while randomly sampling inlier points pairs

RANSAC

Hypothesis generation from 2 points

Will every two points generate the whole inlier set?

LO-RANSAC

Input: $\mathcal{X} = \{\mathbf{x}_j\}_{j=1}^{N}$ data points

$e(S) = \theta$ estimates *model parameters* θ given sample $S \subseteq \mathcal{X}$

$$f(\mathbf{x}, \theta) = \begin{cases} 0, & \text{if distance to model } \leq \text{ threshold } \sigma \\ 1, & \text{otherwise} \end{cases}$$

Cost function for single data point $\mathbf{x}$

$\Rightarrow J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ is #outliers

η – required confidence in the solution, σ – outlier threshold

Output: θ^* parameter of the model minimizing the cost function

$iter \leftarrow 0,\ J^* \leftarrow \infty$
repeat
  Select *random* $S \subseteq \mathcal{X}$ (sample size $m = |S|$) SAMPLING
  Estimate parameters $\theta = e(S)$
  Evaluate $J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ VERIFICATION
  If $J(\theta) < J^*$ then SO-FAR-THE-BEST
    $\theta^* \leftarrow \theta,\ J^* \leftarrow J(\theta)$
  $iter \leftarrow iter + 1$
until $P(\text{better solution exists}) = f(|\mathcal{X}|, J^*, iter) < 1 - \eta$
Compute θ^* from all inliers $\mathcal{X}_{in}$: $\theta^* \leftarrow$ LocalOptimization$(\mathcal{X}_{in}, \theta^*)$

LO-RANSAC

Input: $\mathcal{X} = \{\mathbf{x}_j\}_{j=1}^{N}$ data points

$e(S) = \theta$ estimates *model parameters* θ given sample $S \subseteq \mathcal{X}$

$$f(\mathbf{x}, \theta) = \begin{cases} 0, & \text{if distance to model } \leq \text{ threshold } \sigma \\ 1, & \text{otherwise} \end{cases}$$

Cost function for single data point $\mathbf{x}$

$\Rightarrow J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ is #outliers

η – required confidence in the solution, σ – outlier threshold

Output: θ^* parameter of the model minimizing the cost function

1: $iter \leftarrow 0,\ J^* \leftarrow \infty$
2: **repeat**
3: Select *random* $S \subseteq \mathcal{X}$ (sample size $m = |S|$) SAMPLING
4: Estimate parameters $\theta = e(S)$
5: Evaluate $J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ VERIFICATION
6: If $J(\theta) < J^*$ then SO-FAR-THE-BEST
$\theta^* \leftarrow$ LocalOptimization$(\mathcal{X}_{in}, \theta)$, $J^* \leftarrow J(\theta^*)$
7: $iter \leftarrow iter + 1$
8: **until** $P(\text{better solution exists}) = f(|\mathcal{X}|, J^*, iter) < 1 - \eta$
9: gone

