Robust Fitting in Computer Vision

Jiří Matas

Department of Cybernetics,

Faculty of Electrical Engineering,

Czech Technical University in Prague

http://cmp.felk.cvut.cz/~matas

Joint work with:

Ondra Chum, Daniel Barath, Dmytro Mishkin, Maksym Ivashechkin, ...

Taxonomy of Fitting Problems

- Standard Single Class Single Instance Fitting Problem (SCSI)

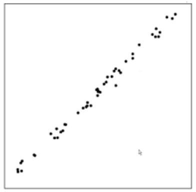 $\Rightarrow$ 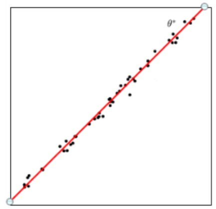

- Robust Single Class Single Instance Fitting Problem (R-SCSI)

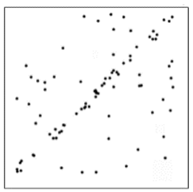 $\Rightarrow$ 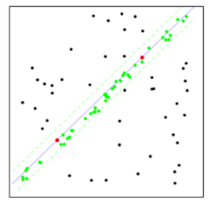

- Single Class Multiple Instance Fitting Problem (SCMI)

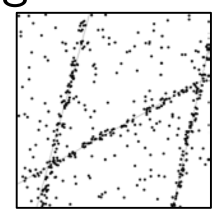 $\Rightarrow$ 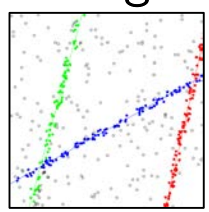

- Multiple Class Multiple Instance Fitting Problem (MCMI)

 $\Rightarrow$ 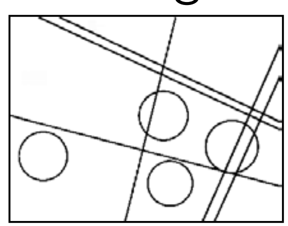

Single/Multi-Class Single/Multi-Instance Fitting Applications

- detection of geometric primitives


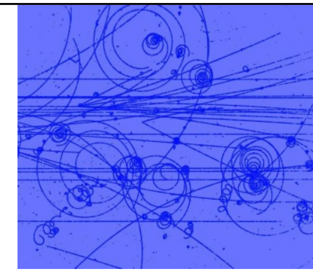

- epipolar geometry estimation
- detection of planar surfaces



- multiple motion segmentation


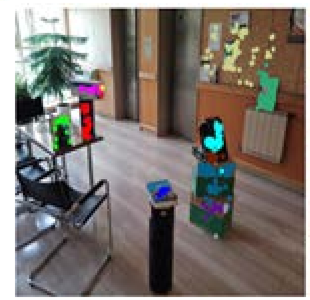
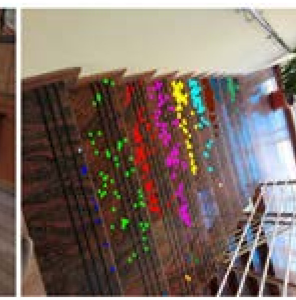

- Interpretation of lidar scans
- Registration of 3D point clouds

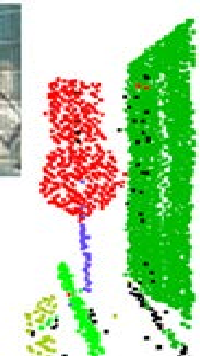
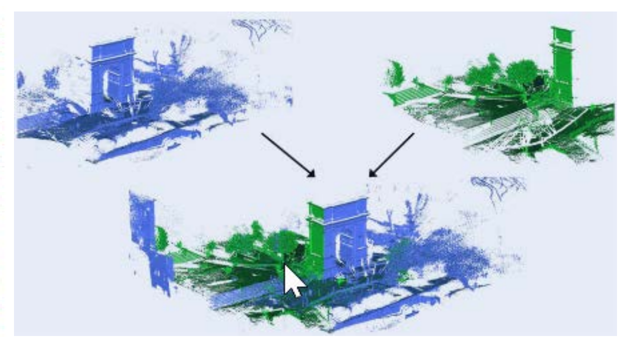

What is RANSAC?

- 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

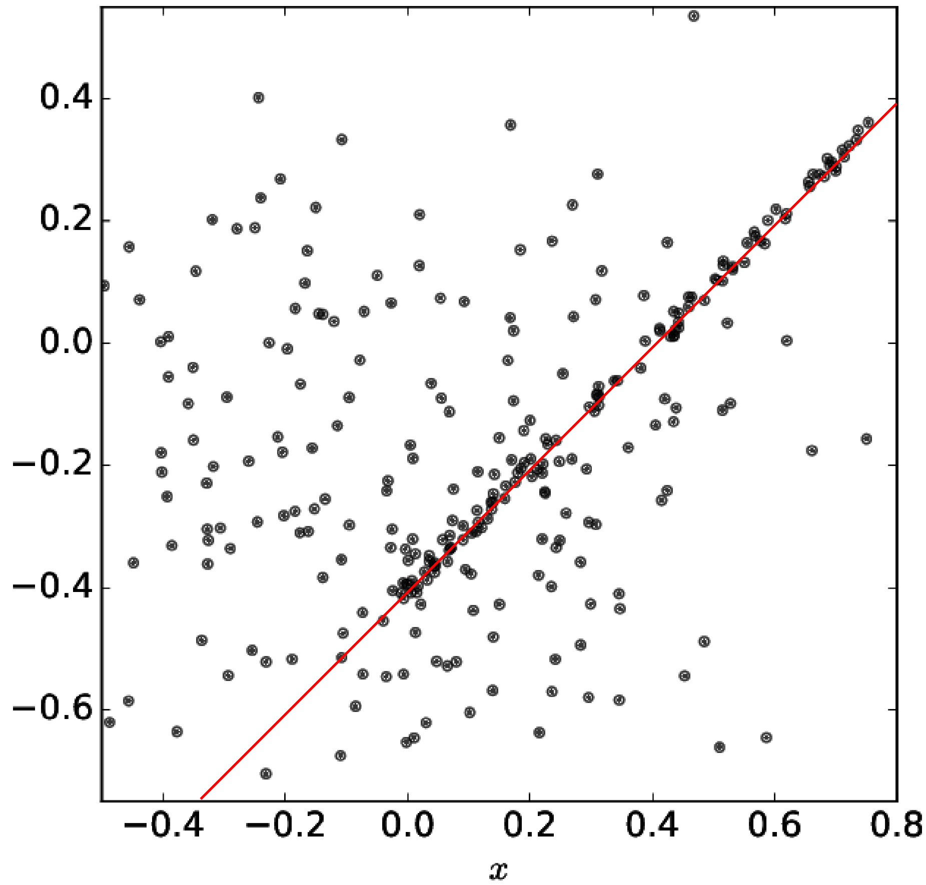

Line Fitting, Inliers Only: Easy!

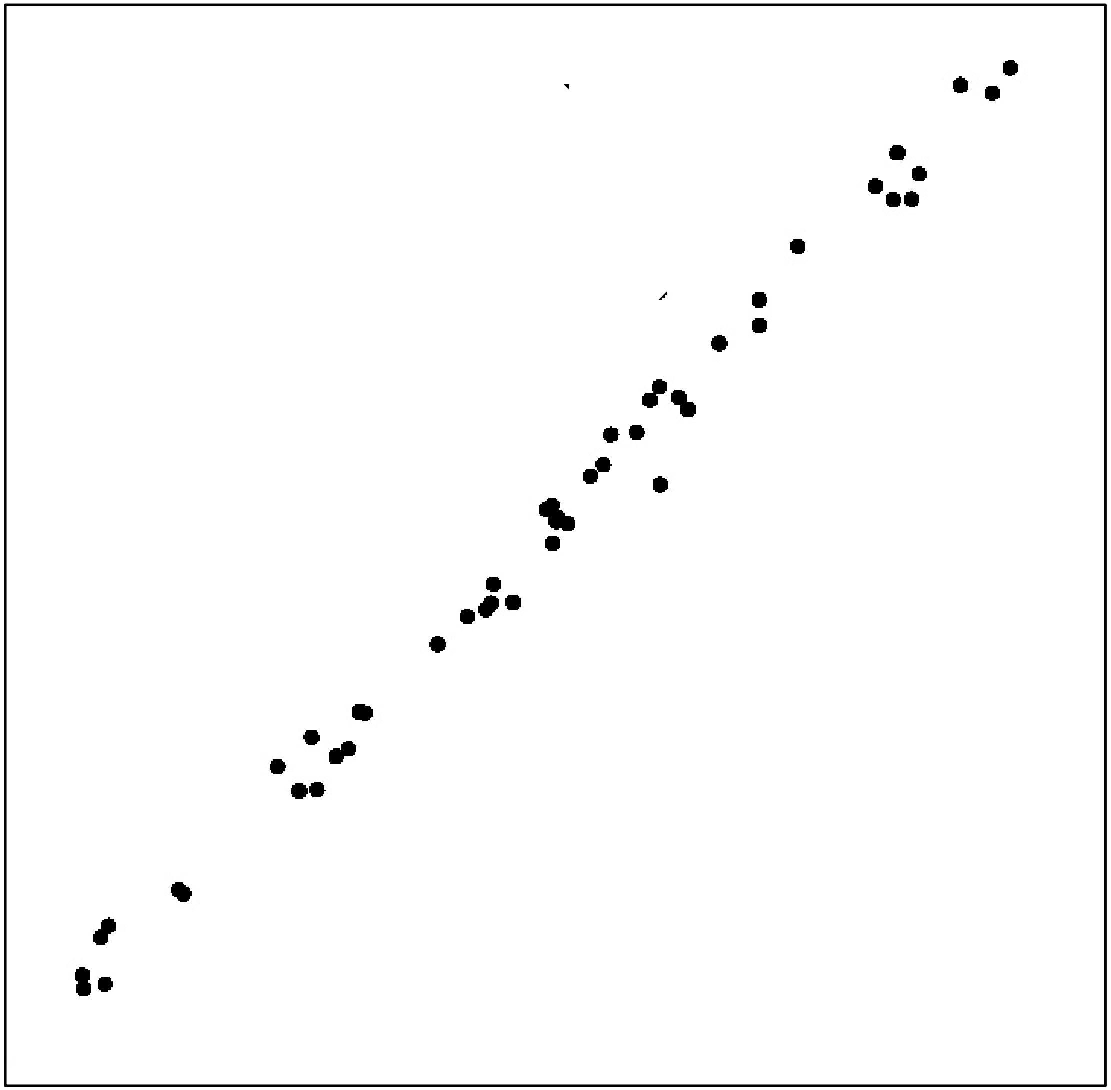

Data points

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

Find the line which "best fits" these points.

Line Fitting, Inliers Only: Easy!

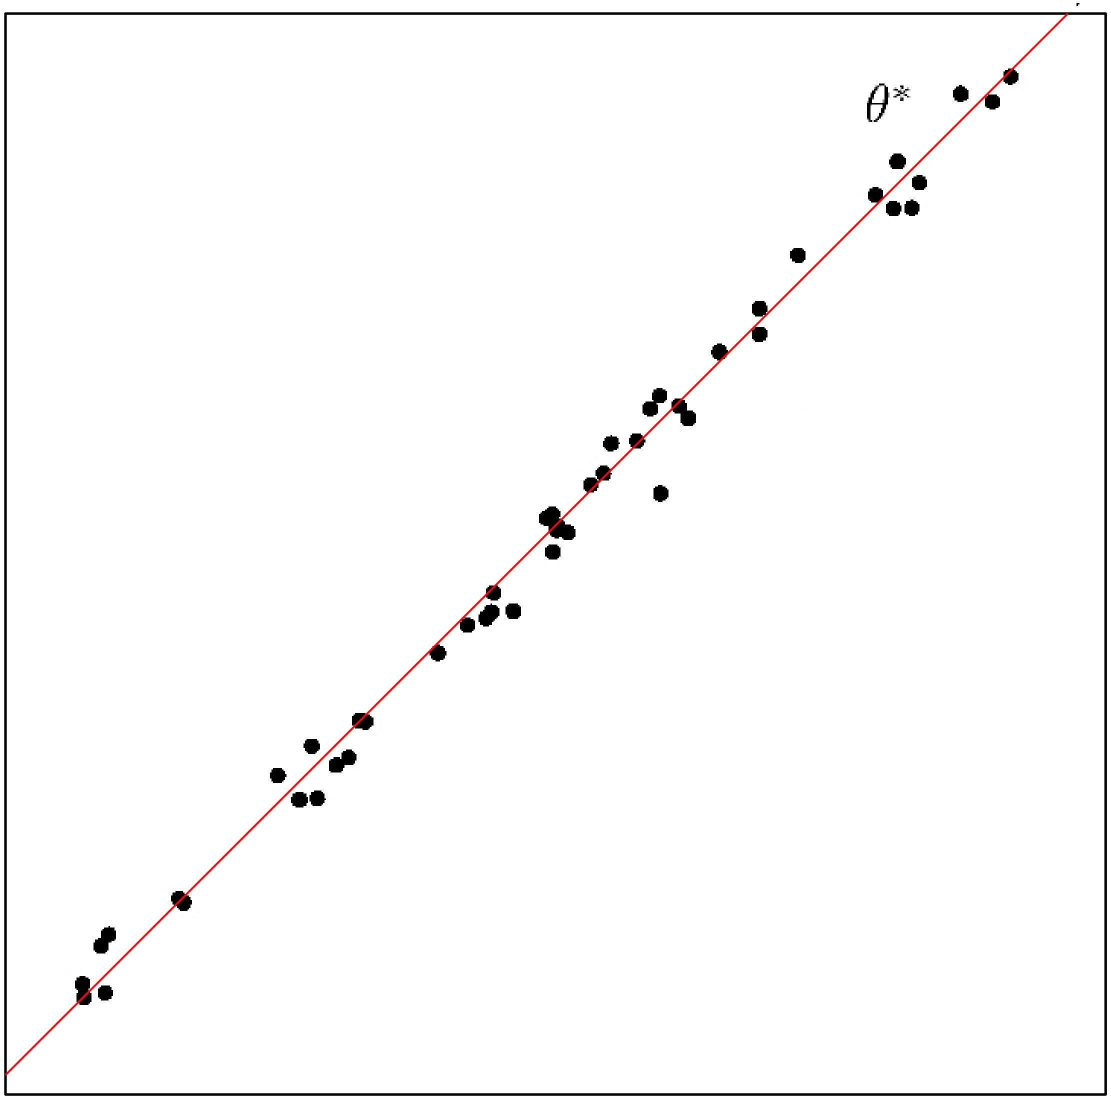

Data points

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

Find the line θ^*hich
"best fits" these points.

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

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

easily solvable by SVD -

Singular Value Decomposition

Least Squares are Well established

Adrien-Marie Legendre

Published least squares (*moindres quarrés*) in 1805.

Carl Friedrich Gauss

Developed least squares in 1795.