LO-RANSAC: Example

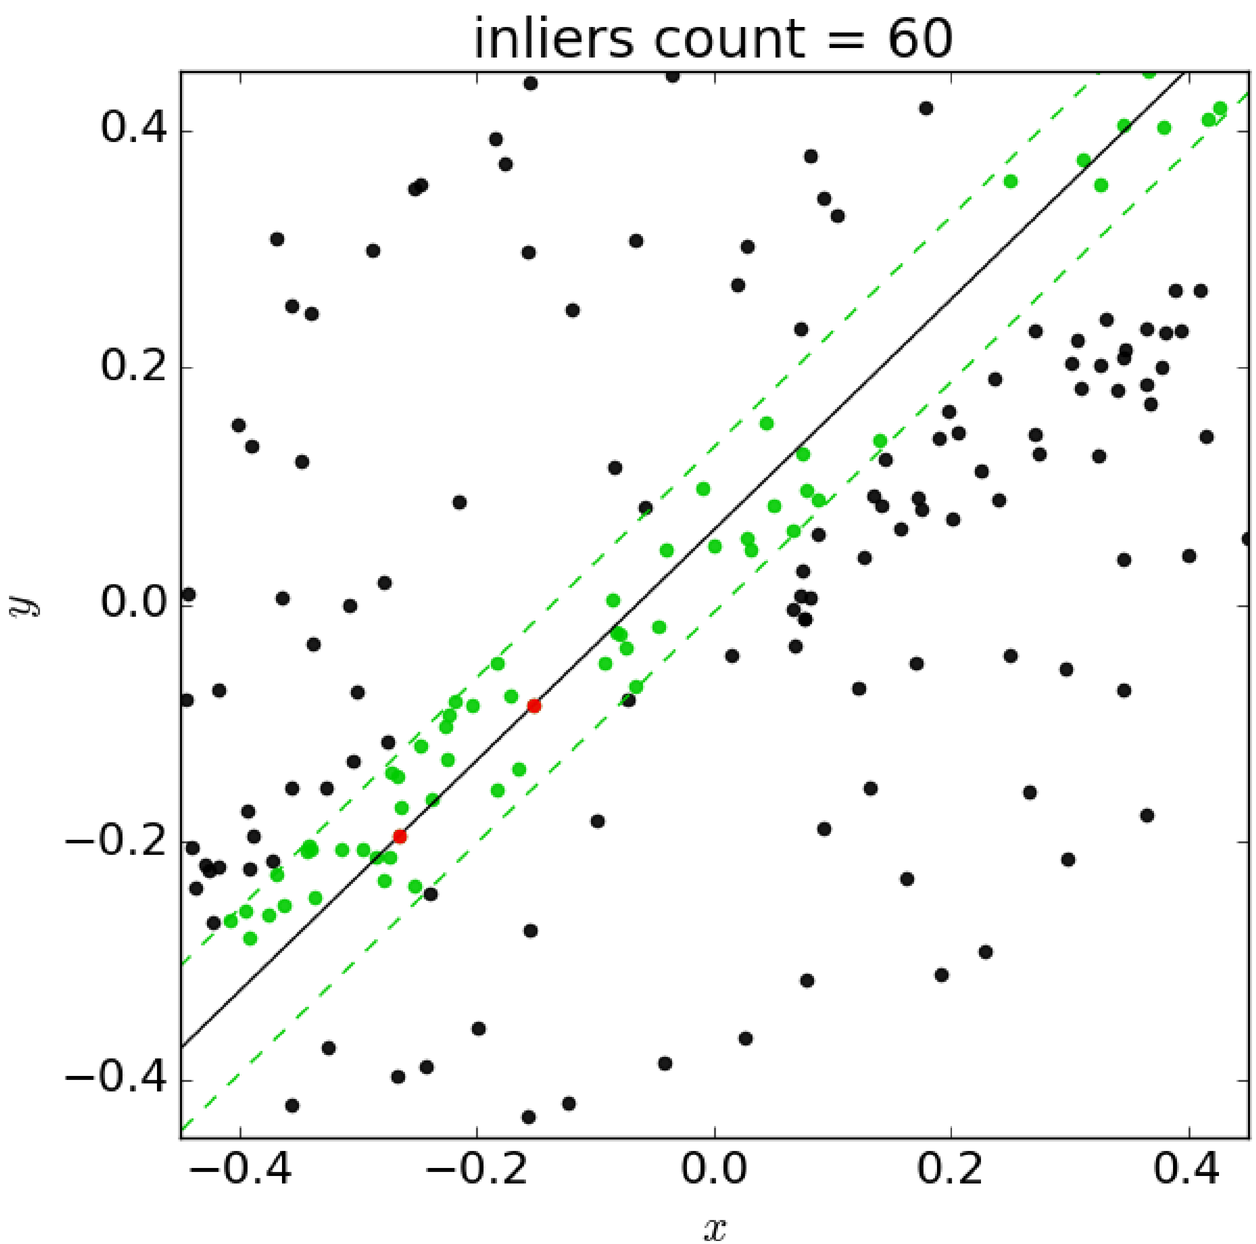

Init

LO-RANSAC: Example

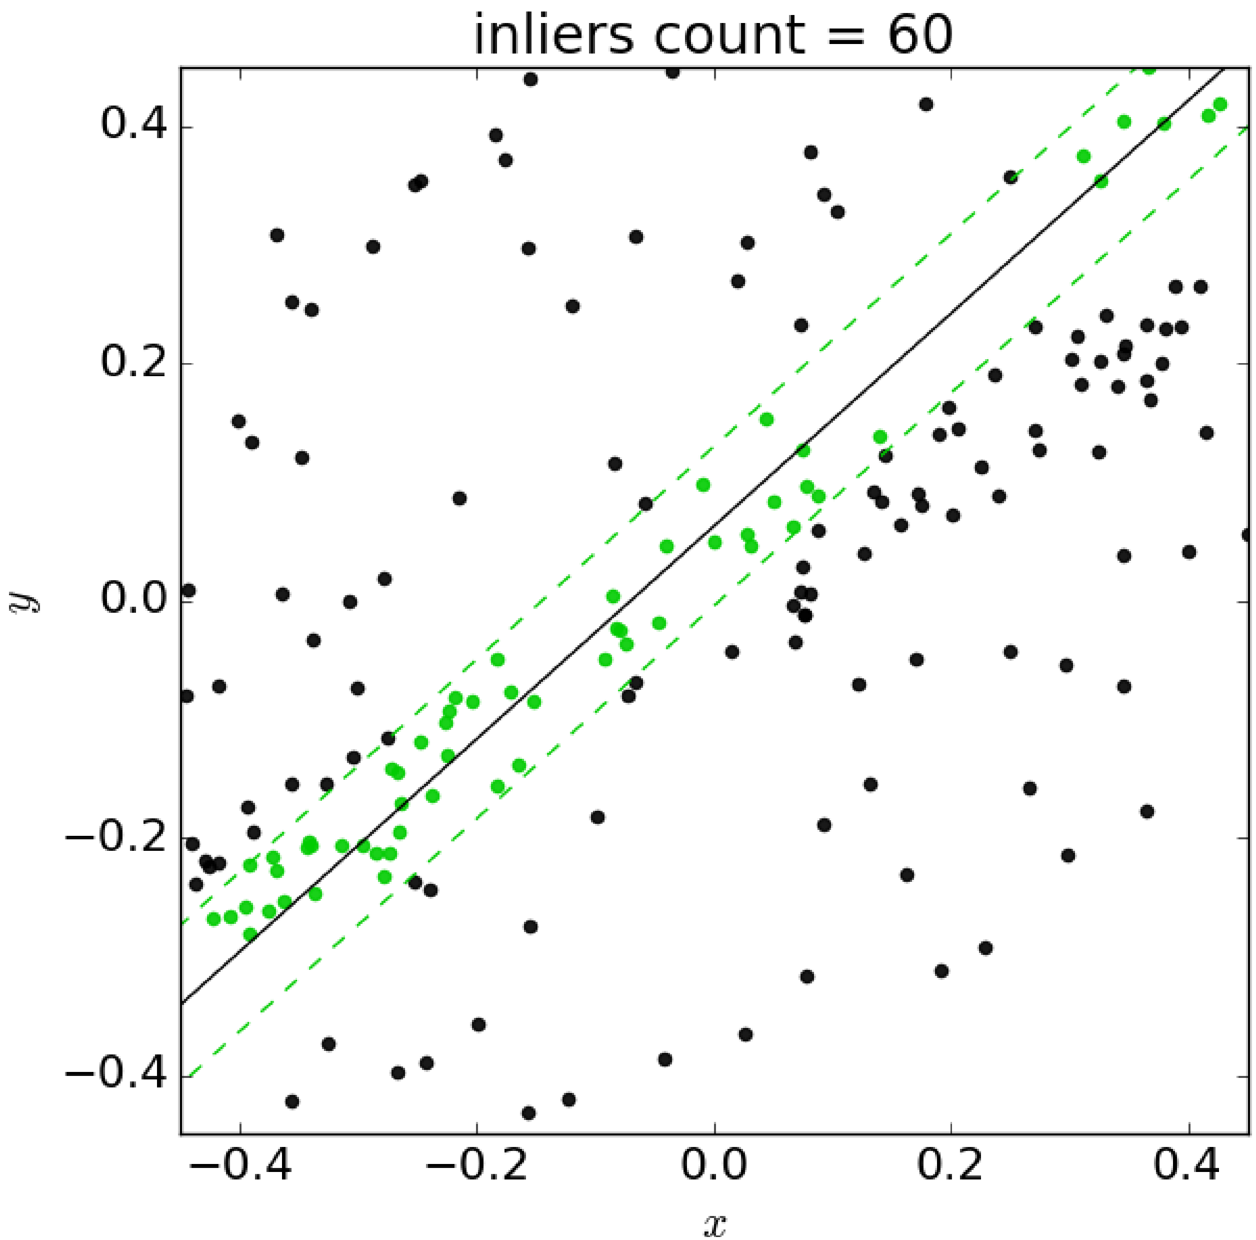

Init

Iteration 1

LO-RANSAC: Example

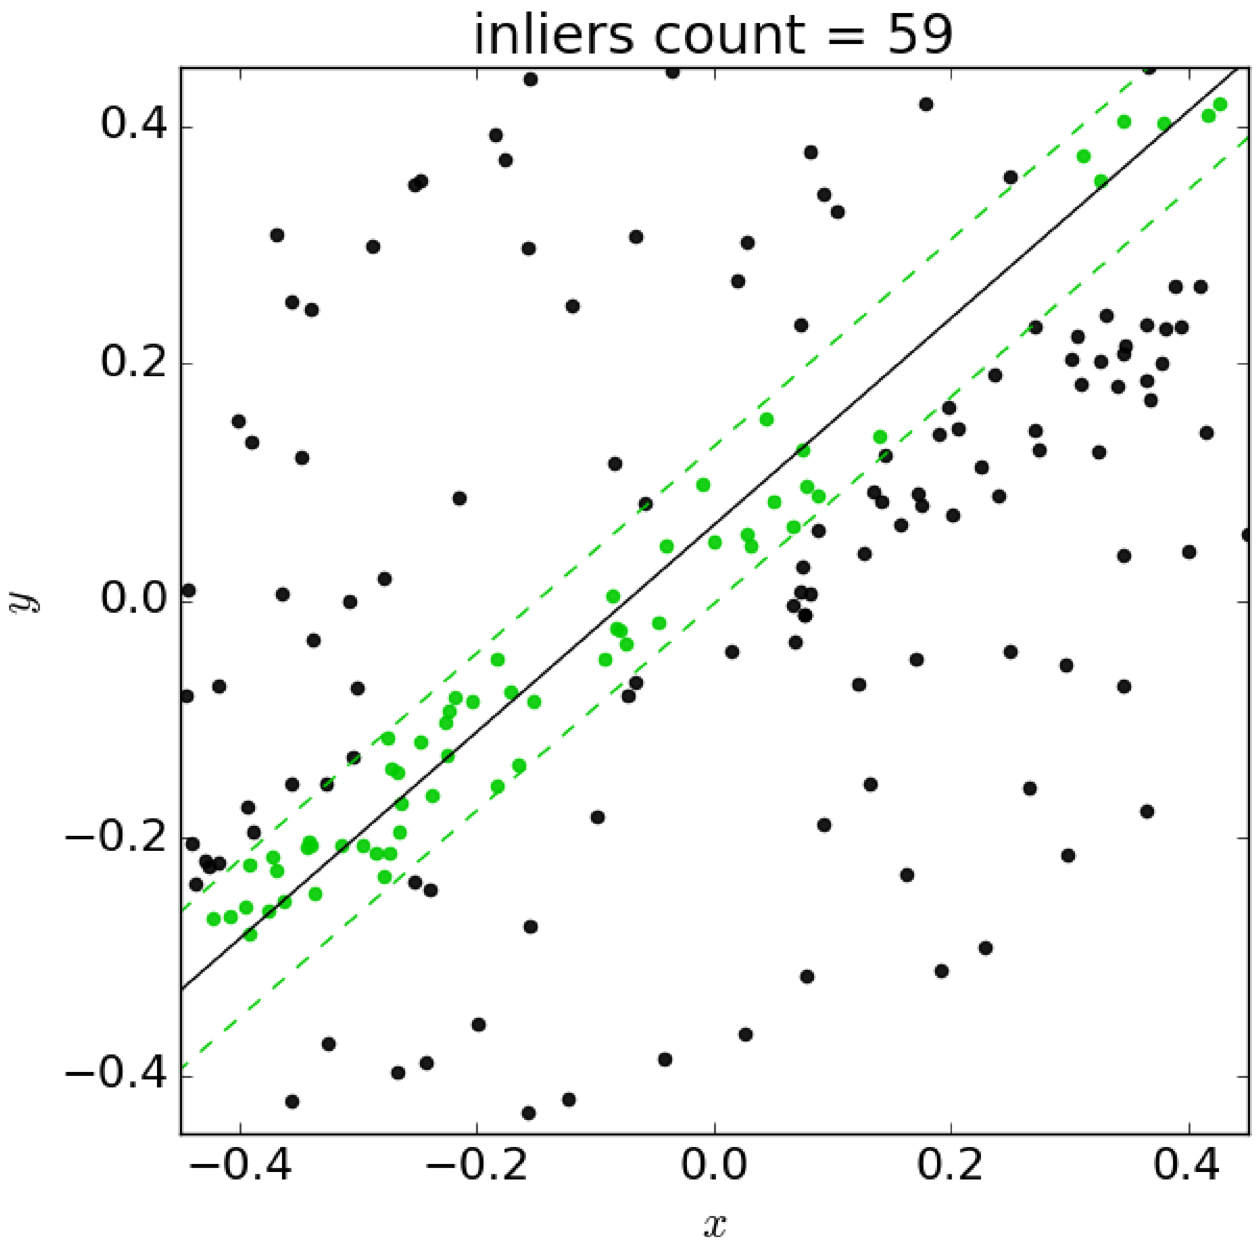

Init

Iteration 1

Iteration 2

LO-RANSAC: Example

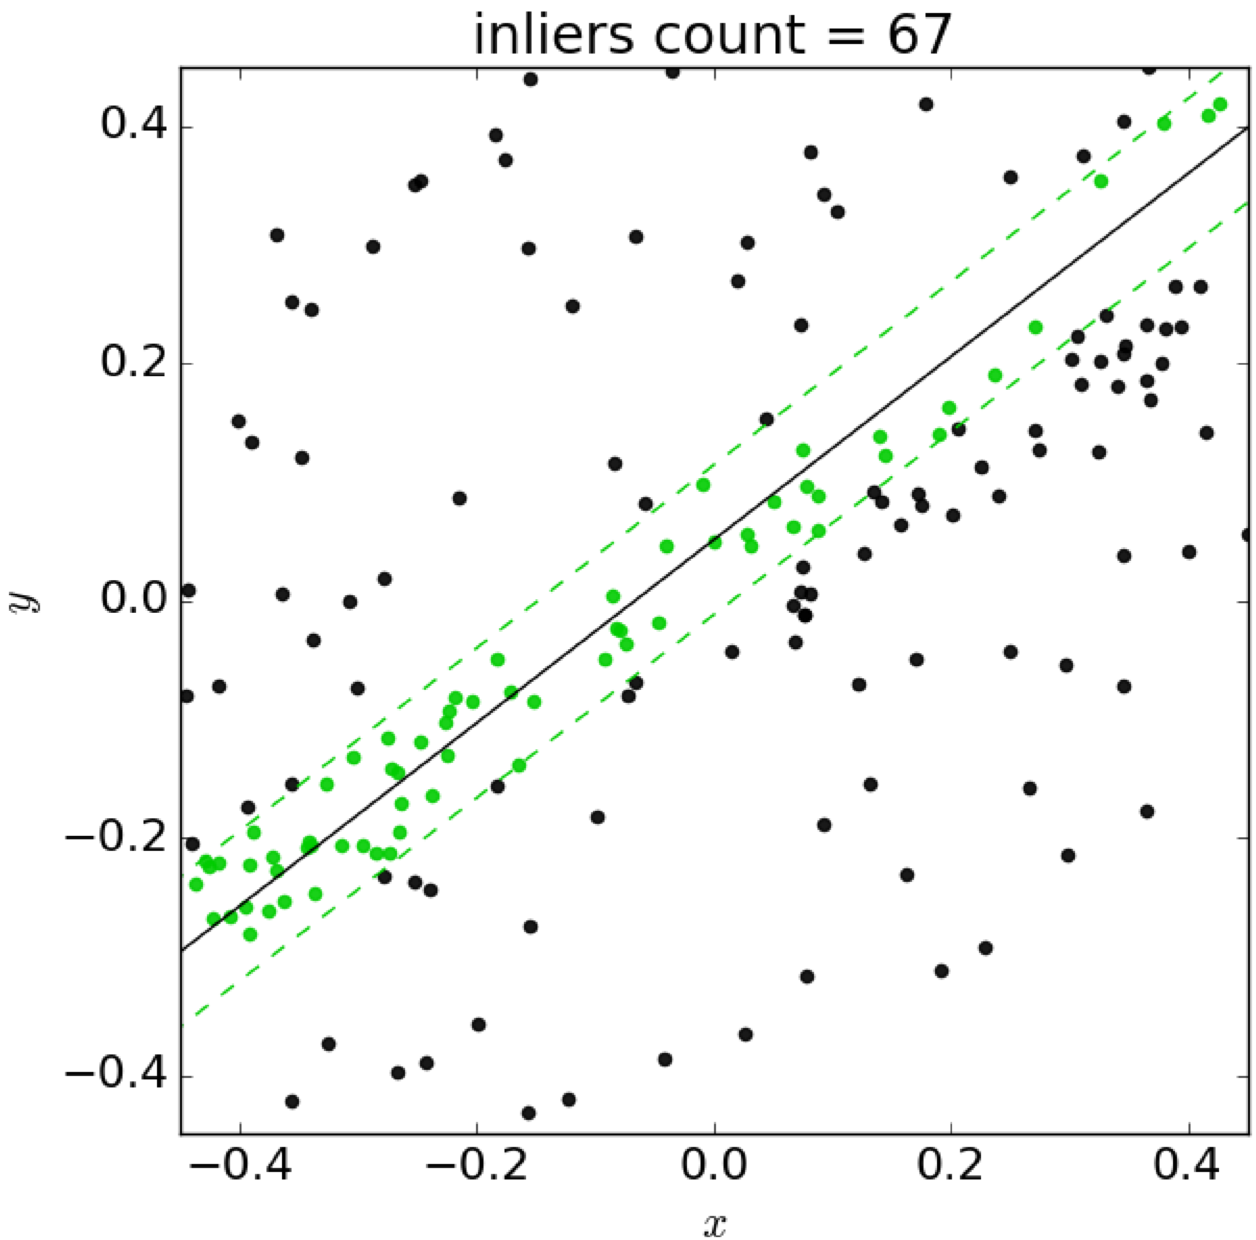

Init

Iteration 1

Iteration 2

...

Iteration 7

LO-RANSAC: Example

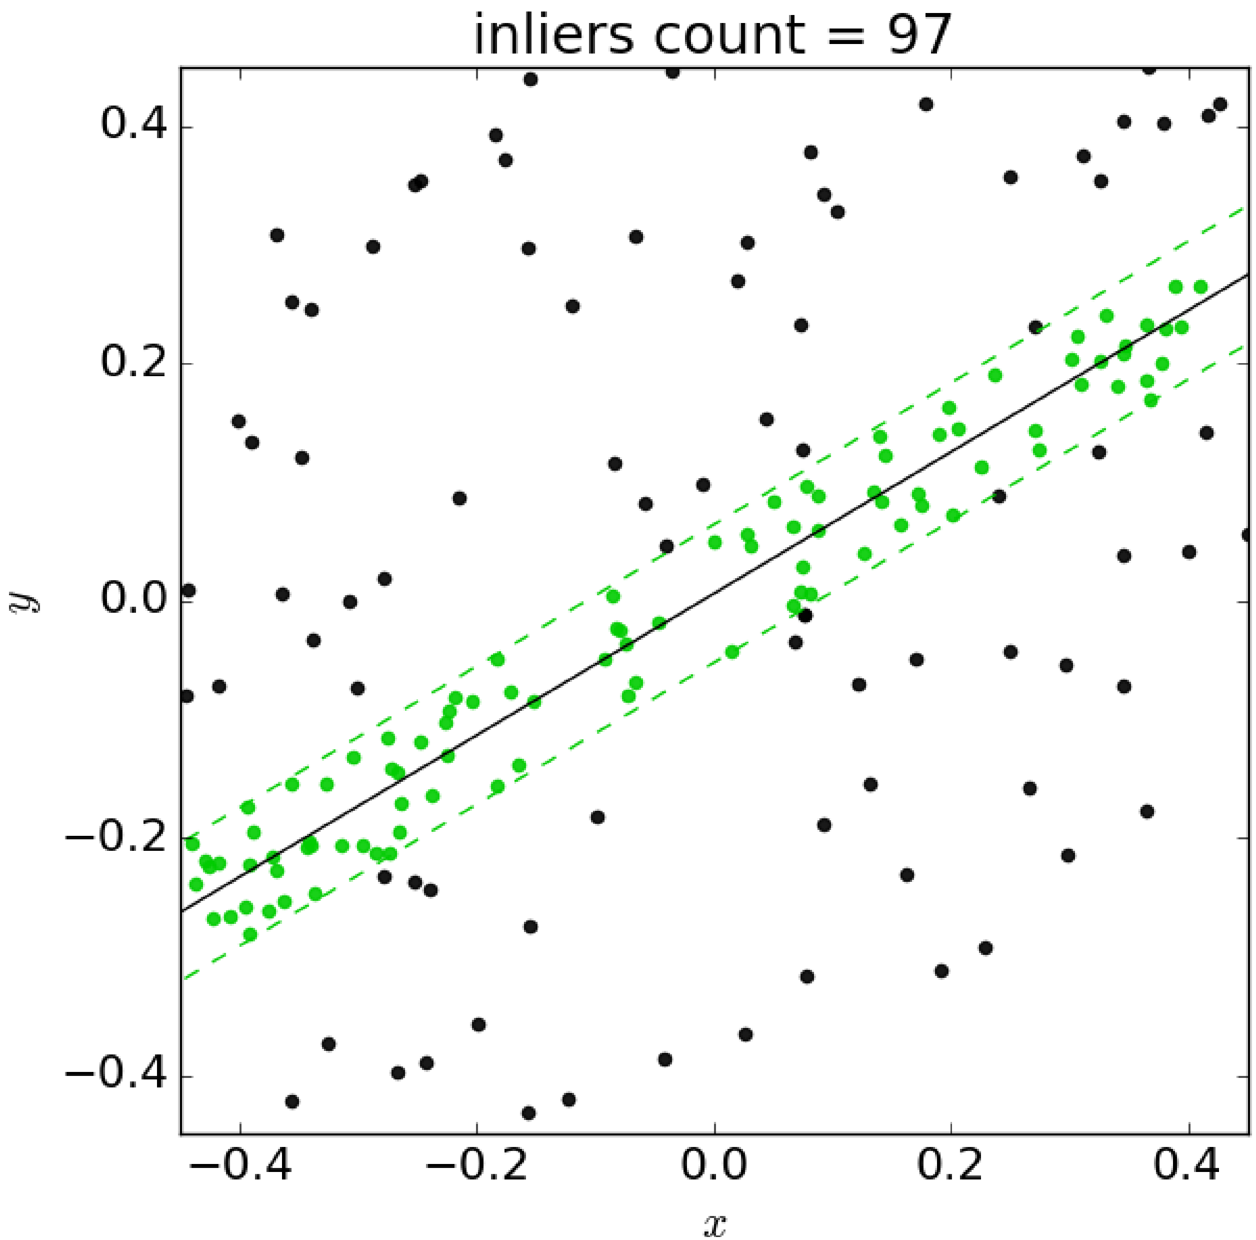

Init

Iteration 1

Iteration 2

...

Iteration 7

...