General Case with Outliers, Example 1

Example 1

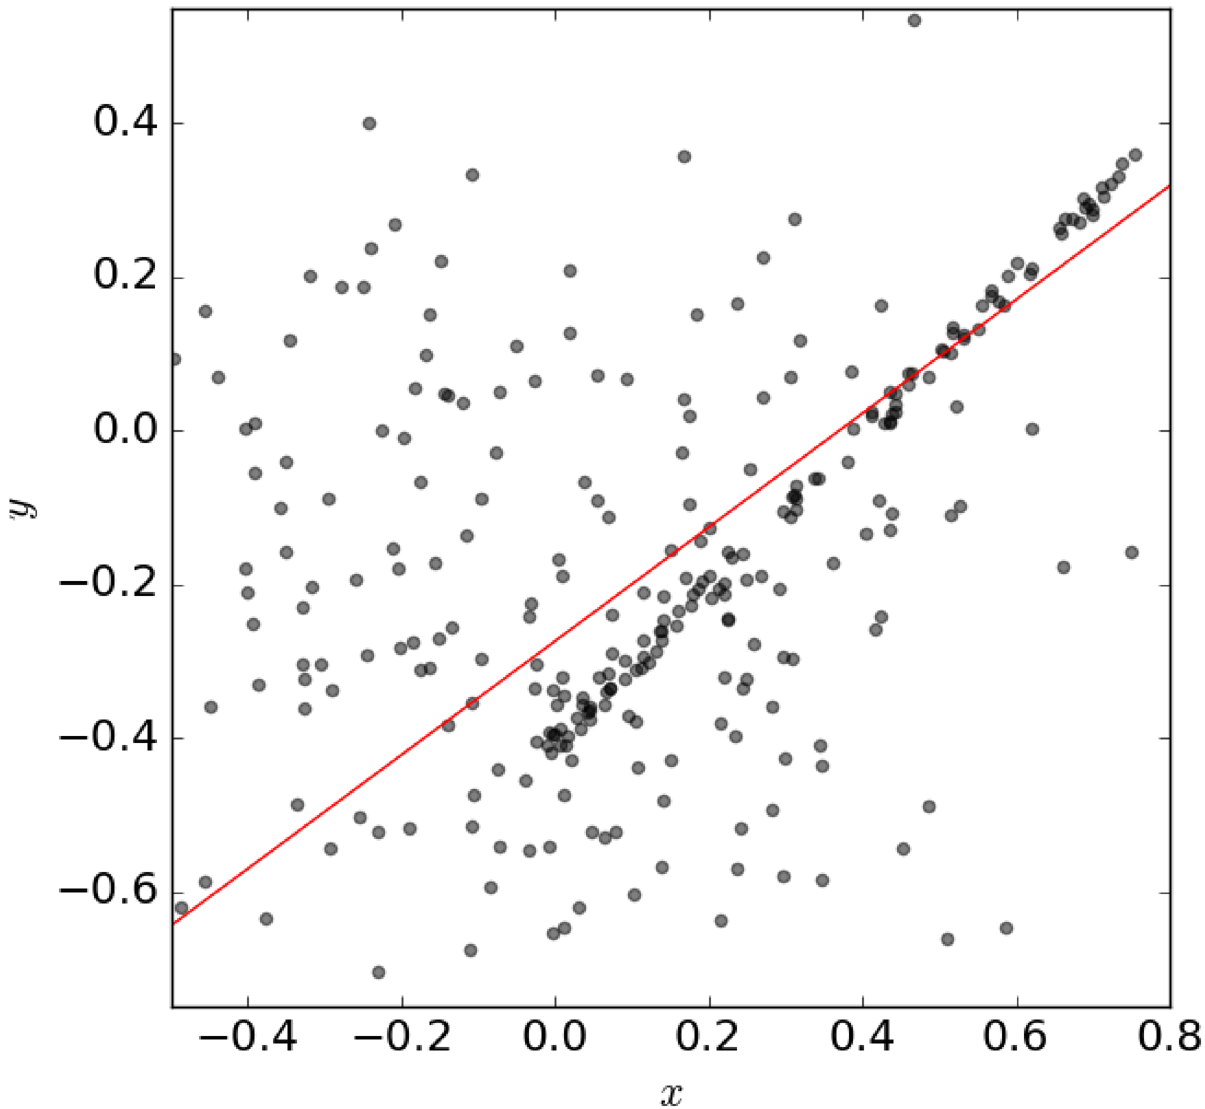

Least squares fit

Example 2

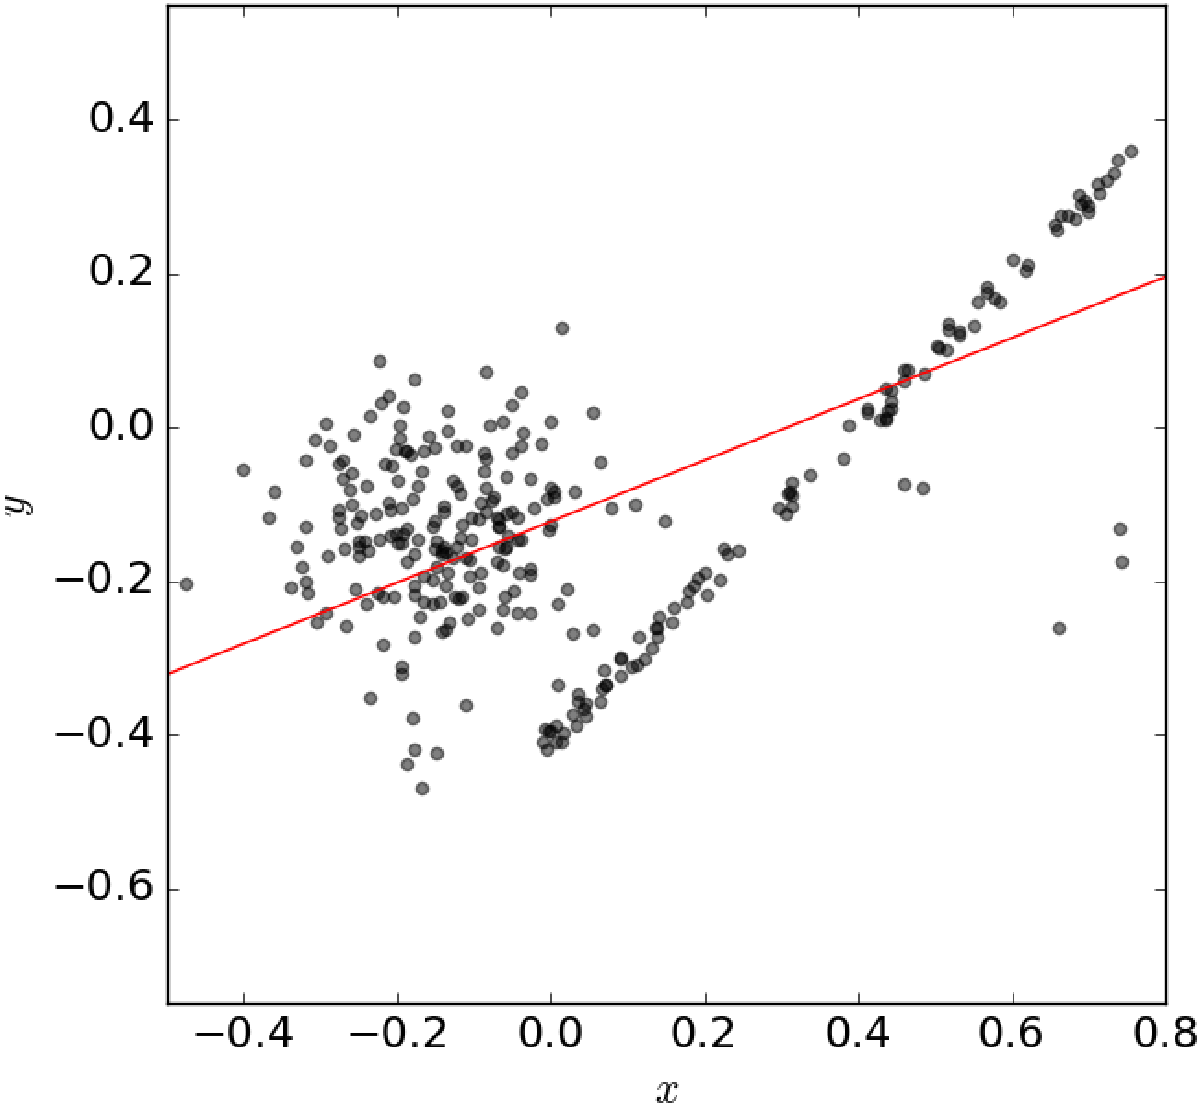

Least squares fit

RANdom SAmple Consensus - RANSAC

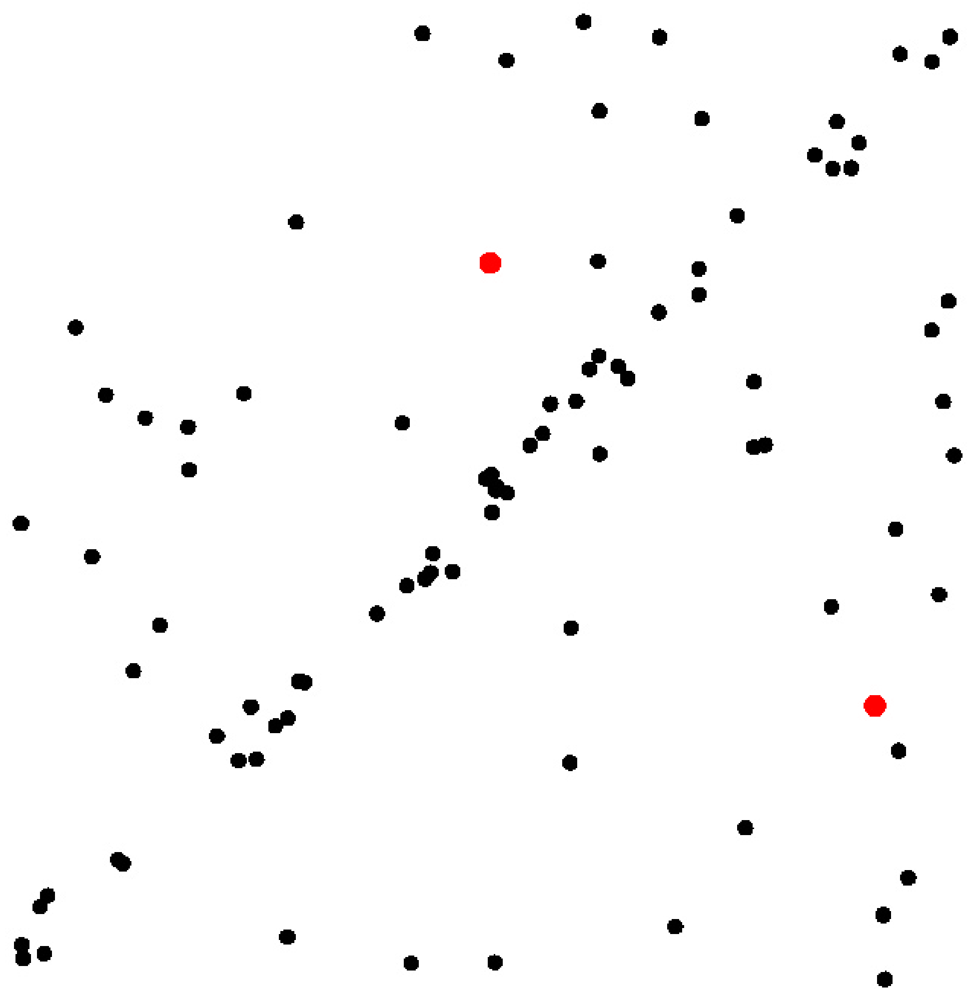

1. Select sample of m points at random (here $m=2$)

RANSAC

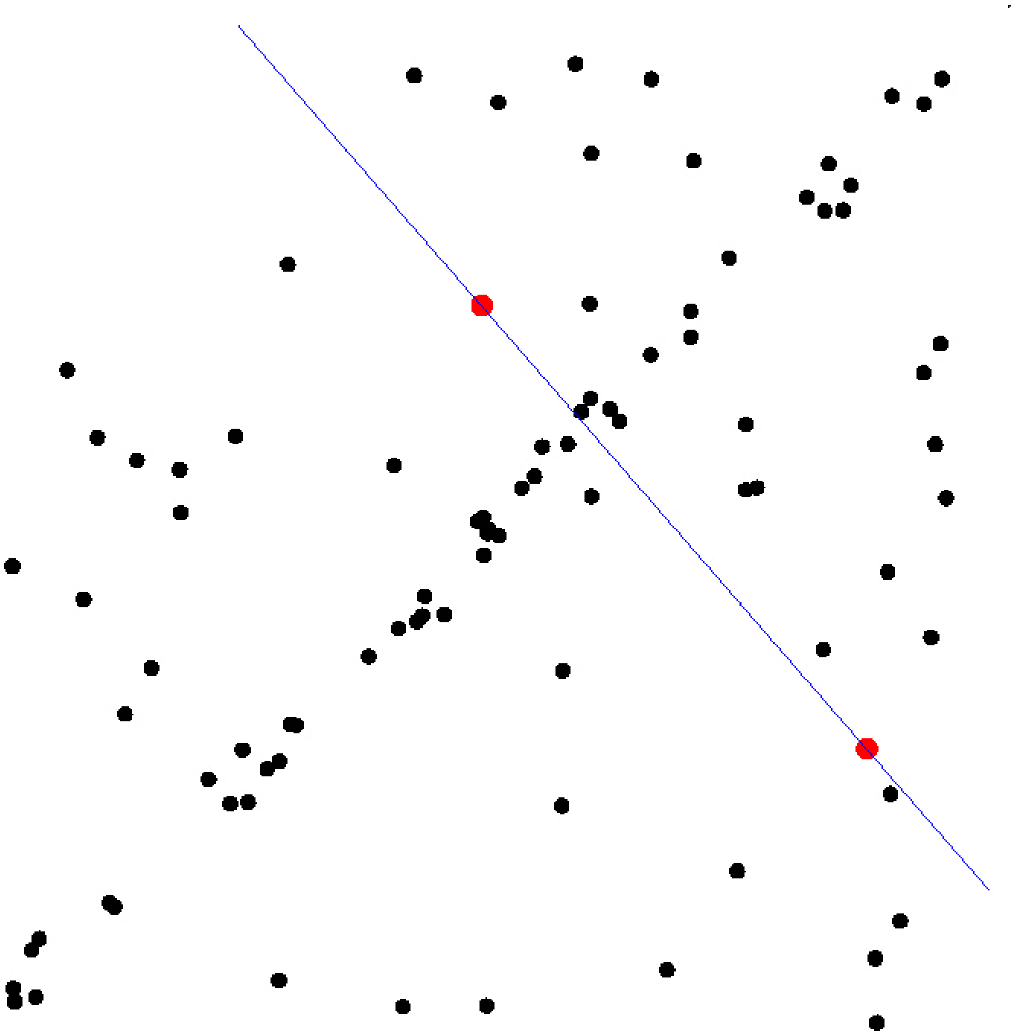

1. Select sample of m points at random

2. Estimate model parameters from the data in the sample

RANSAC

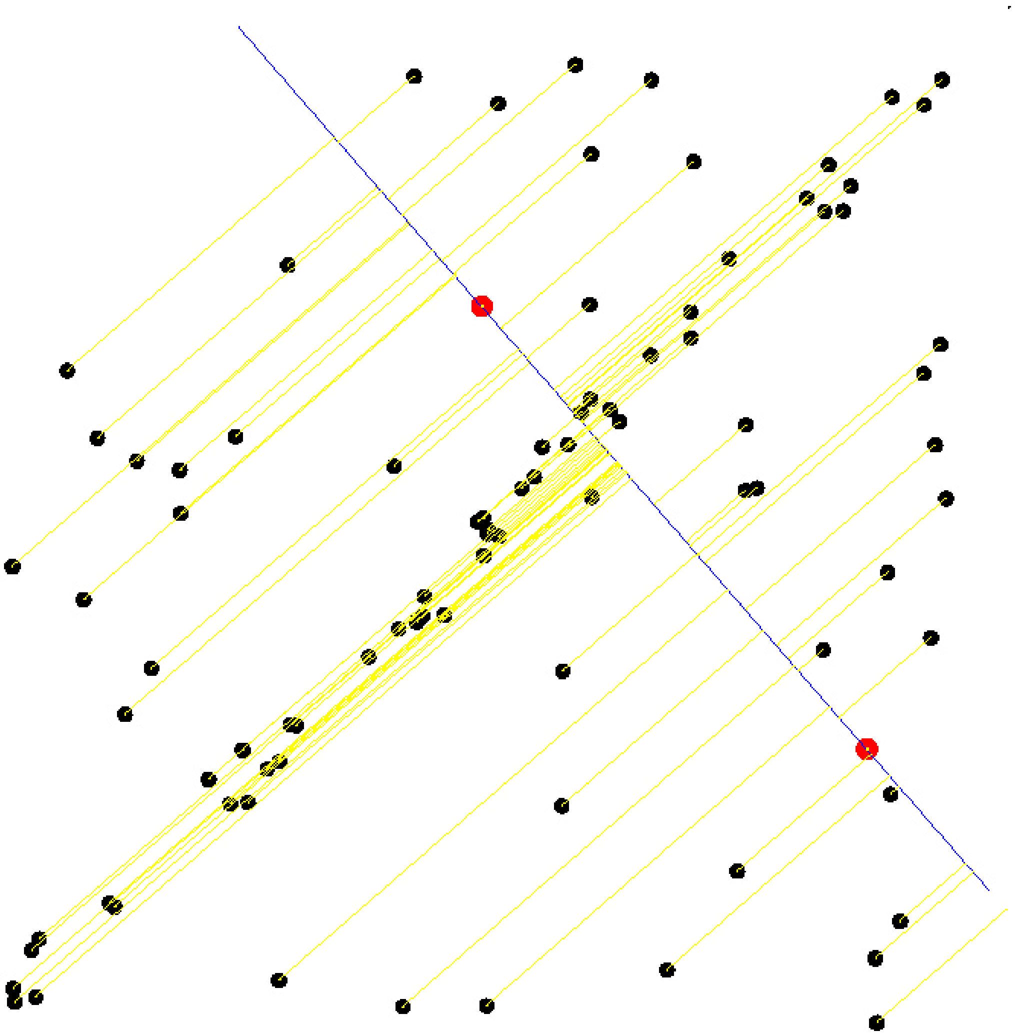

1. Select sample of m points at random

2. Estimate model parameters
from the data in the sample

3. Evaluate the distance from model for each data point

RANSAC

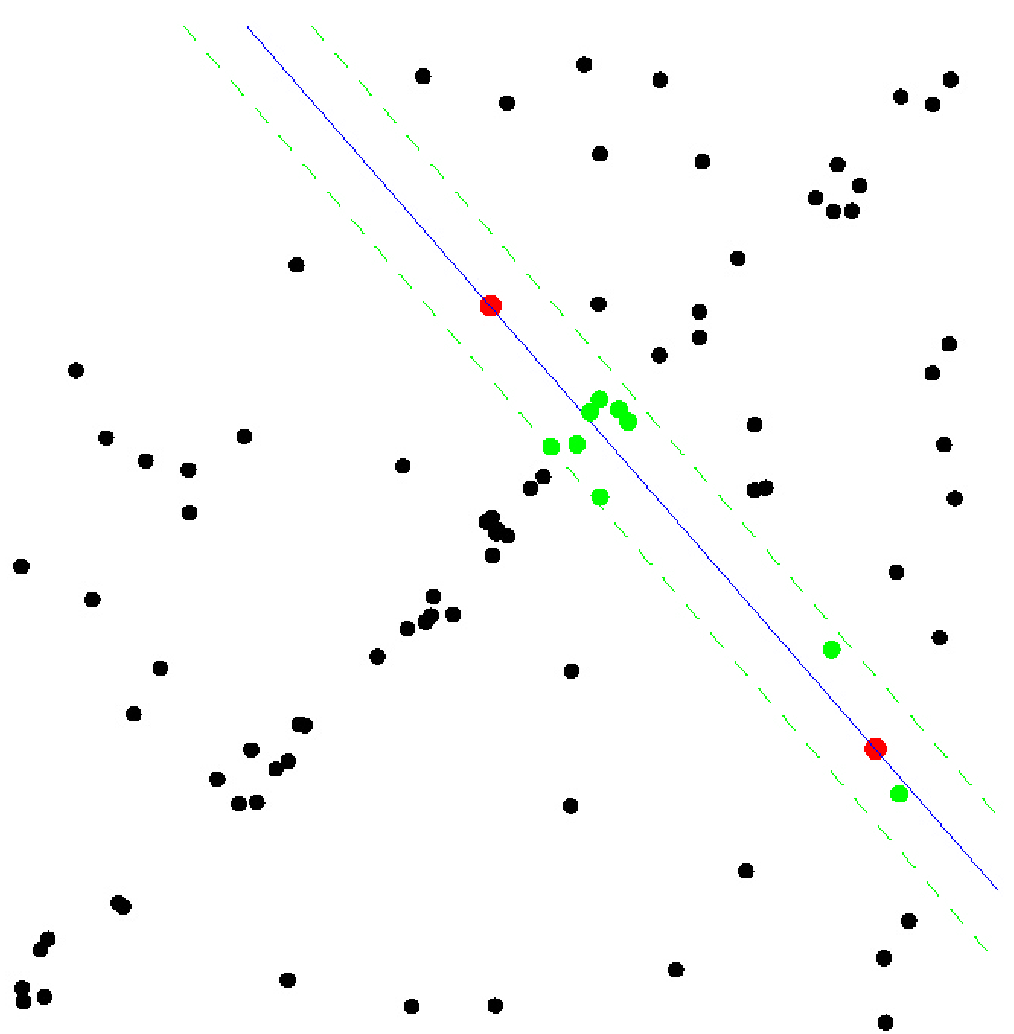

1. Select sample of m points at random

2. Estimate model parameters from the data in the sample

3. Evaluate the distance from model for each data point

4. Select data that support the current hypothesis

RANSAC

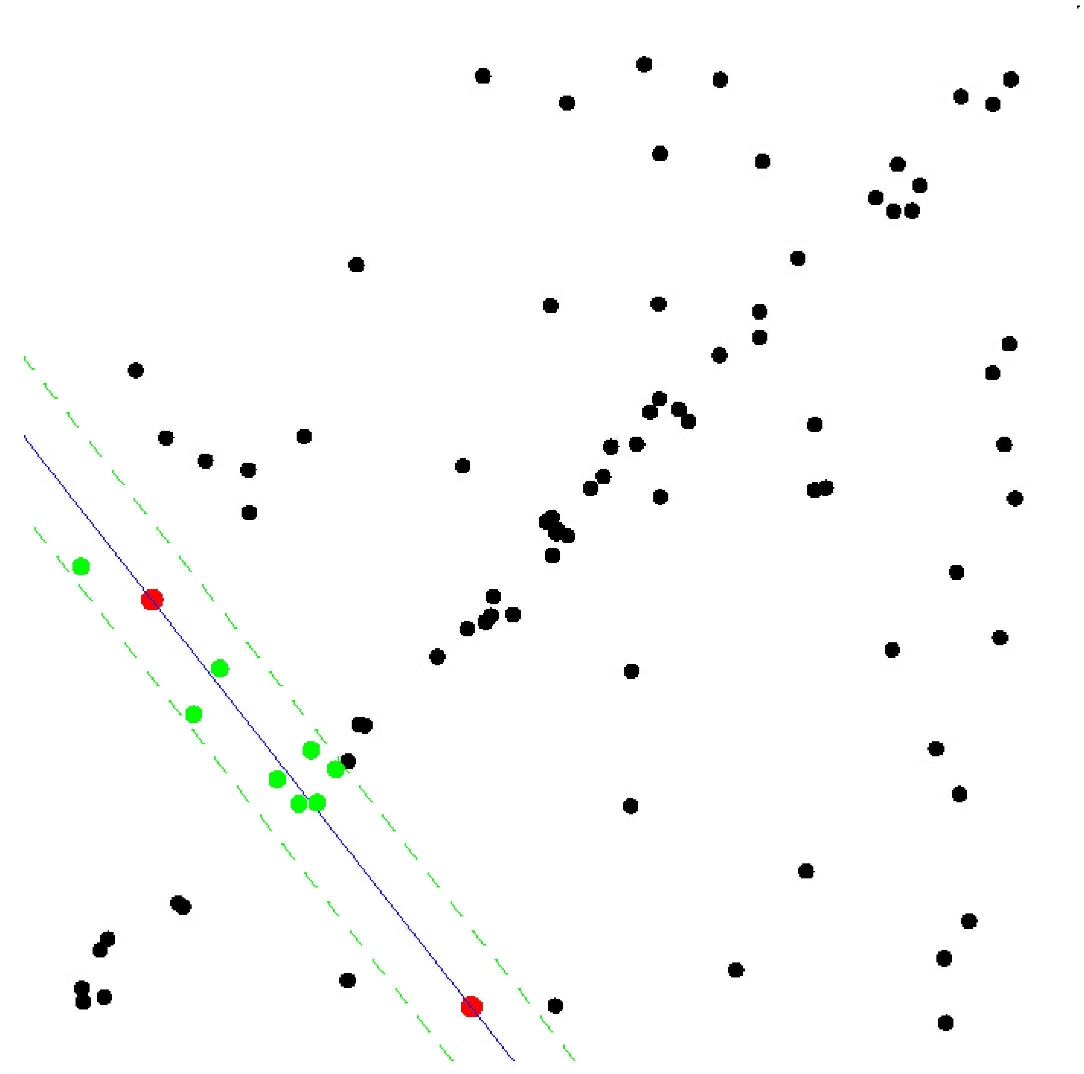

1. Select sample of m points at random
2. Estimate model parameters from the data in the sample
3. Evaluate the distance from model for each data point
4. Select data that support the current hypothesis
5. Repeat sampling

RANSAC

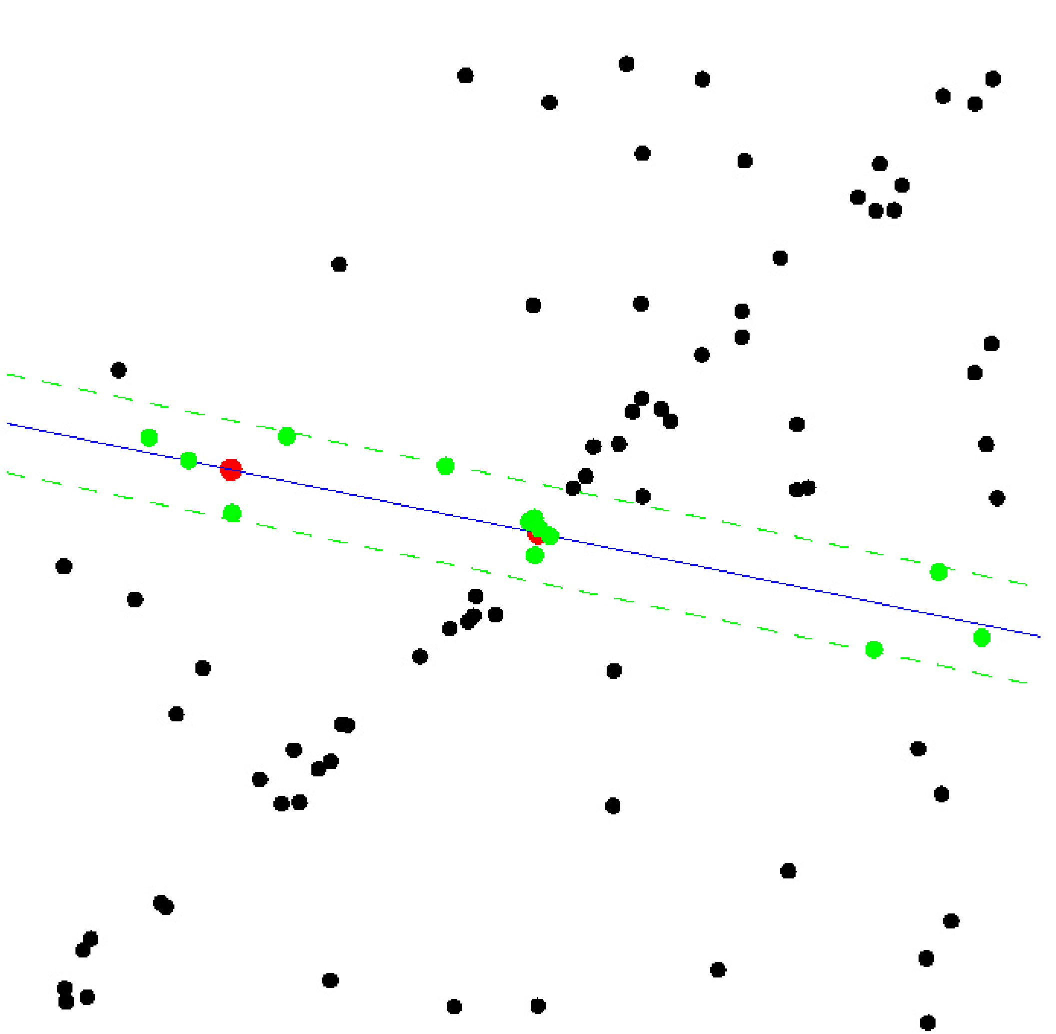

1. Select sample of m points at random
2. Estimate model parameters from the data in the sample
3. Evaluate the distance from model for each data point
4. Select data that support the current hypothesis
5. Repeat sampling

RANSAC

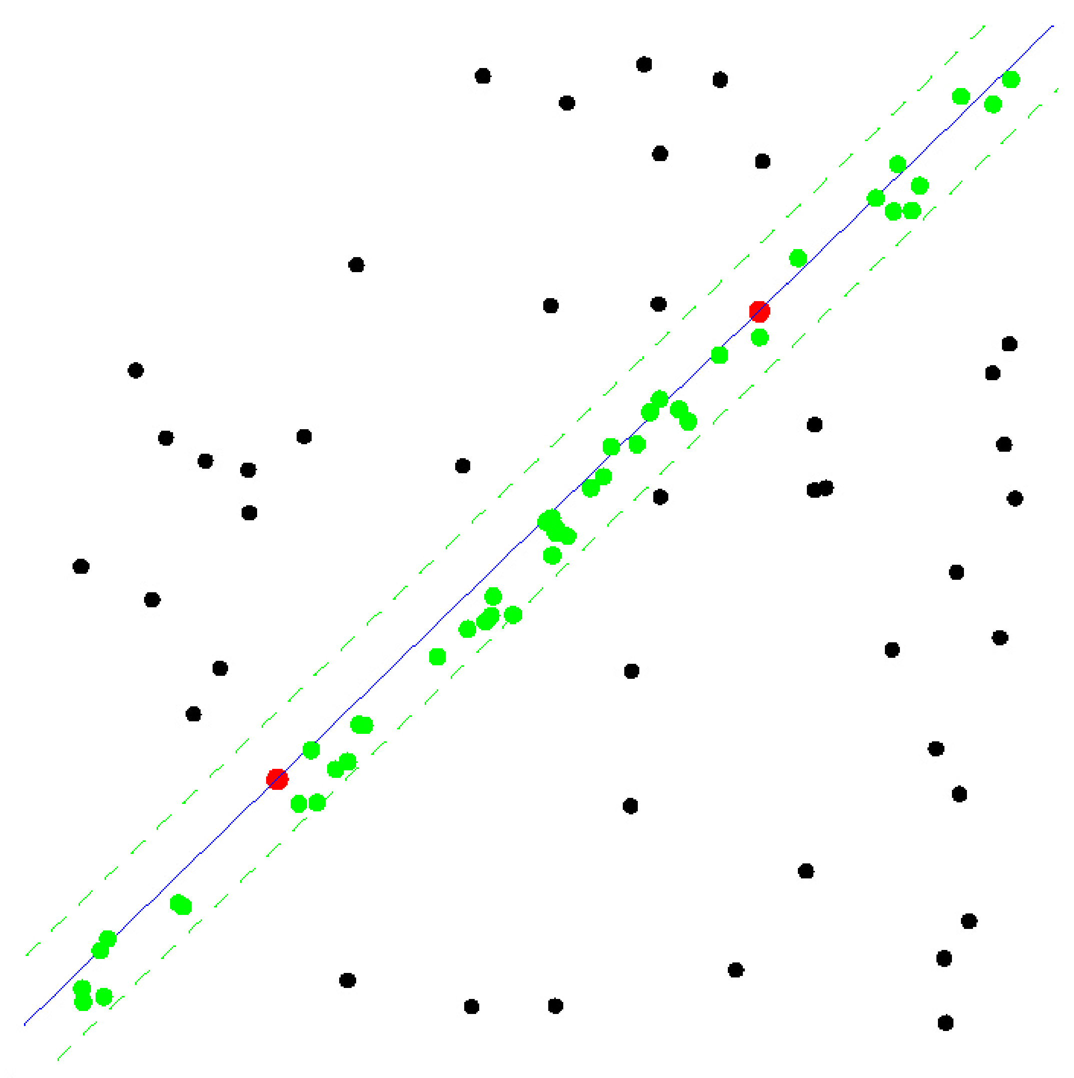

1. Select sample of m points at random
2. Estimate model parameters from the data in the sample
3. Evaluate the distance from model for each data point
4. Select data that support the current hypothesis
5. Repeat sampling until a well-supported model is found

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

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
$\theta^* \leftarrow \theta,\ J^* \leftarrow J(\theta)$
7: $iter \leftarrow iter + 1$
8: **until** $P(\text{better solution exists}) = f(|\mathcal{X}|, J^*, iter) < 1 - \eta$
9: Compute θ^* from all inliers $\mathcal{X}_{in}$: $\theta^* \leftarrow \text{LocalOptimization}(\mathcal{X}_{in}, \theta^*)$

RANSAC termination – how many samples?

- m Size of sample
- $\epsilon = Q/N$ Inlier ratio
- η Confidence (= probability a better solution was overlooked)
- k required number of samples

Finding the solution with confidence η requires at least

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

samples. For small η, ϵ^{m}

$$k \cong \eta/\epsilon^{m} = \eta\epsilon^{-m}$$

since $e^{-x} \cong 1 - x$

RANSAC Properties

- extremely popular (1981 paper has ~ 30 000 citations in Google Scholar)

Pros:

- used in many applications
- knowledge of the percentage of inliers not needed
- works, unlike least median of squares, for and any inlier ratio
- provides a probabilistic guarantee for the solution

Cons:

- slow if inlier ratio ϵ low and model dimensionality m high
- noise level, i.e. the inlier-outlier threshold, assumed known
- assumes unstructured outlier and inlier distribution
- does not validate the model it produces,
- (experimental observation) RANSAC takes several times longer than theoretically expected, due to noise – not every all-inlier sample generates a good hypothesis: $P(\text{inlier sample}) \neq P(\text{good model estimate})$

RANSAC Issues and Fixes

- **Cost function:** MLESAC, Huber loss, ...
- **Accuracy** (parameters are estimated from minimal samples).
  addressed by: **LO-RANSAC**, Locally Optimized RANSAC
- **Outlier threshold** σ (how to set it in advance? Or, how to avoid setting it?):
  addressed by: Least median of Squares, MINPRAN, **MAGSAC**, ...
- **Unstructured outliers and inliers**
  addressed by: GC-RANSAC, Graph Cut RANSAC and NAPSAC
- **Ignores info about data quality:**
  addressed by PROSAC
- **Speed:** Running time grows with number of data points, number of iterations
  addressed by: WaldSAC (Randomized evaluation RANSAC),
- **Correctness of the results. Degeneracy.** Solution: DegenSAC.