Itration 15

LO-RANSAC: Example

Comparison with model (100 inliers):

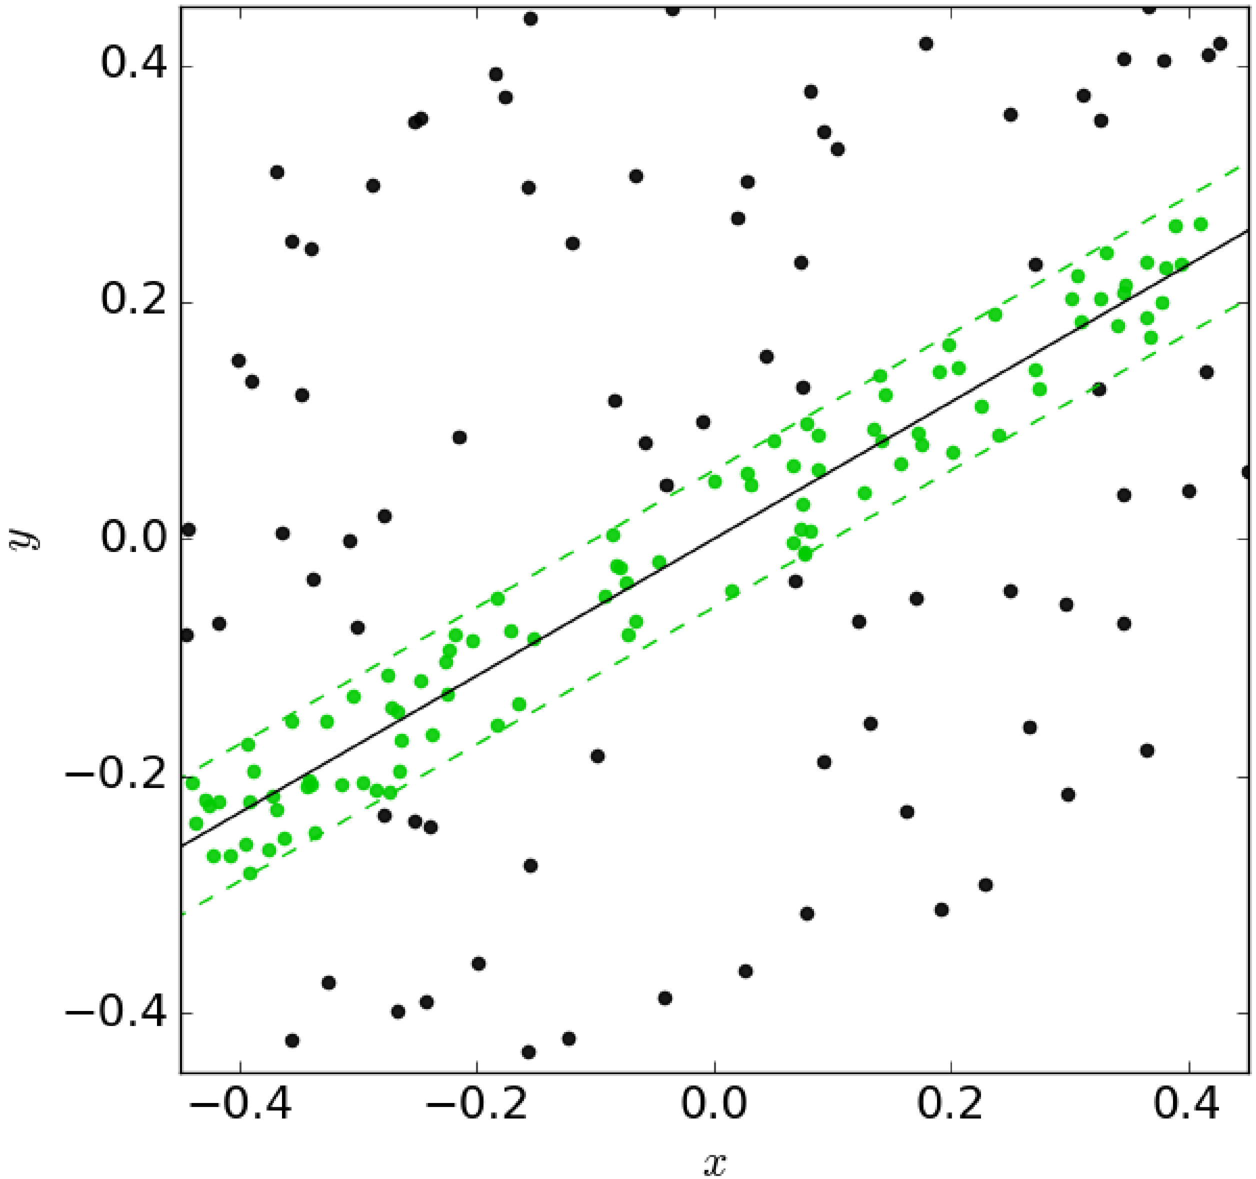

Locally Optimized RANSAC

Estimation of (approximate) models with lower complexity (less data points in the sample) followed by LO step estimating the desired model speeds the estimation up significantly.

The estimation of epipolar geometry is up to 10000 times faster when using 3 region-to-region correspondences rather than 7 point-to-point correspondences.

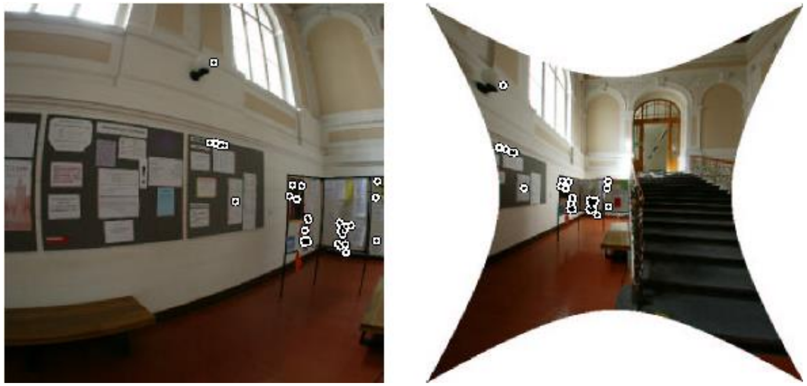

Fish-eye images by Braňo Mičušík

Simultaneous estimation of radial distortion and epipolar geometry with LO is superior to the state-of the art in both speed a precision of the model.

Chum, Matas, Obdržálek: Enhancing RANSAC by Generalized Model Optimization, *ACCV* 2004

LO-RANSAC: Problem Summary

It was observed experimentally that RANSAC takes several times longer than theoretically expected. This is due to the noise – not every all-inlier sample generates a good hypothesis.

By applying local optimization (LO) to the-best-so-far hypotheses:

(i) a near perfect agreement with theoretical performance

(ii) lower sensitivity to noise and poor conditioning.

The LO is shown to be executed so rarely, log(*iter*) times, that it has minimal impact on the execution time.

Chum, Matas, Kittler: Locally Optimized RANSAC, DAGM 2003

RANSAC – Time Complexity

Repeat k times (k is a function of sample size m, number of inliers Q, number of data N, and confidence η)

1. Hypothesis generation
   - Select a sample of m data points
   - Calculate parameters of the model(s)
2. Model verification
   - Find the support (consensus set) by verifying all N data points

t_M – time needed to draw a sample

$\overline{m}_s$ – average number of models per sample

Running time:

$$t = k(t_M + \overline{m}_s N)$$

Note 1: unit of time = time to evaluate 1 point ($\Rightarrow$ evaluating N points takes time N).

Note 2: number of models per sample for our toy, line fitting example, is equal to 1. Some tasks (e.g. epipolar geometry estimation) generate different number of solutions (models) per sample, depending on the sample data. 7-point algorithm, for example, generates up to 3 models.

Randomised RANSAC (R-RANSAC) [Matas, Chum 02]

Repeat until termination condition is met:

1. Hypothesis generation (as before)

2a. Model pre-verification $T_{d,d}$ test:

Evaluate $d \ll N$ data points, reject the model if not all d data points are consistent with the model

2b. Model verification

Verify the rest of the data points if pre-verification test was successful

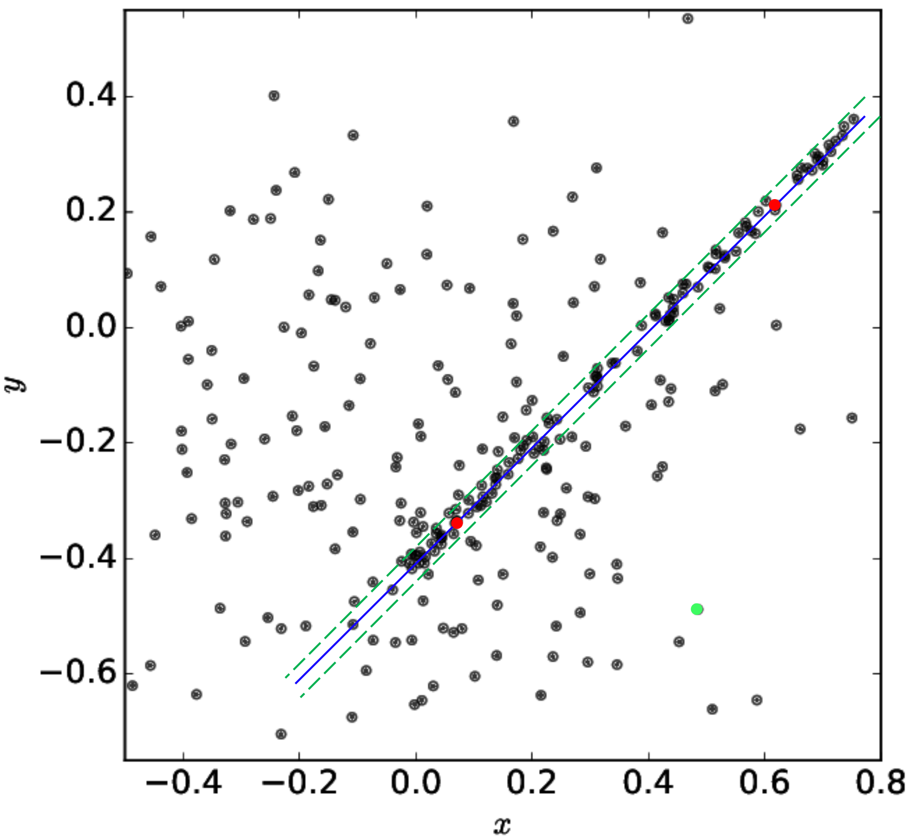

Example (d=1)

1. Generate a model (sample 2 points)

2a. Sample another point ●

Does it fall within threshold?