Accuracy

Locally Optimized RANSAC (LO-RANSAC): Problem Intro

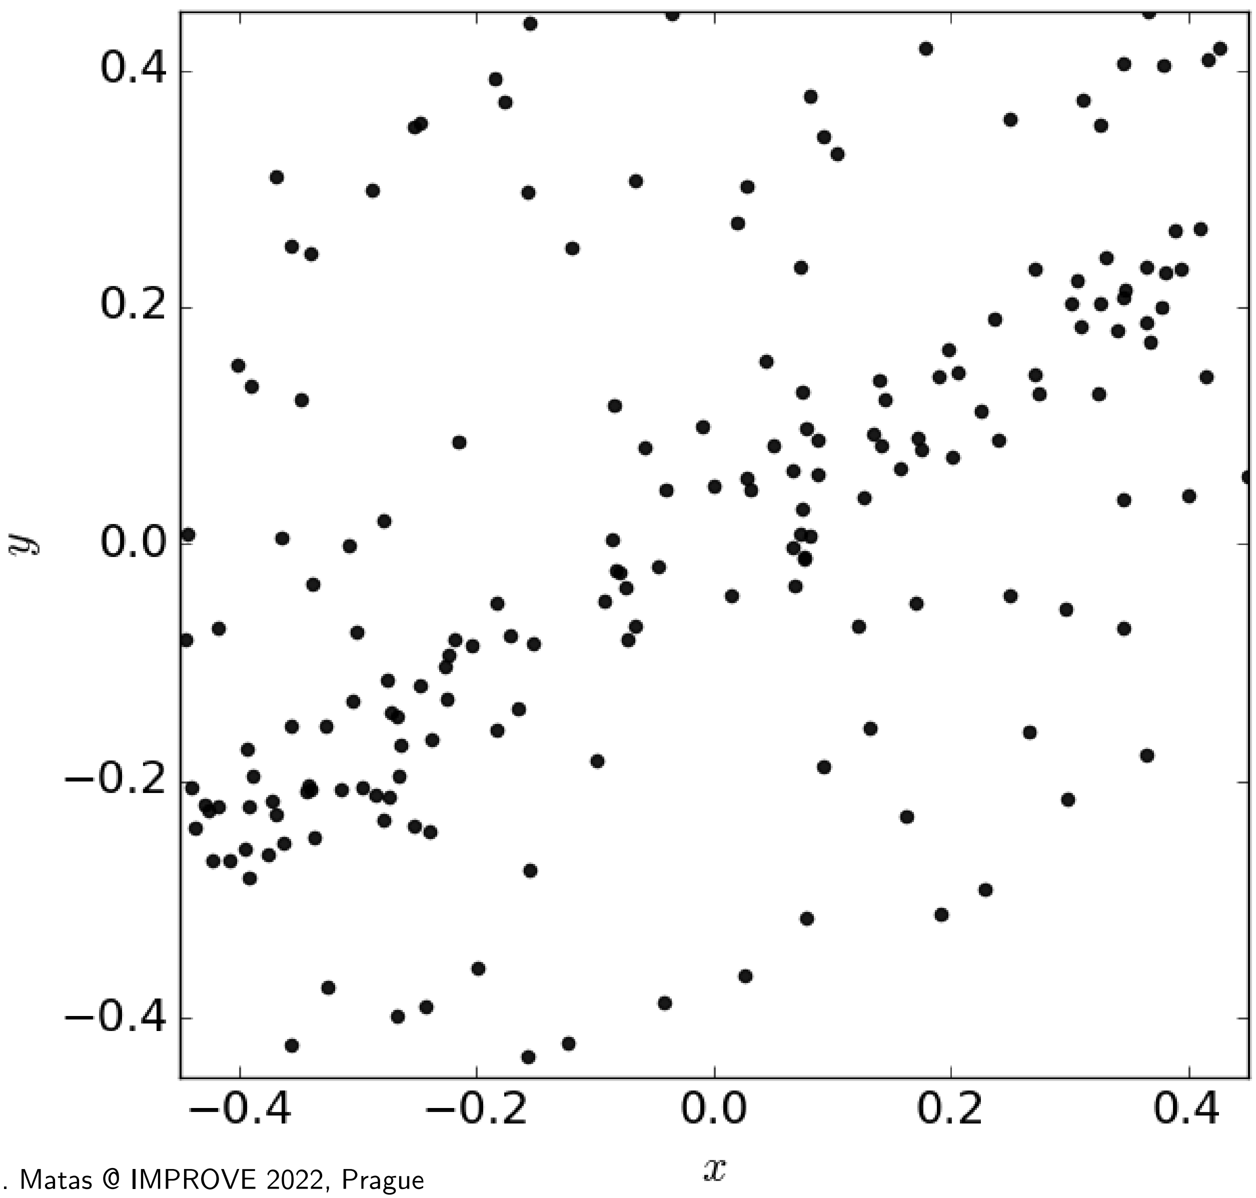

Data: 200 points

LO-RANSAC: Problem Introduction

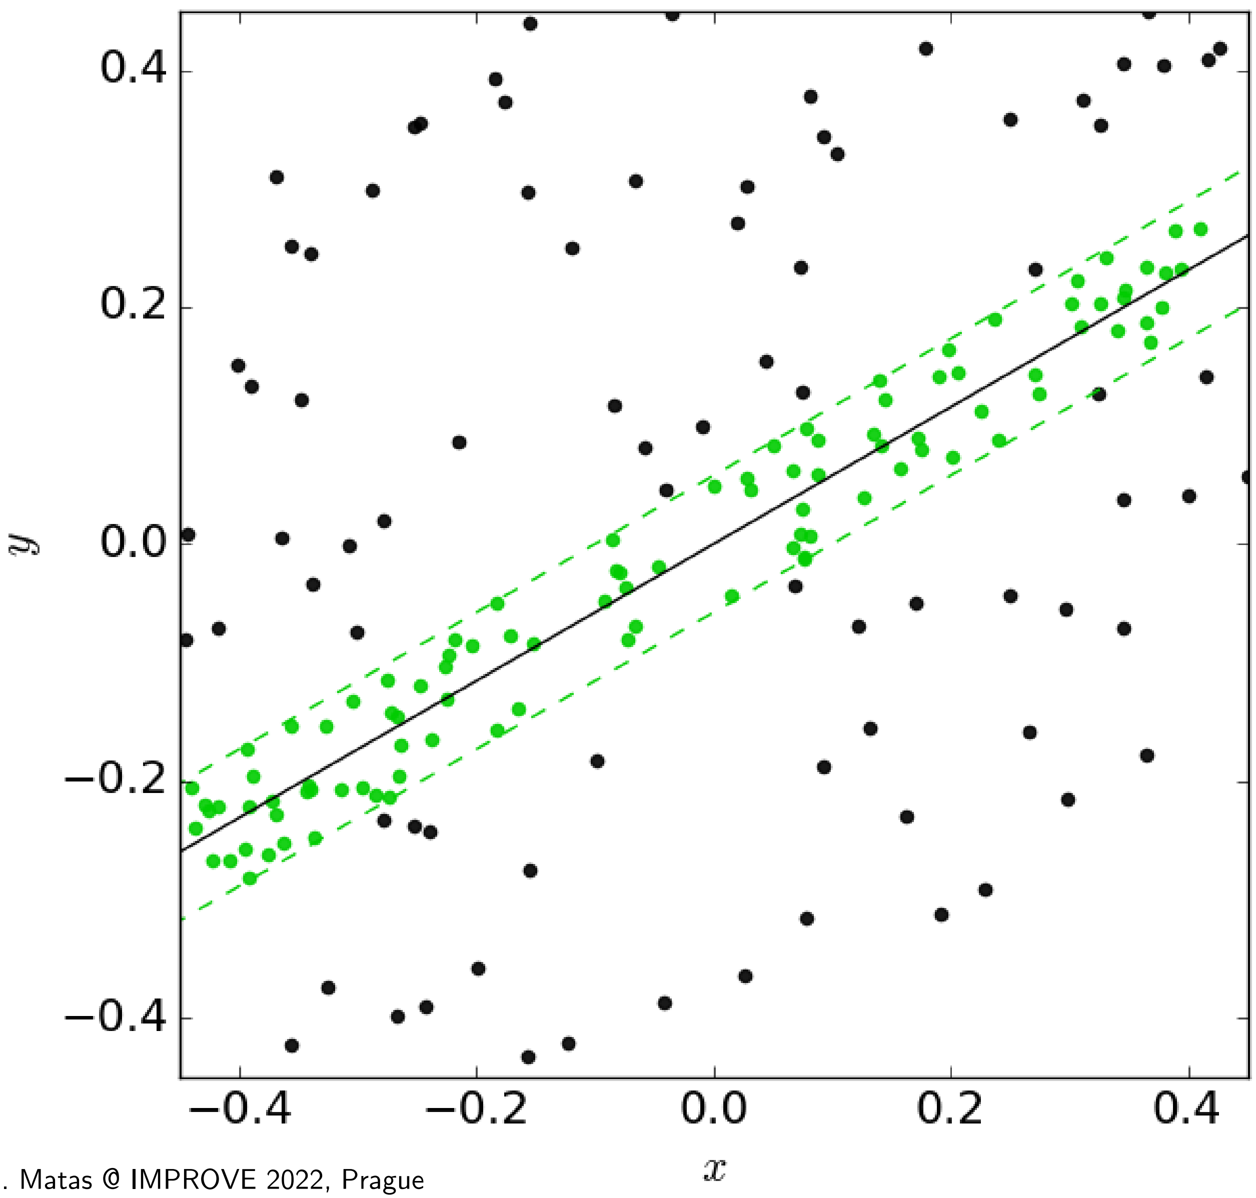

Data: 200 points

Model, 100 inliers

LO-RANSAC: Problem Introduction

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

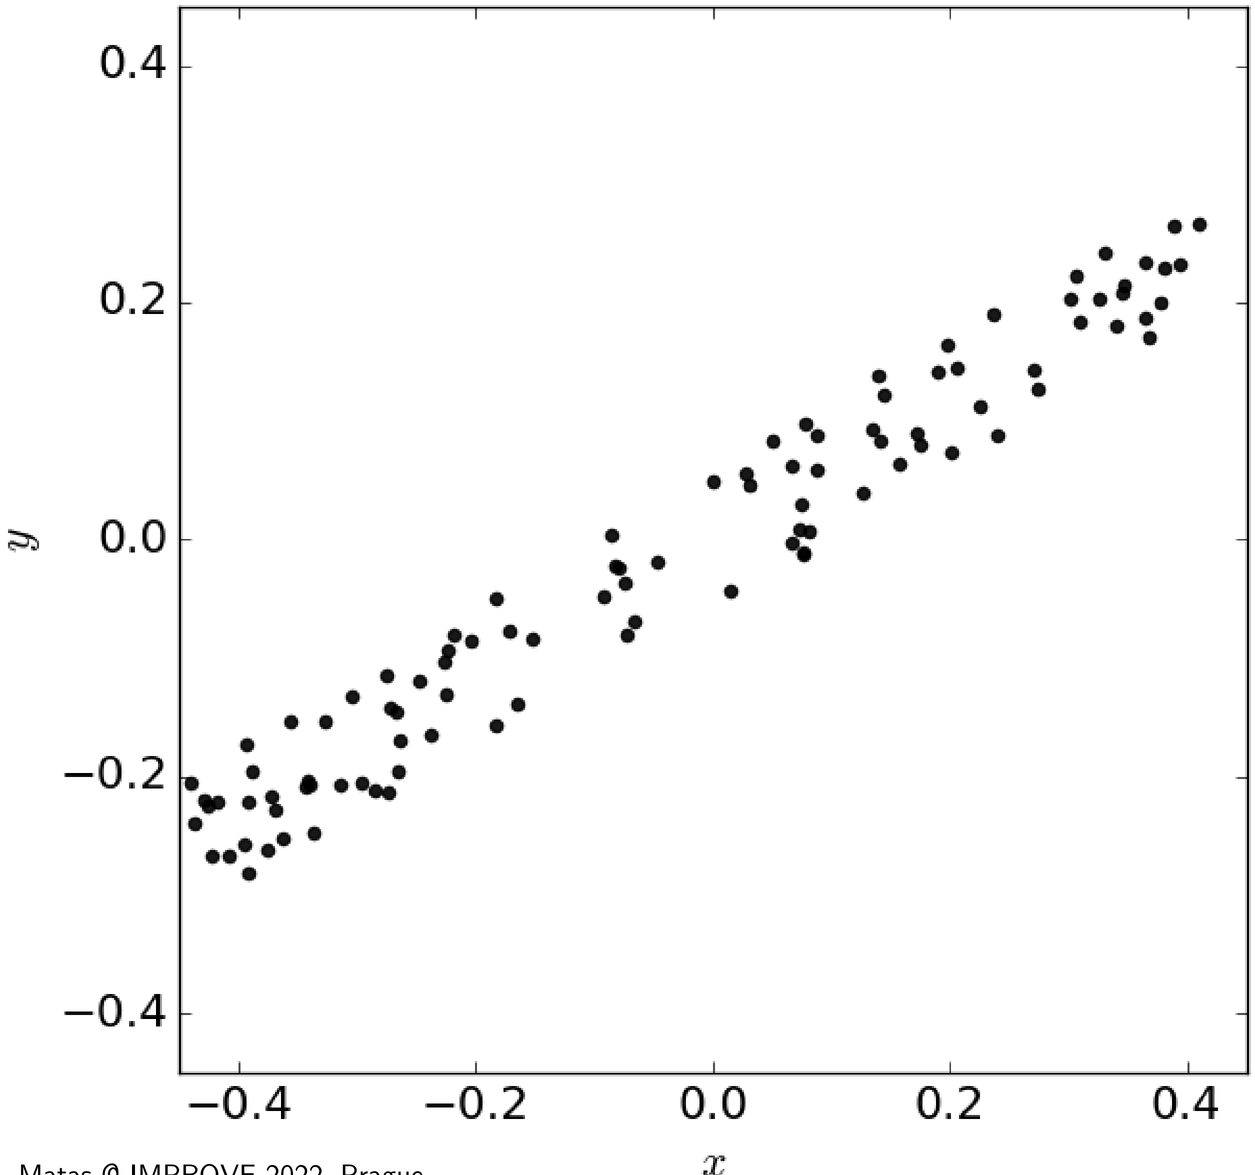

LO-RANSAC: Problem Introduction

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

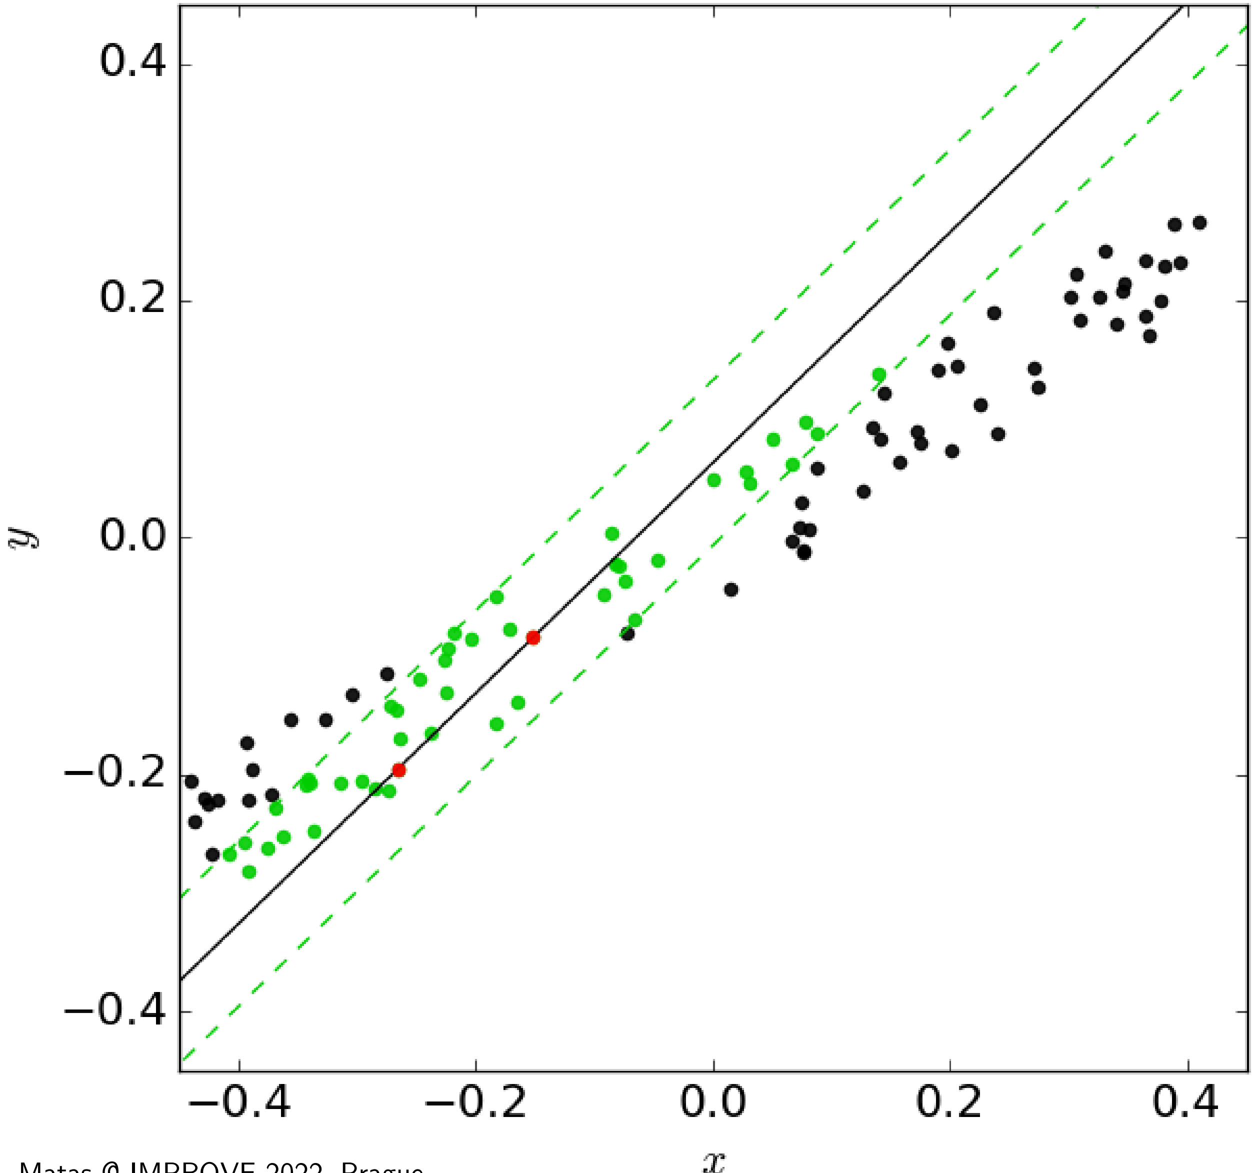

RANSAC

Hypothesis generation from 2 points

Will every two points generate the whole inlier set?

This sample:
NO. 45 inliers.

LO-RANSAC: Problem Introduction

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

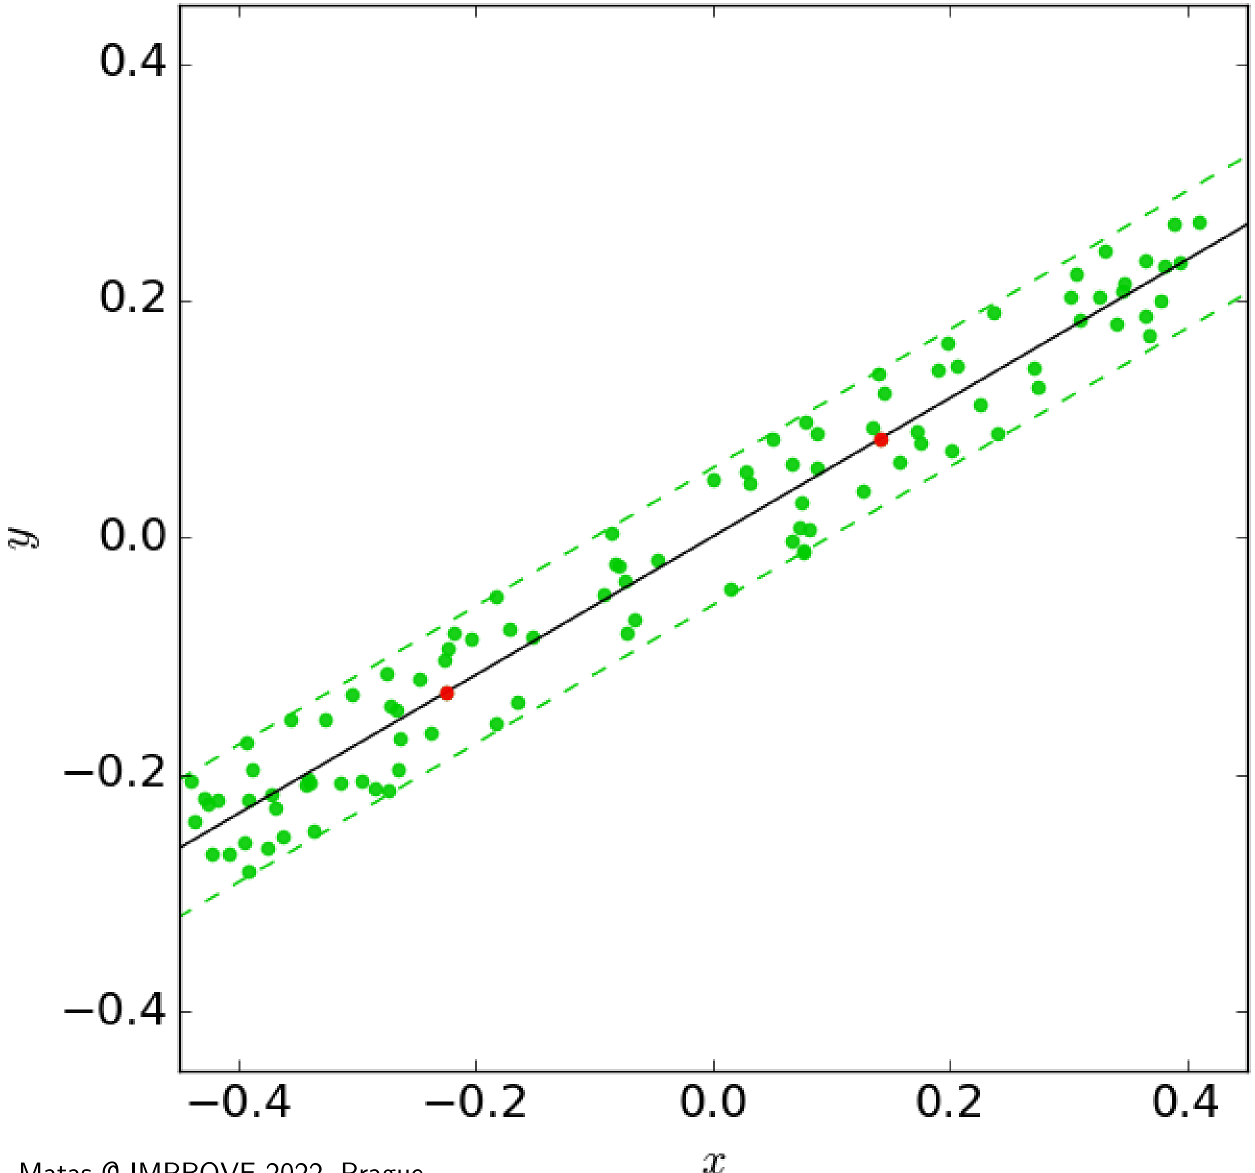

RANSAC

Hypothesis generation from 2 points

Will every two points generate the whole inlier set?

This sample:
YES. 100 inliers.

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
$\theta^* \leftarrow \theta,\ J^* \leftarrow J(\theta)$
7: $iter \leftarrow iter + 1$
8: **until** $P(\text{better solution exists}) = f(|\mathcal{X}|, J^*, iter) < 1 - \eta$
9: Compute θ^* from all inliers $\mathcal{X}_{in}$: $\theta^* \leftarrow$ LocalOptimization$(\mathcal{X}_{in}, \theta^*)$

A Locally Optimized RANSAC

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

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

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

$\Rightarrow J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$ 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 $\theta = e(S)$ — MODEL ESTIMATION
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$(\mathcal{X}_{in}, \theta)$, $J^* \leftarrow J(\theta)$ — **LOCAL OPTIMIZATION**
7: $iter \leftarrow iter + 1$
8: **until** P(better solution exists) $= f(|\mathcal{X}|, J^*, iter) < \eta$