No. Go to 1.

Randomised RANSAC (R-RANSAC) [Matas, Chum 02]

Repeat until termination condition is met:

1. Hypothesis generation (as before)

2a. Model pre-verification $T_{d,d}$ test:

Evaluate $d \ll N$ data points, reject the model if not all d data points are consistent with the model

2b. Model verification

Verify the rest of the data points if pre-verification test was successful

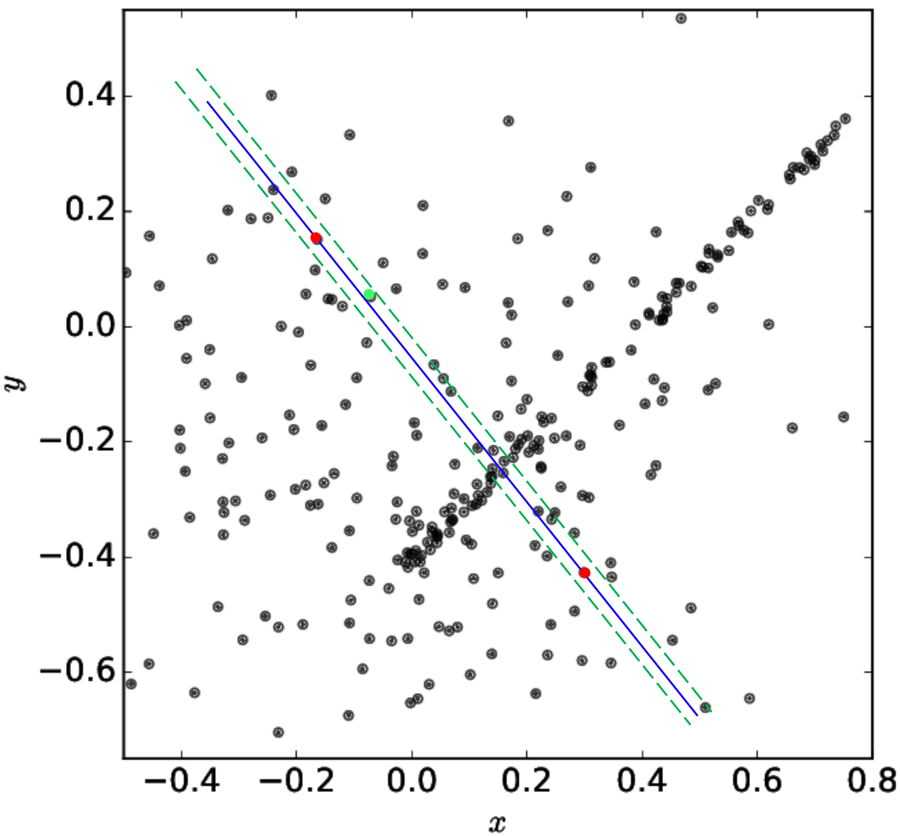

Example (d=1)

1. Generate a model (sample 2 points)

2a. Sample another point ●

Does it fall within threshold?

Yes.

2b. Verify all other points.

R-RANSAC Example, Running Time Analysis

Find a line in 2D points. N=10k, $\epsilon = 0.1$ (10% inliers.)
RANSAC:
Probability of selecting 2 'good' points is ϵ^2.
Average number of samples to find a good model is $1/\epsilon^2 = 100$.
For each model, N points are verified.
Total number of evaluations is $100N =$ **1M**

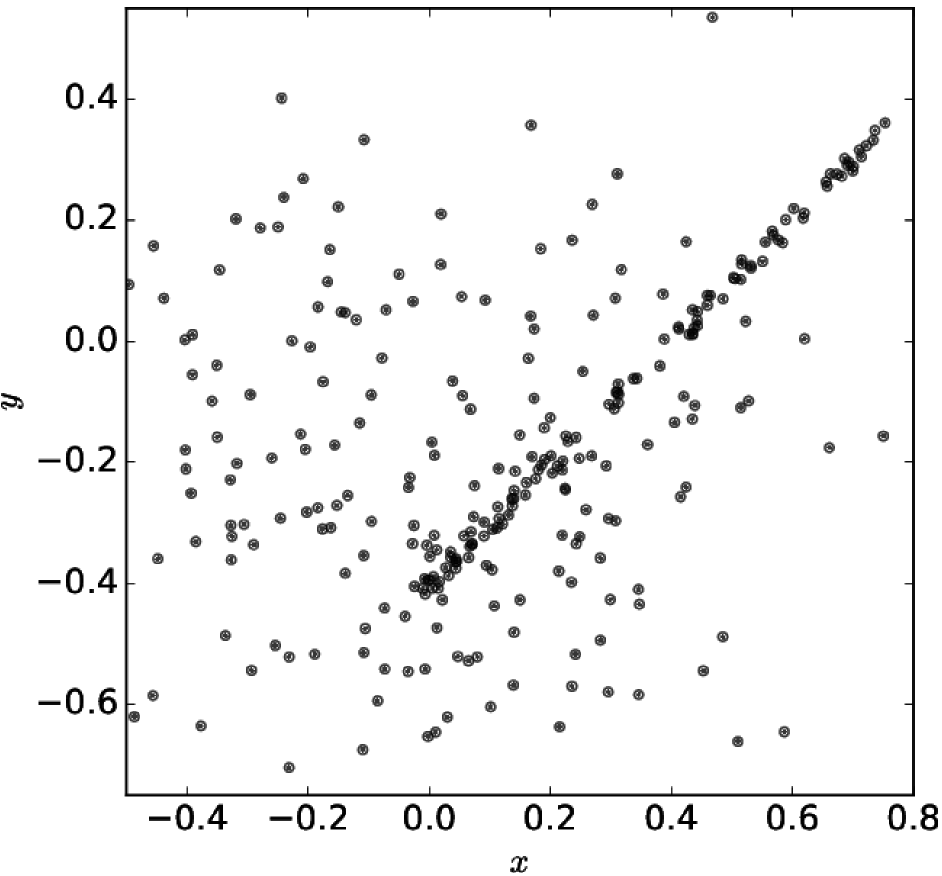

R-RANSAC Example, Running Time Analysis

Find a line in 2D points. N=10k, $\epsilon = 0.1$ (10% inliers.)

R-RANSAC (d=1):

Probability of selecting 2 'good' points is ϵ^2.

Probability of selecting inlier point for pre-verification is ϵ.

Average number of samples to find a good model is $1/\epsilon^3 = 1000$.

Probability of a random point passing pre-verification test for a 'bad' model is $\delta = 0.03$

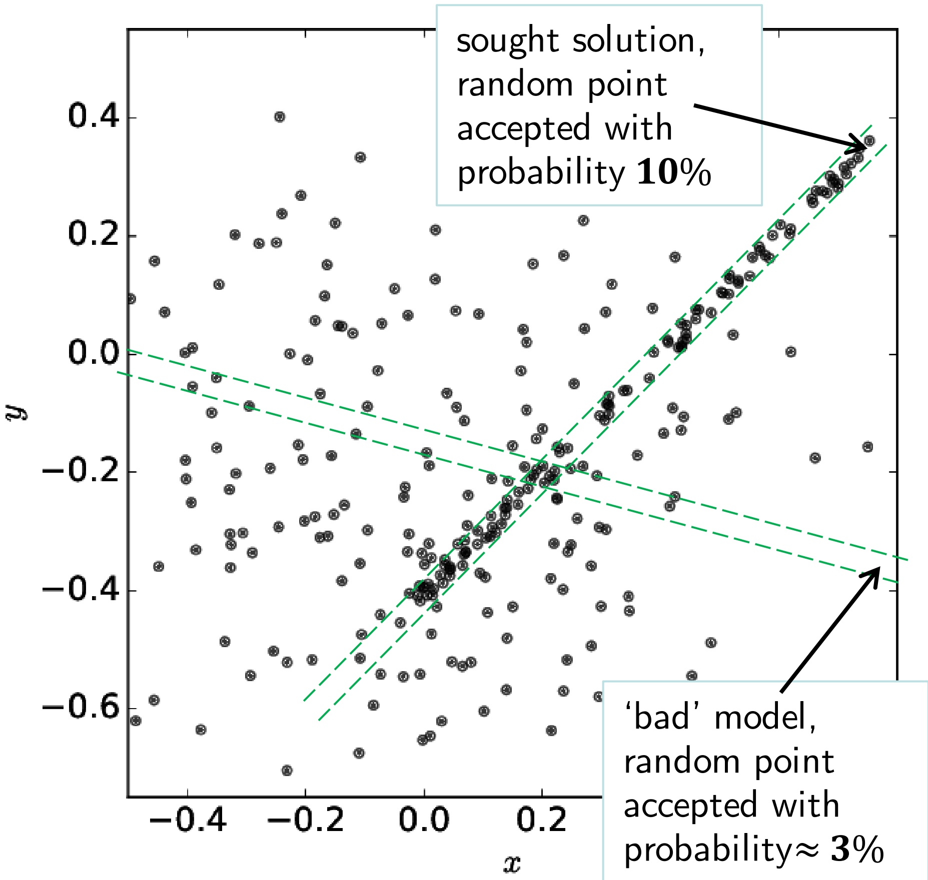

In 1000 samples:

$1000 \cdot \epsilon^2 = 10$ 'good' models

- $10 \cdot \epsilon = 1$ passes pre-verification
- $10 \cdot (1-\epsilon) = 9$ fails pre-verification

$1000 \cdot (1 - \epsilon^2) = 990$ 'bad' models

- $990 \cdot \delta = 30$ passes pre-verification
- $990 \cdot (1-\delta) = 960$ fails pre-verification

Total number of evaluations, on average:

$1N$	(good model, point accepted)
+ 9	(good model, point rejected)
+ $30N$	(bad model, point accepted)
+ 960	(bad model, point rejected)

≈ **311k**

R-RANSAC Example, Running Time Analysis

Find a line in 2D points. N=10k, $\epsilon = 0.1$ (10% inliers.)
R-RANSAC (d=2):
Probability of selecting 2 'good' points is ϵ^2.
Probability of selecting 2 inlier points for pre-verification is ϵ^2.
Average number of samples to find a good model is $1/\epsilon^4 = 10000$.

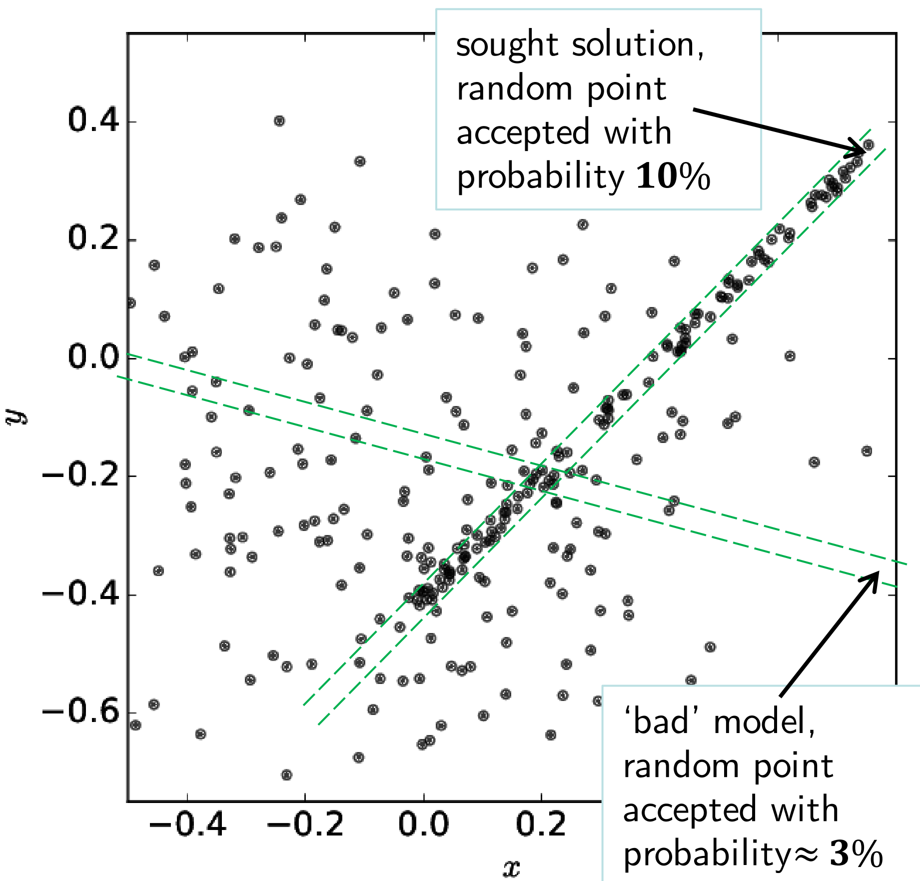

In 10000 samples:
$10000 \cdot \epsilon^2 = 100$ 'good' models
  $100 \cdot \epsilon^2 = 1$ passes pre-verification
  $100 \cdot (1-\epsilon) = 99$ fails pre-verification
$10000 \cdot (1-\epsilon^2) = 9900$ 'bad' models
  $9900 \cdot \delta^2 = 9$ passes pre-verification
  $990 \cdot (1- \delta^2) = 9891$ fails pre-verification

Total number of evaluations, on average:
$1N$ (good model, 2 points accepted)
$+\ 99 \cdot 2$ (good model, point(s) rejected)
$+\ 9N$ (bad model, 2 points accepted)
$+\ 9891 \cdot 2$ (bad model, point rejected)
$\approx$ **120k**

Note: For this case, d=2 is optimal (fastest)

Randomised RANSAC (R-RANSAC) [Matas, Chum 02]

Speeds up RANSAC; "Randomised" stands for *randomised verification*

Running time (RANSAC → R-RANSAC):

$$t = k(t_M + \overline{m}_s N) \quad \rightarrow \quad t = \frac{k}{1-\alpha}(t_M + \overline{m}_s V)$$

V - average number of data points verified
α – probability that a good model is rejected by $T_{d,d}$ test
k – *number of samples* (function of sample size, inlier ratio and confidence)

Optimal Randomised Strategy

Model Verification employing Sequential Decision Making

H_g: $P(x_i = 1|H_g) \geq \varepsilon$
H_b: $P(x_i = 1|H_b) = \delta$
$x_i = 1$ x_i is consistent with the model

where
H_g - hypothesis of a 'good' model (≈ from an uncontaminated sample)
H_b - hypothesis of a 'bad' model (≈ from a contaminated sample)
δ - probability of a data point being consistent with an arbitrary model

Optimal (the fastest) test that ensures with probability α that that H_g is not incorrectly rejected is the Sequential probability ratio test (SPRT) [Wald47]

SPRT [simplified from Wald 47]

Likelihood ratio $$\lambda_i = \prod_{j=1}^{i} \frac{P(x_j|H_b)}{P(x_j|H_g)}$$

Set (compute) threshold A. Set j=1

1. Select a point and check whether it is consistent with model
2. Update likelihood ratio
3. If $\lambda_j > A$ decide the model is 'bad', else increment j
4. If $j > N$ (total number of points) decide model is 'good', else go to 1.

Properties of SPRT:

1. probability of rejecting a "good" model $\alpha < 1/A$
2. average number of verifications $V = C \log(A)$

$$C \approx \left(P(0|H_b) \log \frac{P(0|H_b)}{P(0|H_g)} + P(1|H_b) \log \frac{P(1|H_b)}{P(1|H_g)} \right)^{-1}$$

$$C \approx \left((1-\delta) \log \frac{1-\delta}{1-\varepsilon} + \delta \log \frac{\delta}{\varepsilon} \right)^{-1}$$

SPRT properties

Probability of rejecting a "good" model $\alpha = 1/A$

$$\lambda_i = \prod_{j=1}^{i} \frac{P(x_j|H_b)}{P(x_j|H_g)} = \frac{P(x|H_b)}{P(x|H_g)}, x = (x_1, \ldots, x_i)$$

If $\lambda_i > A$ then $P(x|H_g) < P(x|H_b)/A$, therefore

$$\alpha = \int_{\lambda_i > A} P(x|H_g) dx < \int_{\lambda_i > A} P(x|H_b)/A dx =$$

$$= \frac{1}{A} \int_{\lambda_i > A} P(x|H_b) dx \leq \frac{1}{A} \int P(x|H_b) dx = \frac{1}{A}$$

WaldSAC

Running time

$$t(A) = \frac{k}{(1 - 1/A)}(t_M + \overline{m}_S C \log A)$$

In sequential statistical decision problem decision errors are traded off for time. These are two incomparable quantities, hence the constrained optimization.

In WaldSAC, decision errors cost time (more samples) and there is a single minimised quantity, time $t(A)$, a function of a single parameter A.