LO-RANSAC: Example

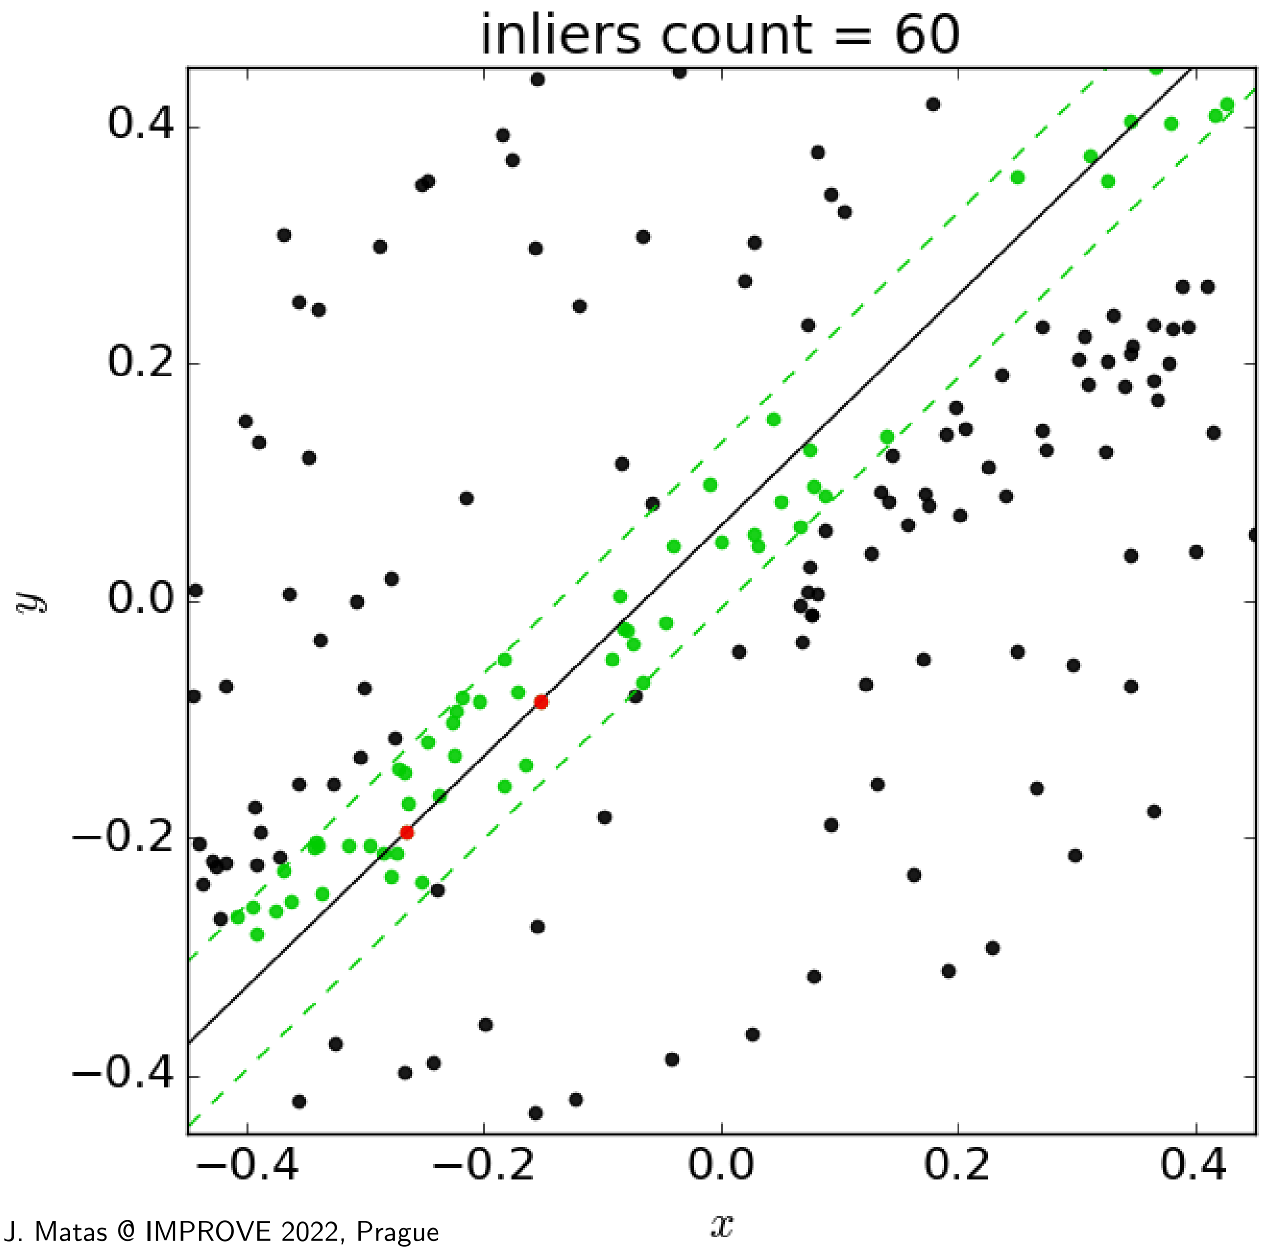

Init

LO-RANSAC: Example

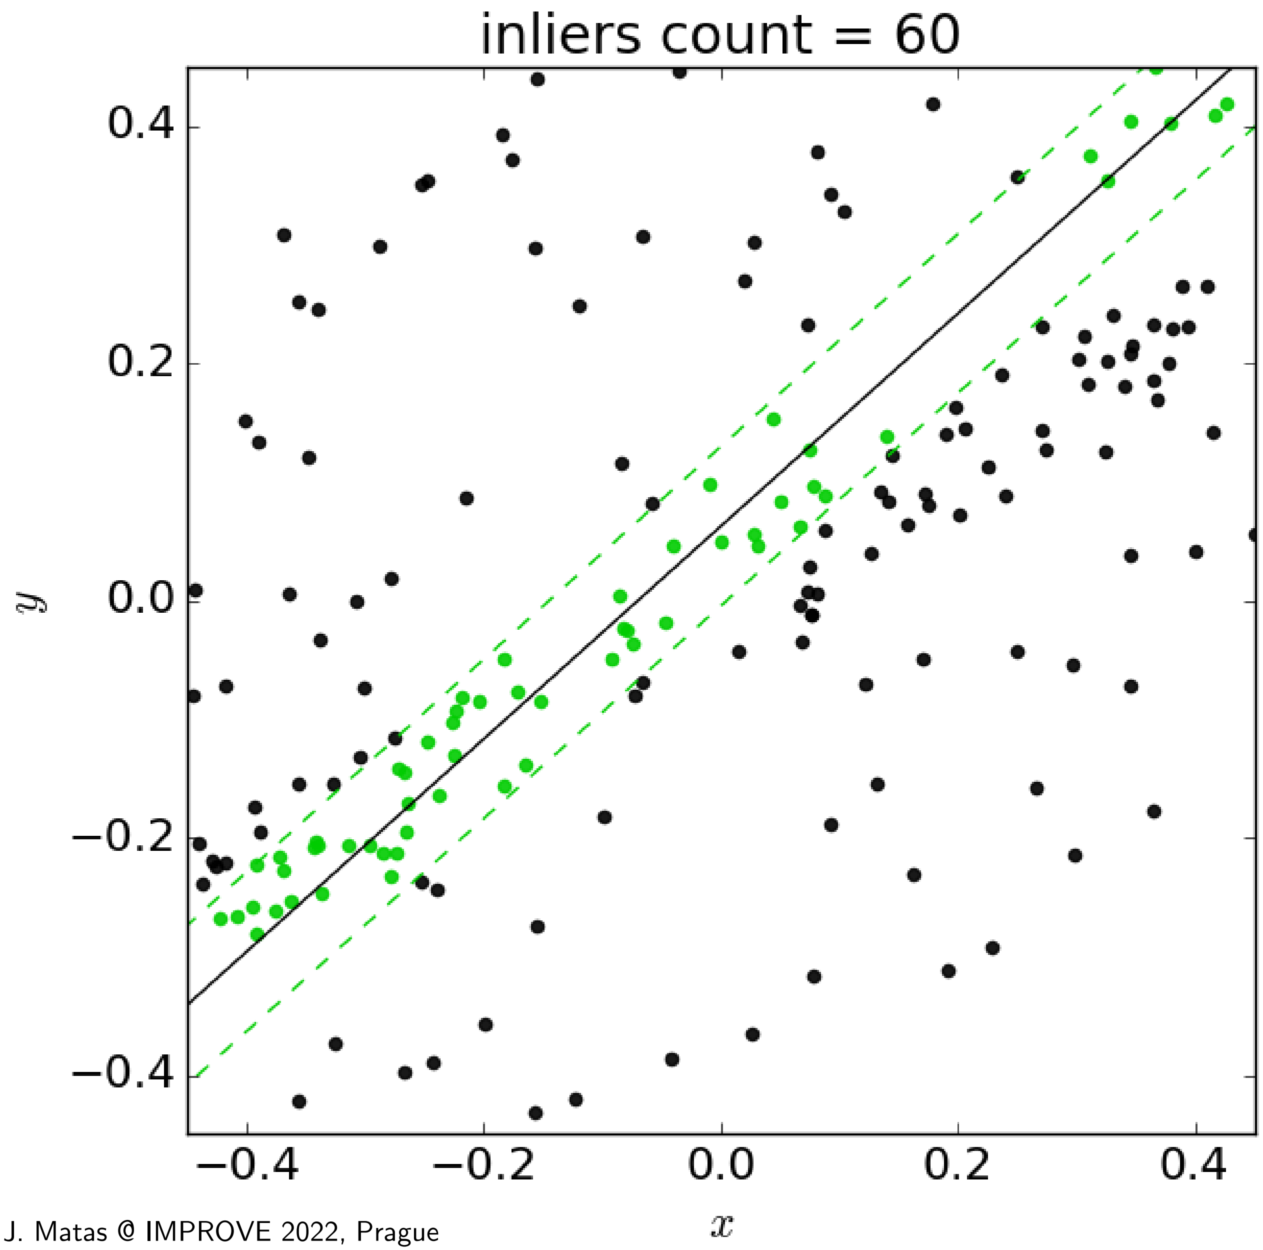

Init

Iteration 1

LO-RANSAC: Example

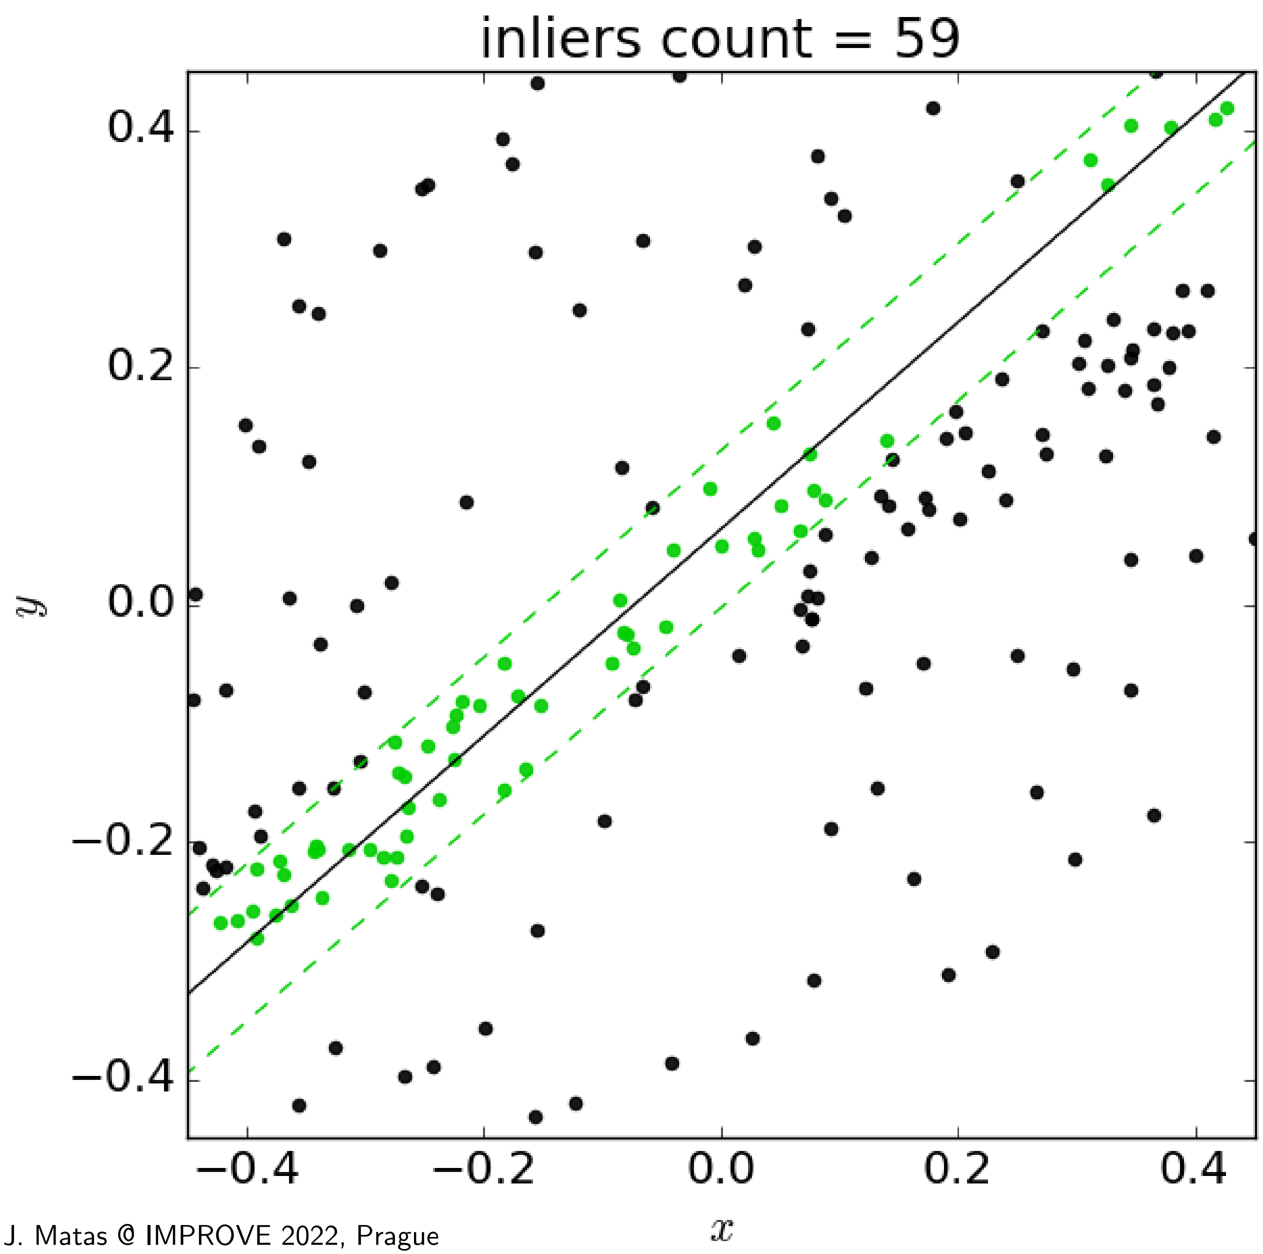

Init

Iteration 1

Iteration 2

LO-RANSAC: Example

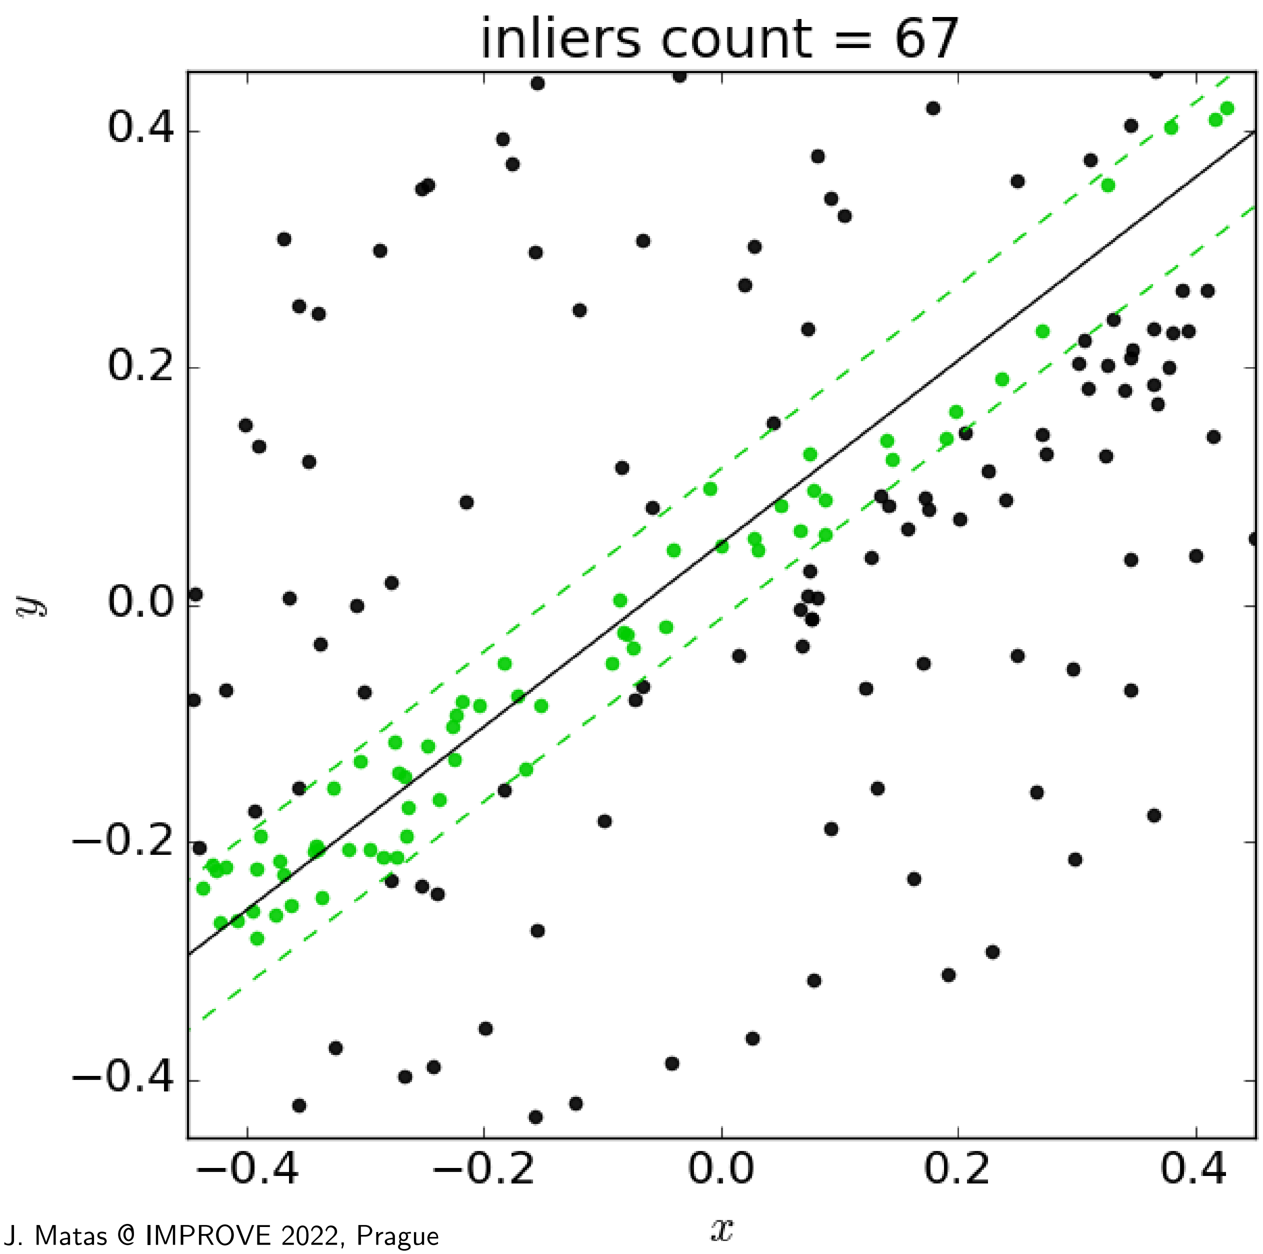

Init

Iteration 1

Iteration 2

...

Iteration 7

LO-RANSAC: Example

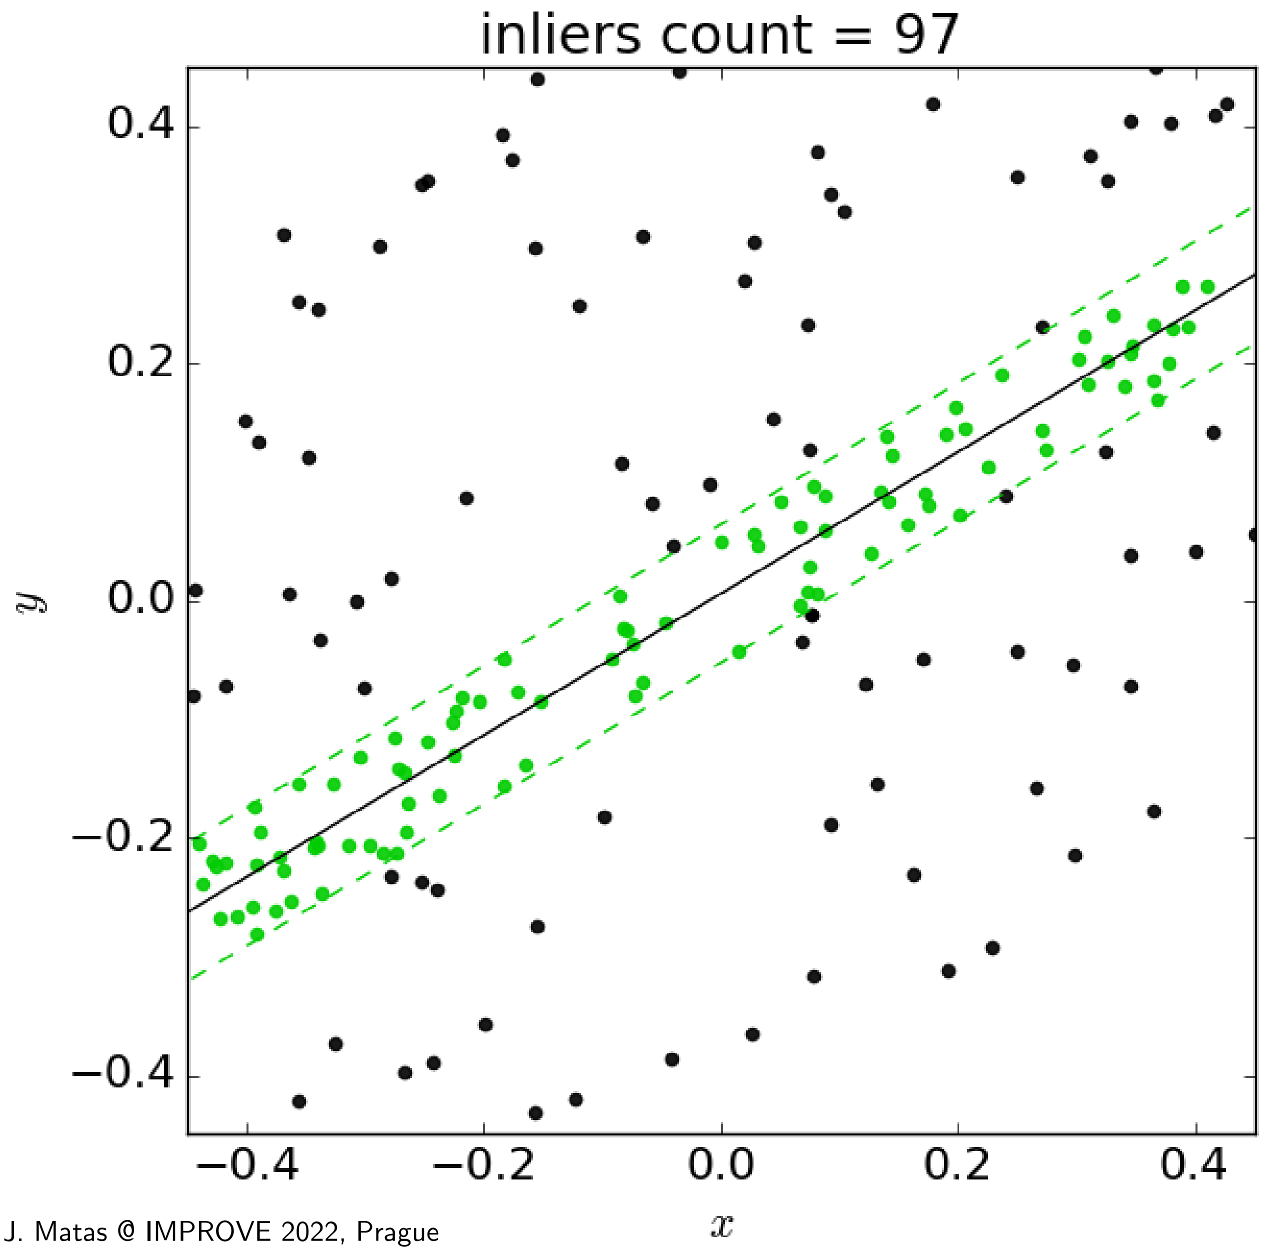

Init

Iteration 1

Iteration 2

...

Iteration 7

...

Itration 15

LO-RANSAC: Example

Comparison with model (100 inliers):

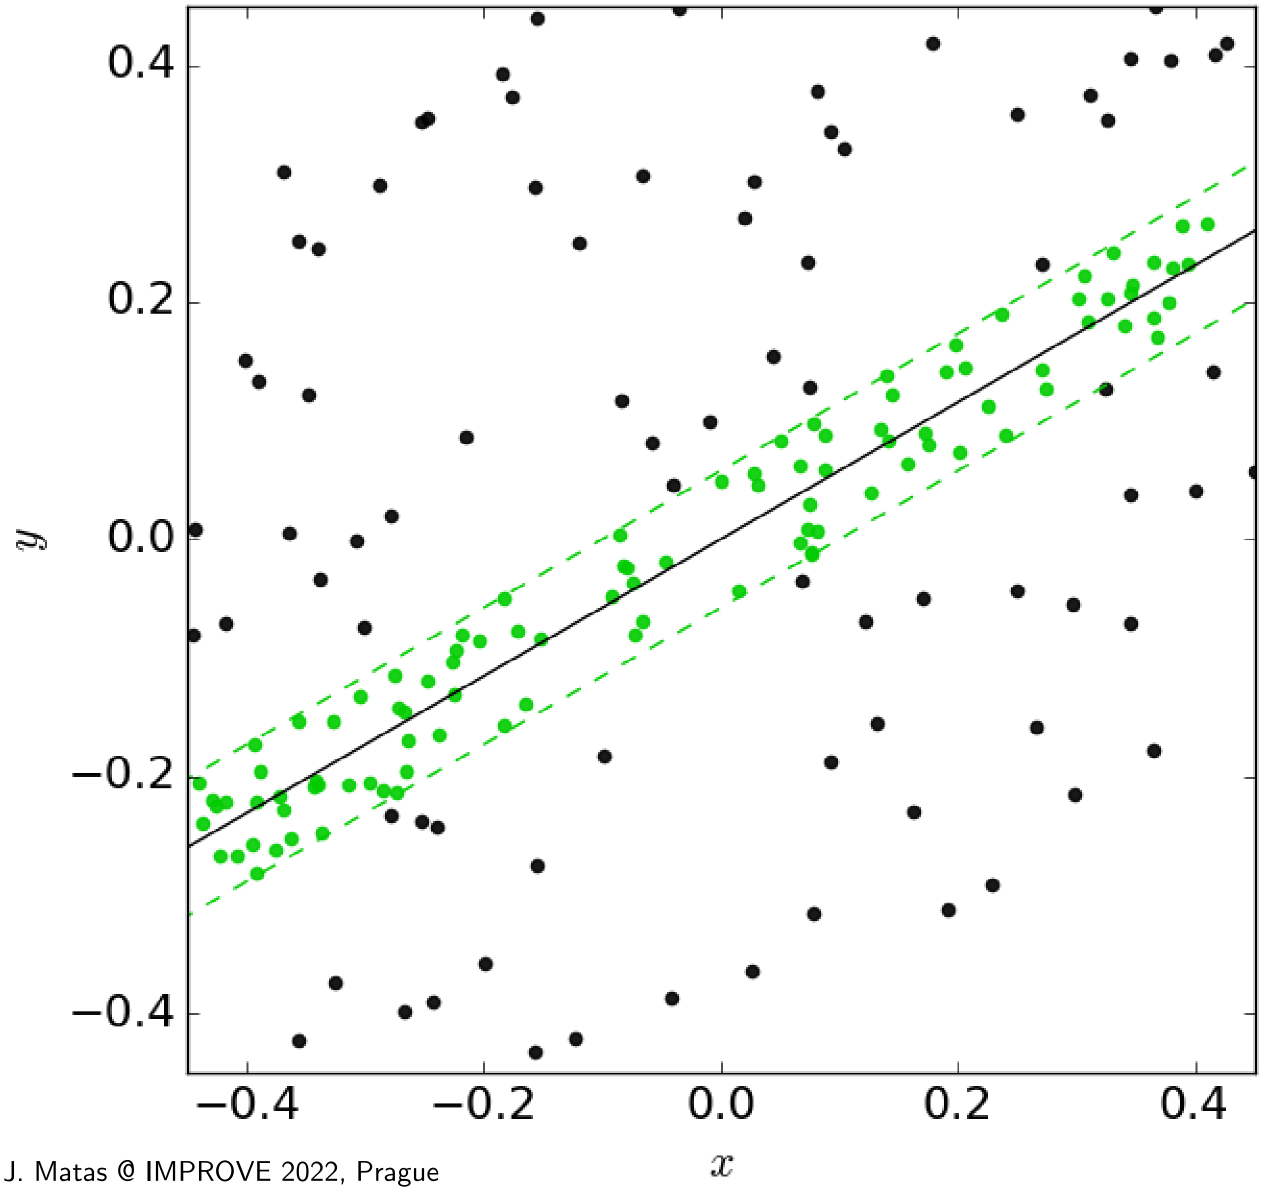

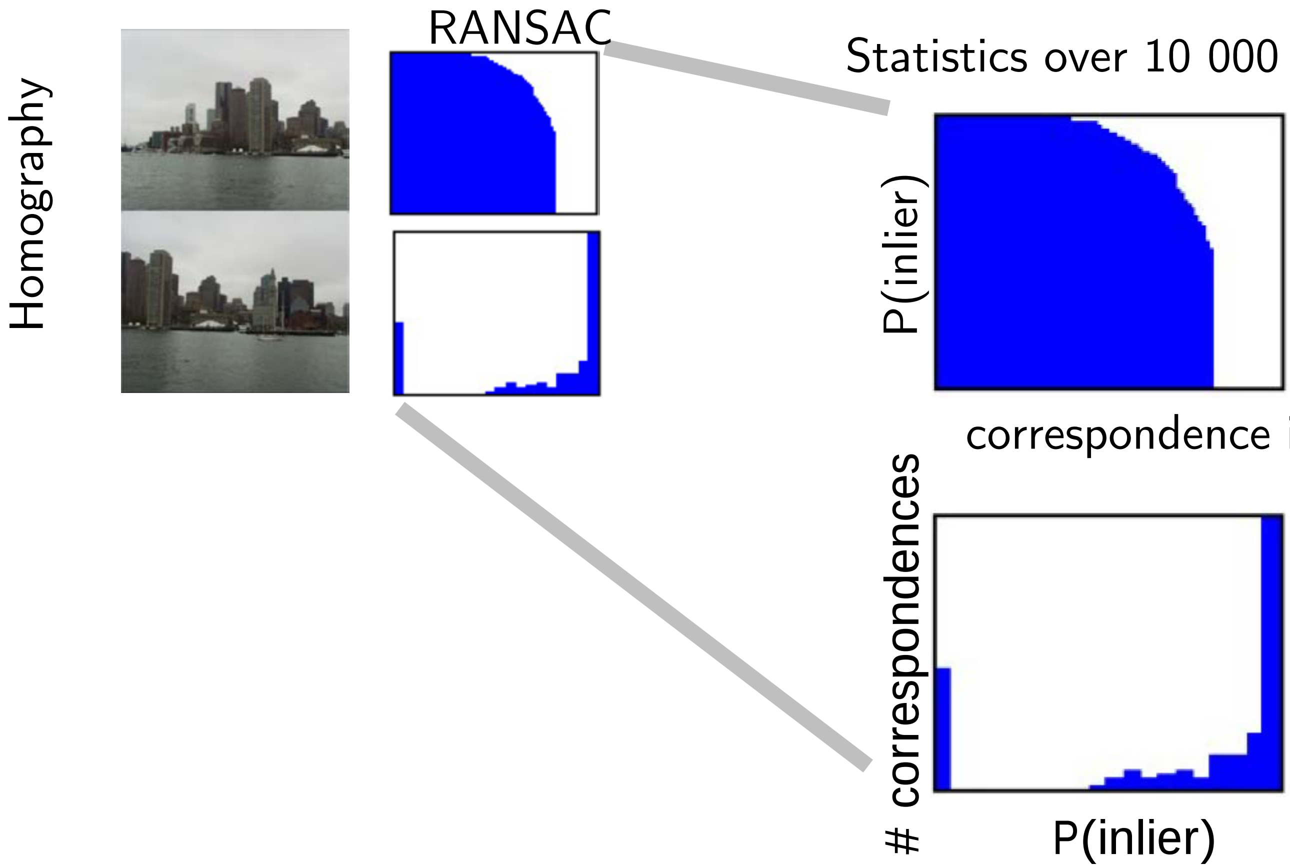

Ideally, inliers are always detected as inliers, outliers as outliers

Ideally, inliers have probability one, outliers probability zero

Lebeda, Matas, Chum: Fixing the Locally Optimized RANSAC. BMVC 2012

Stability of LO-RANSAC

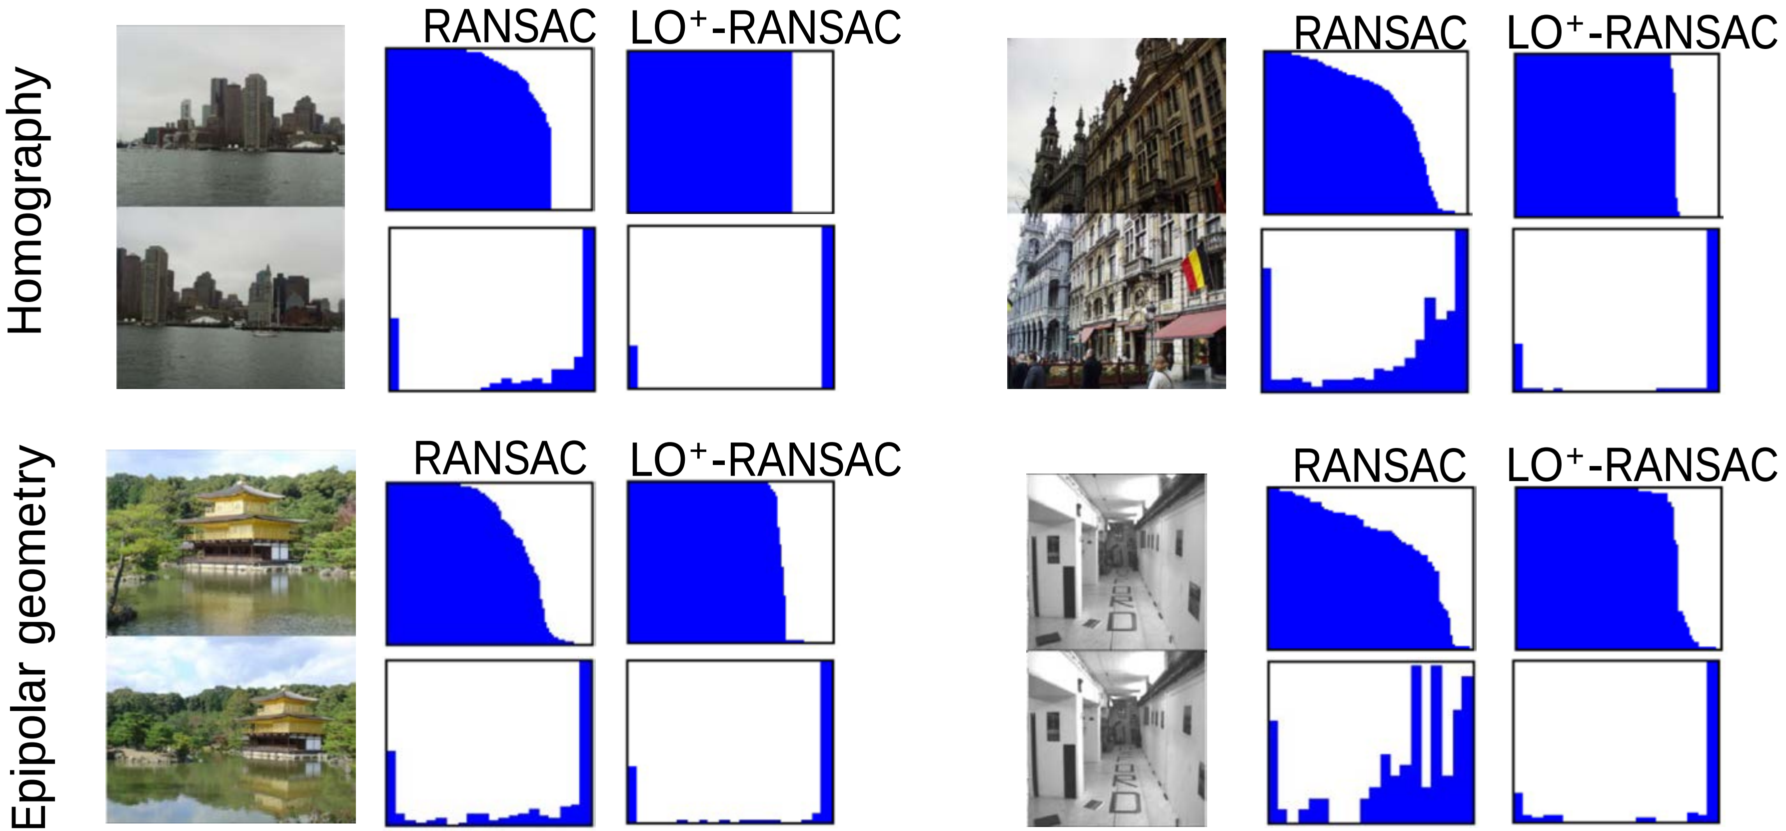

Lebeda, Matas, Chum: Fixing the Locally Optimized RANSAC. BMVC 2012

LO-RANSAC: Problem Summary

By applying local optimization (LO) to the-best-so-far hypotheses [1]:

(i) a very good agreement with theoretical performance achieved

(ii) lower sensitivity to noise and poor conditioning.

The LO is executed rarely, log(*iter*) times, it has small impact on the execution time.

If Bundle Adjustment (non-linear minimization of an appropriate loss function) used in as LO, state-of-the-art precision is achieved [2].

[1] Chum, Matas, Kittler: Locally Optimized RANSAC, DAGM 2003

[2] Ivashechkin, Barath, Matas: VSAC: Efficient and Accurate Estimator for H and F, ICCV 2021

Spatial Coherence of Geometric Data

Exploiting the Spatial Coherence of Geometric Data

Motivation: In vision, we usually have geometric data, e.g., 3D points, where the points often originate from spatially coherent structures.

Two-view geometry.

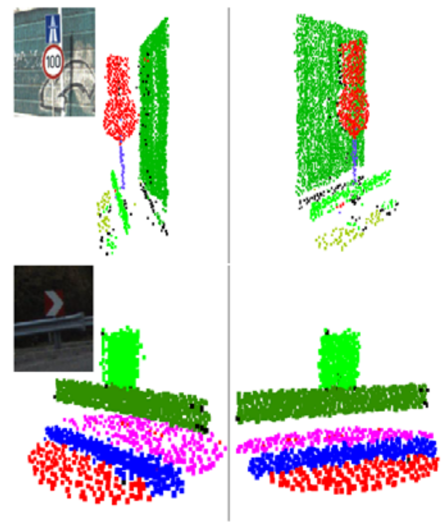

Planes in LiDAR data.

Vanishing points

(*Left*) Rigid motions in two views.
(*Right*) 1st images of image pairs with the inliers of homographies.

Exploiting the Spatial Coherence of Geometric Data

Approaches

- Exploit spatial coherence in the *local optimization* step to find local structures accurately (Graph-Cut RANSAC).
- Find spatial structures early by localized *sampling* (Progressive NAPSAC).

Graph-Cut RANSAC

Idea:

- Consider, in the local optimization of LO-RANSAC, that geometric data often form spatially coherent structures.
- The spatial coherence is used when selecting the inliers of a model.

Minimal sample initializing a rigid motion.

Inliers by standard RANSAC-like thresholding.

Inliers by considering spatial coherence.

Barath and Matas, Graph-Cut RANSAC; CVPR 2018

Robust fitting as a Labeling Problem

Idea: formalize robust fitting as a binary labeling problem.

Given

- a set of data points $\mathcal{P}$,
- an inlier-outlier threshold ϵ, and
- the parameters of a model θ

Objective: find labeling $\mathcal{L}$ where

- point $p_i \in \mathcal{P}$ with residual $r \leq \epsilon$ is labeled *inlier*.
- point $p_o \in \mathcal{P}$ with residual $r > \epsilon$ is labeled *outlier*.

Robust fitting as a Labeling Problem

The problem: find labeling $\mathcal{L}^* = \arg_{\mathcal{L}} \min E_{\{0;1\}}(\mathcal{L})$

$$\mathcal{L} \in \{0, 1\}^n$$

Inlier label Outlier label

Unary energy

$$E_{\{0;1\}}(\mathcal{L}) = \sum_{p \in \mathcal{P}} \left| \mathcal{L}_p \right|_{\{0;1\}}$$

Energy from each data point.

Energy from point p.

$$\left| \mathcal{L}_p \right|_{\{0;1\}} = \begin{cases} 0 & \text{if} \left(\mathcal{L}_p = 0 \wedge r_p \leq \epsilon \right) \\ & \vee \left(\mathcal{L}_p = 1 \wedge r_p > \epsilon \right) \\ 1 & \text{otherwise.} \end{cases}$$

„Close" and labeled inlier

„Far" and labeled outlier

Notes:

- Problem $\mathcal{L}^* = \arg_{\mathcal{L}} \min E_{\{0;1\}}(\mathcal{L})$ can be easily solved in polynomial time by the standard min-cut/max-flow, i.e. graph-cut, algorithm.
- Labeling $\mathcal{L}^*$ is exactly what RANSAC does, so this is „just" a reformulation of the RANSAC problem, but needs an initial estimate of model parameter θ
- Advantage: new energy terms can be added.
- Using a different loss function than RANSAC's is easy and beneficial.

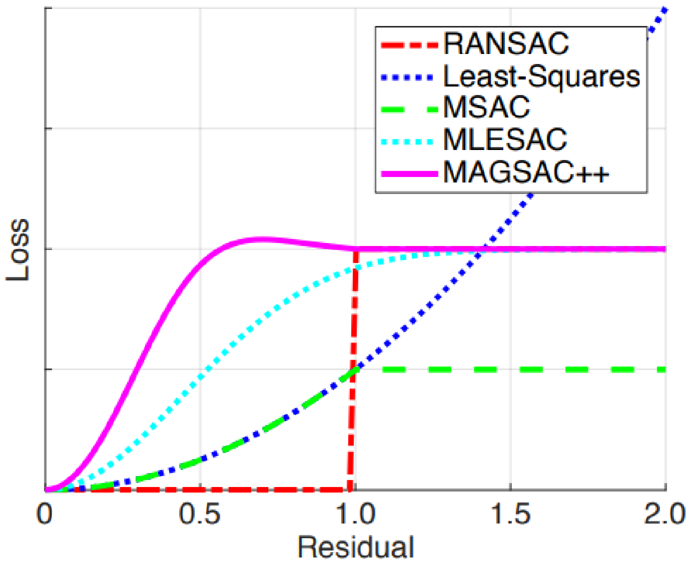

Spatial Coherence in Robust fitting (Potts model)

Modelling the spatial coherence by the Potts model.

Binary energy

$$E_{\text{Potts}}(\mathcal{L}) = \sum_{(p,q)\in\mathcal{E}} \begin{cases} 1 & \text{if } \mathcal{L}_p \neq \mathcal{L}_q \\ 0 & \text{otherwise} \end{cases}$$

Edge set of $\mathcal{G}$

Neighborhood graph $\mathcal{G}$

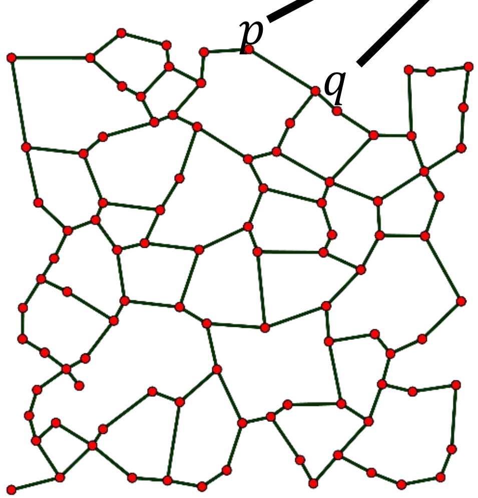

In short, if neighboring points

- have different labels, penalize.
- have similar labels, do nothing.