Optimal test (optimal A) given ε and δ

Optimal A^* $$A^* = \arg\min_A t(A)$$

Optimal A^* found by solving $$\frac{\partial t}{\partial A} = 0$$

$$A^* = \frac{t_M}{\overline{m}_s C} + 1 + \log A^*$$

$$A^* = \lim_{n \to \infty} A_n$$

Computed in several iterations:

$$A_0 = \frac{t_M}{\overline{m}_s C} + 1, \quad A_{n+1} = \frac{t_M}{\overline{m}_s C} + 1 + \log A_n$$

SPRT

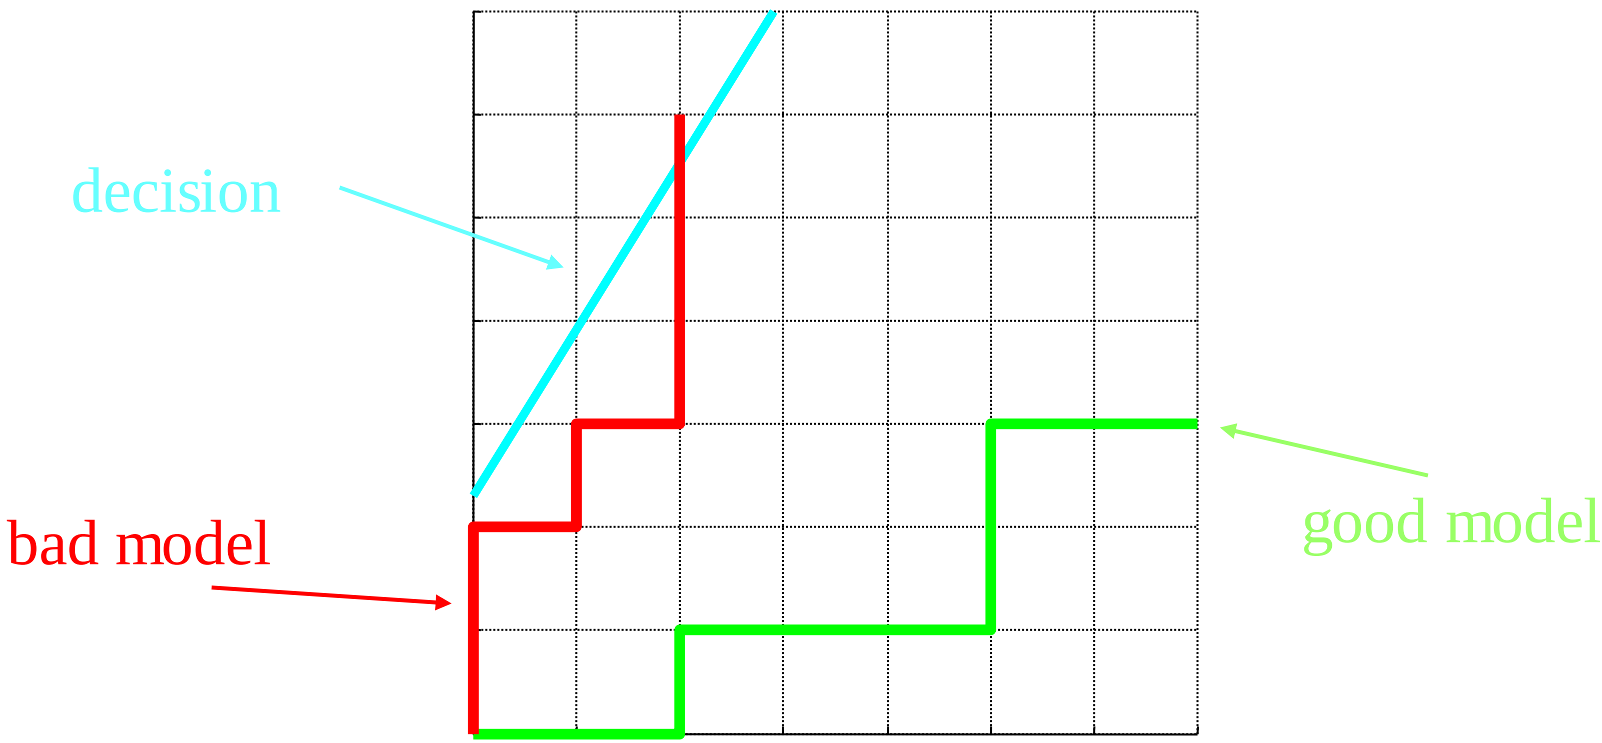

Note: the Wald's test is equivalent to series of $T(d, c)$, where $c = \lceil (\log A - d \log \lambda_1)/\log \lambda_0 \rceil$

Exp. 1: Wide-baseline matching

	samples	models	V	time	spd-up
R	2914	7347	110.0	1099504	1.0
R-R	7825	19737	3.0	841983	1.3
Wald	3426	8648	8.2	413227	2.7

Exp. 2 Narrow-baseline stereo

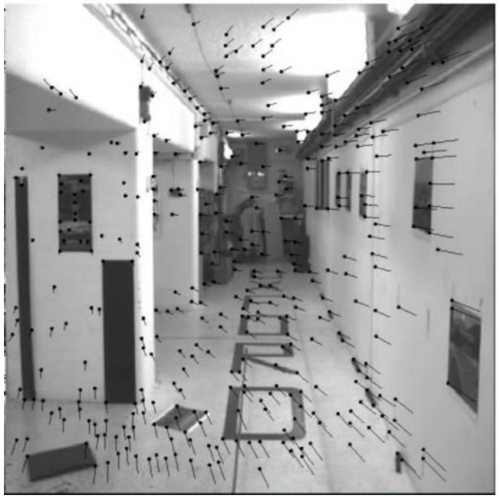

	samples	models	V	time	spd-up
R	155	367	600.0	235904	1.0
R-R	247	587	86.6	75539	3.1
Wald	162	384	23.1	25032	9.4

Randomised Verification in RANSAC: Conclusions

- The same confidence η in the solution reached faster (data dependent, $\approx$ 10x)
- No change in the character of the algorithm, it was randomised anyway.
- Optimal strategy derived using Wald's theory for known ε and δ.
- Results with ε **and** δ estimated during the course of RANSAC are not significantly different. Performance of SPRT is insensitive to errors in the estimate.
- δ can be learnt, an initial estimate can be obtained by geometric consideration
- Lower bound on e is given by the best-so-far support

PROSAC – PROgressive SAmple Consensus

- Not all correspondences are created equally
- Some are better than others
- Sample from the best candidates first

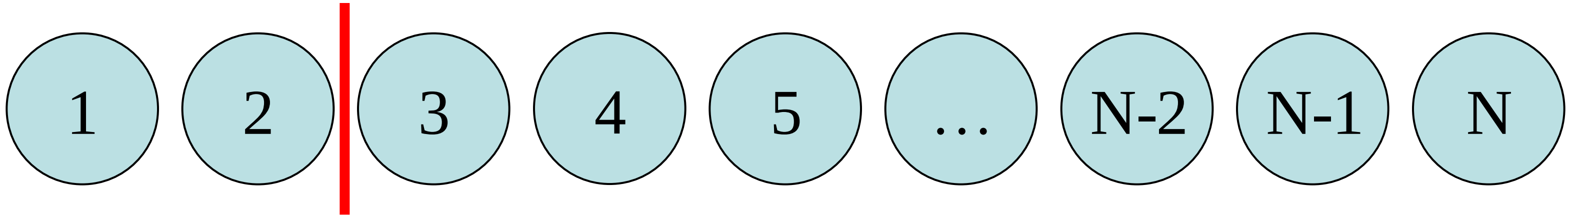

Sample from here

PROSAC Samples

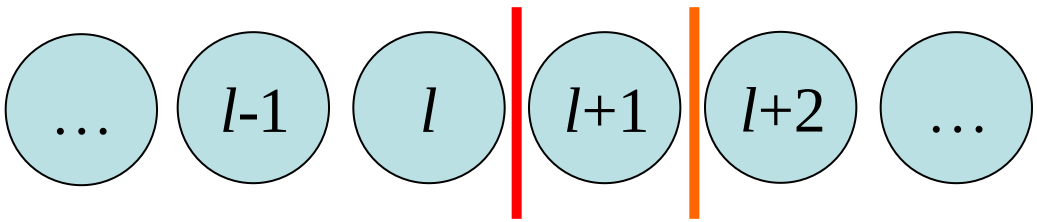

Draw T_l samples from $(1 \ldots l)$
Draw T_{l+1} samples from $(1 \ldots l+1)$

Samples from $(1 \ldots l)$ that are not from $(1 \ldots l+1)$ contain $l+1$

Draw T_{l+1} - T_l samples of size m-1 and add $l+1$

Degenerate Configurations

The presence of degenerate configuration causes RANSAC to fail in estimating a correct model, instead a model consistent with the degenerate configuration and some outliers is found.

The DEGENSAC algorithm handles scenes with:

- all points in a single plane
- majority of the points in a single plane and the rest off the plane
- no dominant plane present

No a-priori knowledge of the type of the scene is required

Chum, Werner, Matas: Epipolar Geometry Estimation unaffected by dominant plane, *CVPR* 2005

GC-RANSAC [Barath and Matas, CVPR 2018]

Input: $\mathcal{X} = \{\mathbf{x}_j\}_{j=1}^{N}$

$e(S) = \theta$ estimates model parameters θ, given sample $S \subseteq \mathcal{X}$

$$f(\mathbf{x}, \theta) = \begin{cases} 0, \text{if distance to model } \leq \text{threshold} \\ 1, \text{otherwise} \end{cases}$$

Output: θ^* parameter of the model minimizing the cost function

1. $iter = 0,\ J^* = \infty$
2. **repeat**
3. Select random $S \subseteq \mathcal{X}$ (sample size $m = |S|$)
4. Estimate parameter $\theta = e(S)$
5. Evaluate $J(\theta) = \sum_{x \in \mathcal{X}} f(\mathbf{x}, \theta)$
6. If $J(\theta) < J^*$ then
7. $\theta^*, \mathcal{L}^* \leftarrow \arg\min_{\theta, \mathcal{L}} \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta) + \lambda \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{A}} [\![\mathcal{L}(\mathbf{x}) \neq \mathcal{L}(\mathbf{y})]\!]$
8. $J^* \leftarrow J(\theta^*)$
9. $iter \leftarrow iter + 1$
10. **until** $P(\text{better solution exists}) = f(|\mathcal{X}|, J^*, iter) < \mu$

Run graph-cut, if a so-far-the-best solution is found.

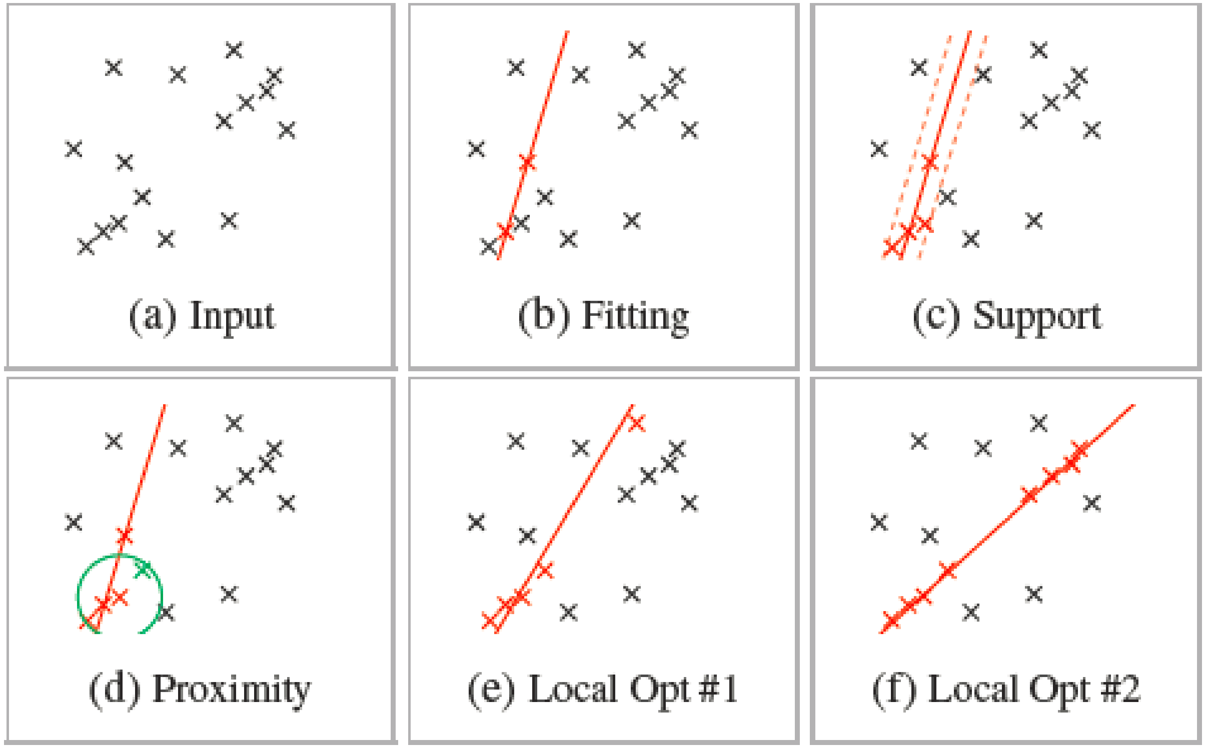

Figure 1: The proposed graph-cut based local optimization converging from a "not-all-inlier" sample, i.e. it is contaminated by an outlier, to the desired model. (a) The input data points, (b) RANSAC-like sampling and model fitting, (c) computation of model support, e.g. counting the inliers, (d) considering spatial proximity by graph-cut, (e-f) iterated local optimization using least-squares fitting and graph-cut.

GC-RANSAC

			Min. 60 FPS with 99% confidence					Confidence 95%				
			PLAIN	LO	LO+	LO'	GC	PLAIN	LO	LO+	LO'	GC
kusvod2	F, #24	LO	–	2	2	2	1 (3)	–	1	1	1	2 (3)
		Error (px)	5.01	4.95	4.97	5.02	**4.65**	5.18	5.08	5.03	5.22	**4.69**
		Time (ms)	6.17	6.09	6.26	5.85	**4.58**	4.94	5.19	5.14	4.86	**3.64**
		Samples	117.00	96.00	99.00	111.00	**70.00**	93.00	76.00	78.00	87.00	**53.00**
Adelaide	F, #19	LO	–	2	2	2	1 (3)	–	2	2	3	2 (4)
		Error (px)	0.55	0.53	0.52	0.55	**0.50**	0.44	0.45	**0.43**	0.44	**0.43**
		Time (ms)	14.20	14.83	14.85	**14.13**	18.94	262.73	**194.18**	210.85	237.09	227.12
		Samples	124.00	**113.00**	**113.00**	122.00	116.00	1363.00	1126.00	1205.00	1305.00	**1115.00**
Multi-H	F, #4	LO	–	1	1	1	1 (3)	–	2	1	2	1 (3)
		Error (px)	0.35	0.34	0.34	0.34	**0.32**	0.33	0.33	0.33	0.34	**0.32**
		Time (ms)	**10.34**	11.53	11.11	**10.34**	14.64	12.81	15.11	14.11	**12.37**	36.01
		Samples	83.00	76.00	76.00	82.00	**74.00**	107.00	89.00	90.00	100.00	**78.00**
EVD	H, #15	LO	–	2	2	2	2 (2)	–	4	4	4	3 (6)
		Error (px)	1.53	1.63	**1.51**	1.58	1.53	0.96	0.95	0.95	0.96	**0.92**
		Time (ms)	16.84	18.34	18.04	**16.82**	19.19	247.25	248.01	251.31	**246.95**	249.89
		Samples	320.00	**298.00**	301.00	318.00	301.00	4303.00	**4203.00**	4248.00	4291.00	4204.00
homogr	H, #16	LO	–	2	2	2	1 (3)	–	2	2	2	1 (4)
		Error (px)	0.53	0.53	0.53	0.53	**0.51**	0.50	0.50	0.49	0.50	**0.47**
		Time (ms)	7.13	10.37	9.83	**7.10**	7.56	17.10	10.09	9.89	8.52	**7.94**
		Samples	193.00	175.00	175.00	189.00	**159.00**	450.00	212.00	214.00	226.00	**165.00**
avg.	#78	LO	–	2	2	2	1 (3)	–	2	2	2	2 (3)
		Error (px)	2.10	2.09	2.07	2.11	**1.96**	1.98	1.95	1.93	2.00	**1.81**
		Time (ms)	10.59	11.73	11.60	**10.46**	12.01	115.90	**98.36**	102.88	107.89	107.06
		Samples	171.00	154.00	156.00	168.00	**144.00**	1286.00	1154.00	1183.00	1224.00	**1134.00**

GC RANSAC - Speed

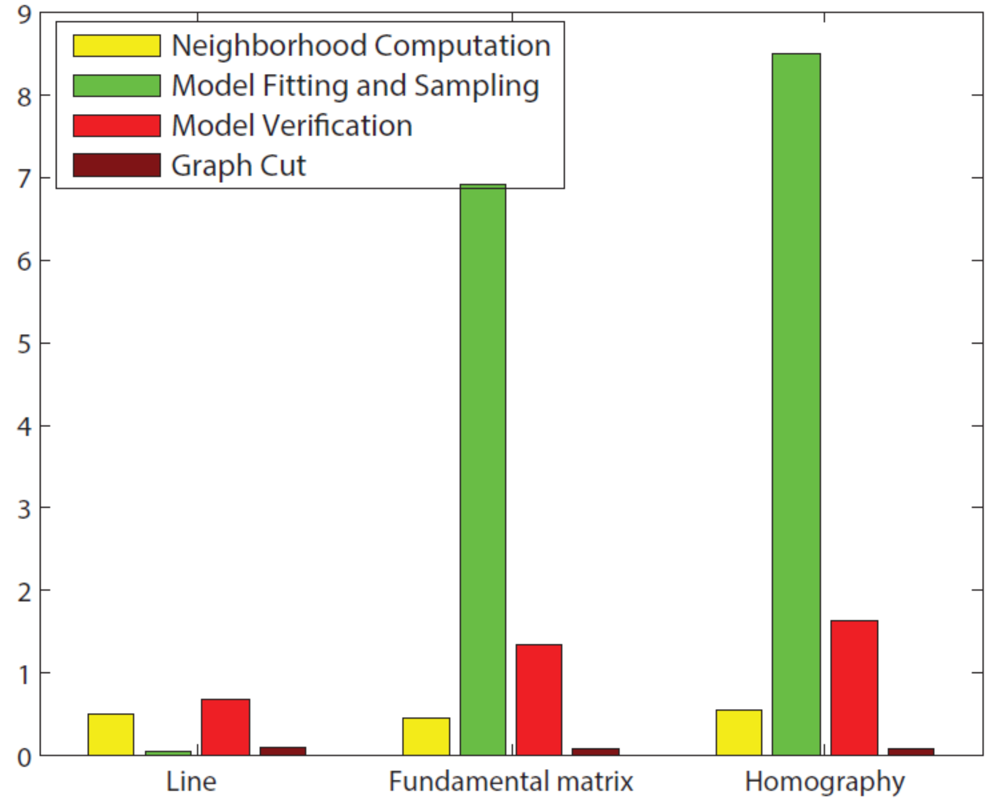

Figure 7: The breakdown of the processing times in milliseconds. Computed as the mean of all tests. *Best viewed in color.*