R. Zabih and V. Kolmogorov. ***Spatially coherent clustering using graph cuts***. *CVPR 2004.*

Spatial Coherence in Robust fitting (Potts model)

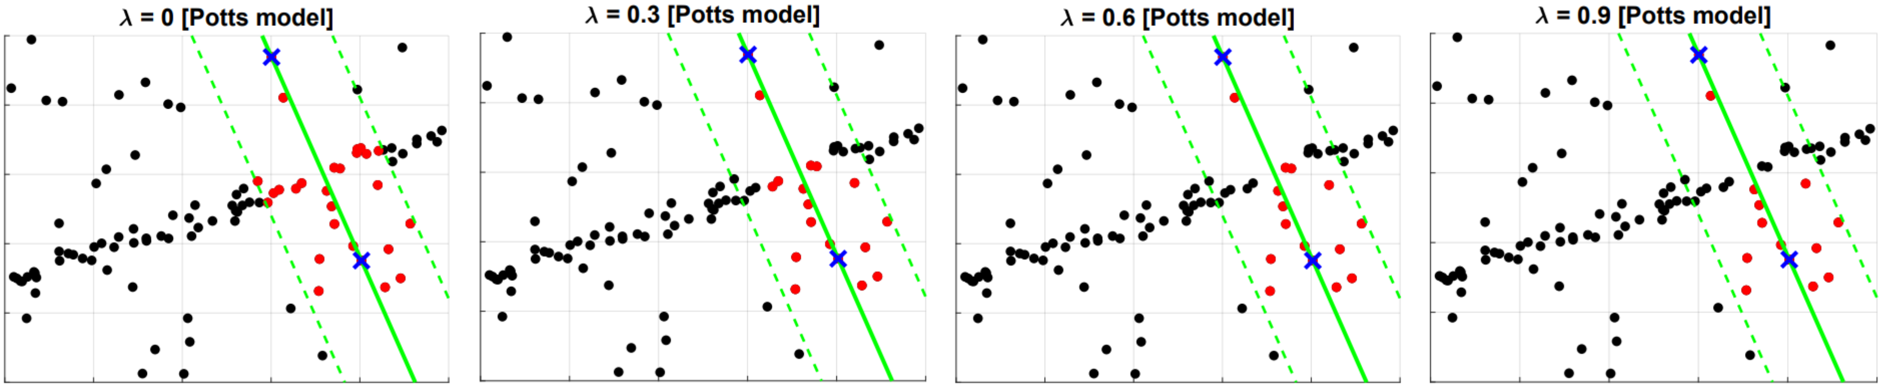

Example labeling using the Potts model to interpret spatial coherence. Parameter $\lambda \in [0, 1]$ is the weight of the term.

Issue: Outliers are considered similarly structured as the inliers. Thus, the Pots model forces all points in the structure to be outliers even if they are „close" to the line.

R. Zabih and V. Kolmogorov. ***Spatially coherent clustering using graph cuts****. CVPR 2004.*

Spatial Coherence in RANSAC (GC model)

Modelling the spatial coherence by the GC model.

$$E_{\mathrm{GC}}(\mathcal{L}) = \sum_{(p,q)\in\mathcal{E}} \begin{cases} 1 & \text{if } \mathcal{L}_p \neq \mathcal{L}_q, \\ 0 & \text{if } \mathcal{L}_p = \mathcal{L}_q = 0, \\ 1 - \dfrac{f(0,p) + f(0,q)}{2} & \text{if } \mathcal{L}_p = \mathcal{L}_q = 1. \end{cases}$$

Function f is the robust loss, e.g., that of RANSAC or MSAC

Example robust loss

$$f_{\mathrm{MSAC}}(\mathcal{L}_p, p) = \begin{cases} r_p^2/\epsilon & if\ (\mathcal{L}_p = 0 \wedge r \leq \epsilon) \\ 0 & if\ (\mathcal{L}_p = 1 \wedge r > \epsilon) \\ 1 & \text{otherwise.} \end{cases}$$

In short, if neighboring points

- have different labels, penalize.
- are labeled inliers, do nothing.
- are labeled outliers labels, penalize depending on their loss.

Spatial Coherence in RANSAC (GC model)

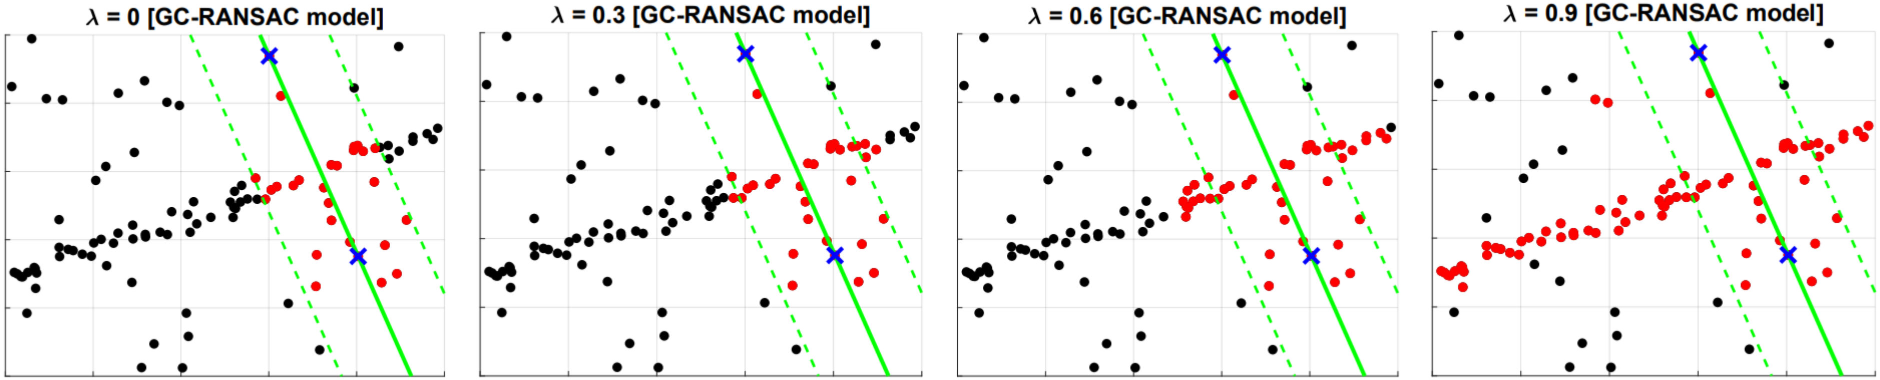

Example labeling using the GC model to interpret spatial coherence.
Parameter $\lambda \in [0, 1]$ is the weight of the term.

Notes:

- „Close" points are inliers no matter the spatial coherence.
- The inlier label is spread along the spatial structure.

GC-RANSAC Local Optimization

Algorithm 2 GC-RANSAC Local Optimization.

Input: $\mathcal{P}$ – data points, L^* – labeling, q^* – quality, θ^* – model;

Output: L^*_{GC} – labeling, w^*_{GC} – support, θ^*_{GC} – model;

$q^*_{\text{GC}}, L^*_{\text{GC}}, \theta^*_{\text{GC}} \leftarrow q^*, L^*, \theta^*$.

$\texttt{terminate}, \leftarrow \texttt{false}$.

while $\neg$ $\texttt{terminate}$ **do**

  $G \leftarrow \texttt{ConstructGraph}(\mathcal{P}, \mathcal{A}, \theta^*_{\text{GC}}, \lambda)$.

  $L \leftarrow \texttt{GraphCut}(G)$

  $\theta, \widehat{L}, q \leftarrow \texttt{RANSAC}(L)$.

  if $q > q^*_{\text{GC}}$ **then**

    $q^*_{\text{GC}}, \theta^*_{\text{GC}}, L^*_{\text{GC}} \leftarrow q, \theta, \widehat{L}$.

  else

    $\texttt{terminate} \leftarrow \texttt{true}$.

  $i \leftarrow i + 1$.

  if $\texttt{ShouldTerminate}(\theta^*_{\text{GC}}, L^*_{\text{GC}}, i)$ **then**

    $\texttt{terminate} \leftarrow \texttt{true}$.

Selecting the inliers by graph-cut

Inner RANSAC on the spatially coherent inliers

Update if new so-far-the-best

GC RANSAC - Speed

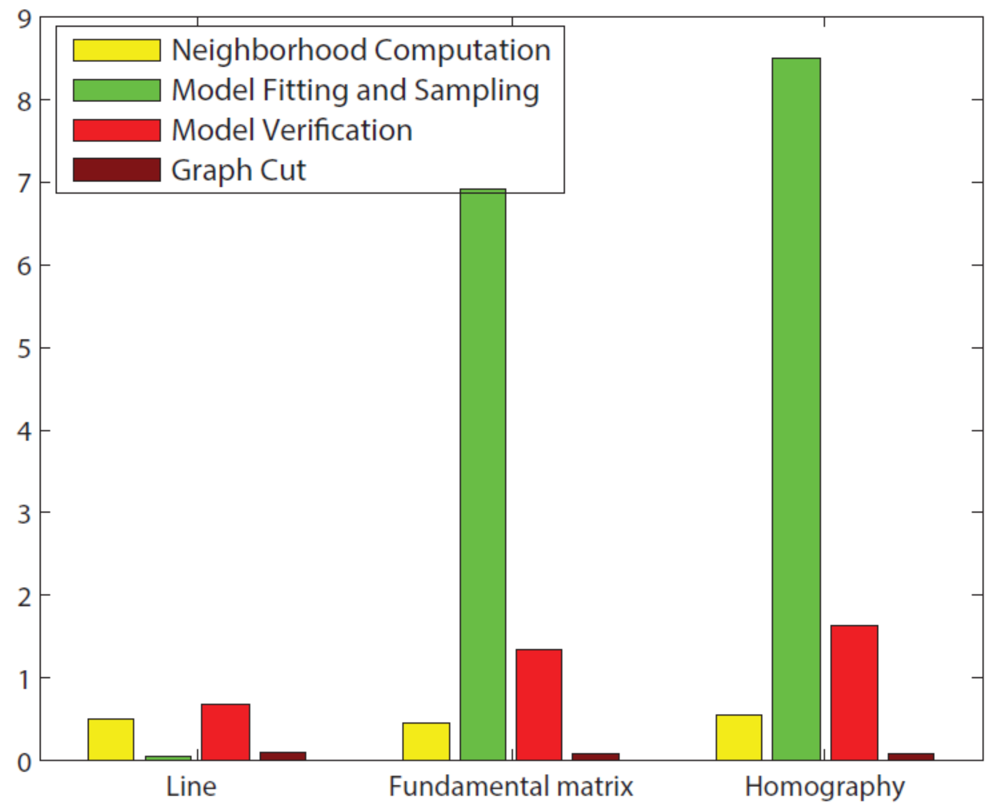

Figure 7: The breakdown of the processing times in milliseconds. Computed as the mean of all tests. *Best viewed in color.*

NAPSAC Sampler

Nasuto et al. Napsac: High noise, high dimensional robust estimation-it's in the bag., BMVC 2002

Idea:

- Points close to an inlier are likely to be inliers.
- Selecting the minimal sample from a hypersphere leads likely to „all-inlier" samples.

Algorithm:

- First point is selected at random.
- The rest of the sample is from the hypersphere, of fixed radius, around the first one.

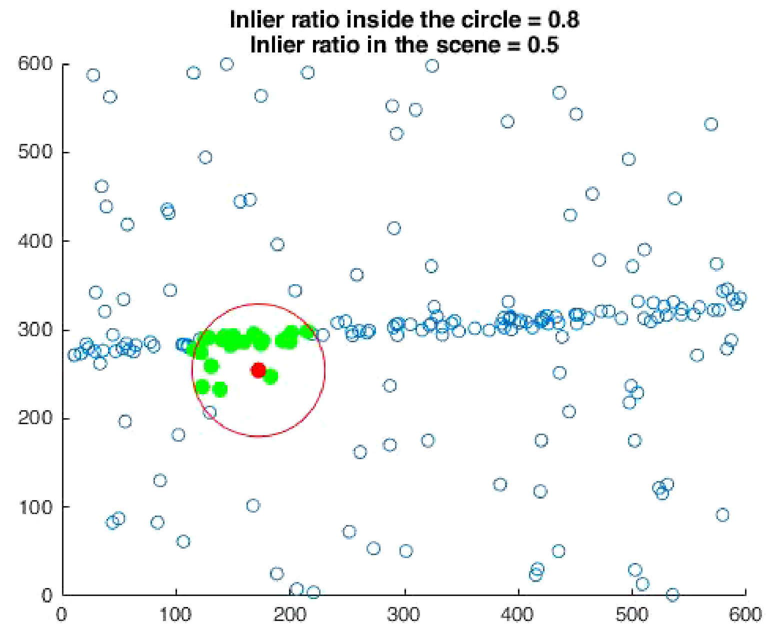

The 1st point (red dot) is selected at random. The points (green) inside the assigned circle (red) are shown.

NAPSAC Sampler

Advantage: high inlier ratio for local samples

Issues:

- Localized samples are often imprecise.
- For some models, local samples are likely degener
  E.g., for fundamental matrix estimation, the poi
  should originate from more planes.
- If the model is not localized,
  it is never found.
  The method is sensitive
  to the radius.

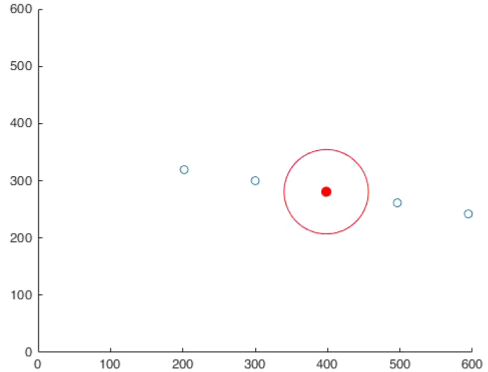

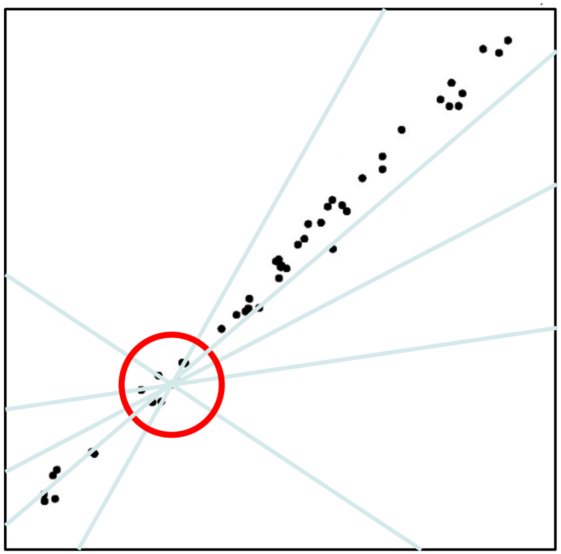

Example line fitting to NAPSAC samples. A red circle is centered on the 1st point. All lines fit using a 2nd point from the circle are inaccurate.

PROSAC – PROgressive SAmple Consensus

O. Chum , J. Matas: Matching with PROSAC - Progressive Sample Consensus, CVPR 2005

- Not all datapoints are created equal, some are better than other, e.g.

  D. Lowe's SIFT distance ratio,
  E. Brachmann, C. Rother: Neural-Guided RANSAC: Learning Where to Sample Model Hypotheses. ICCV 2019

- Sample from the best candidates first

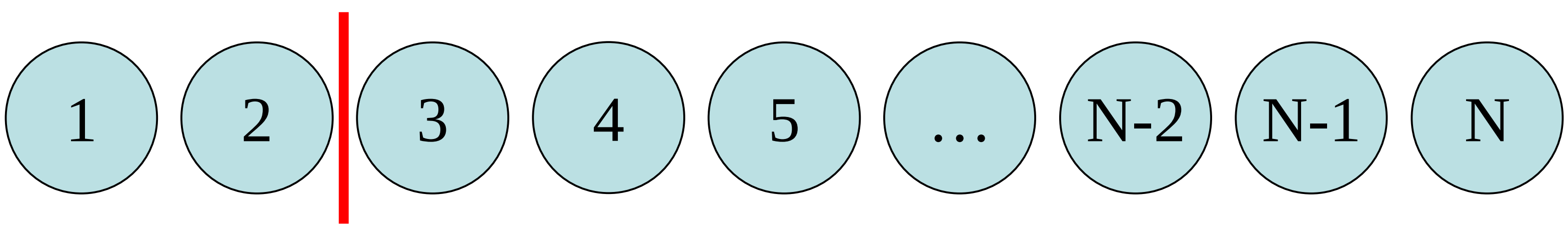

Sample from here

PROSAC Samples

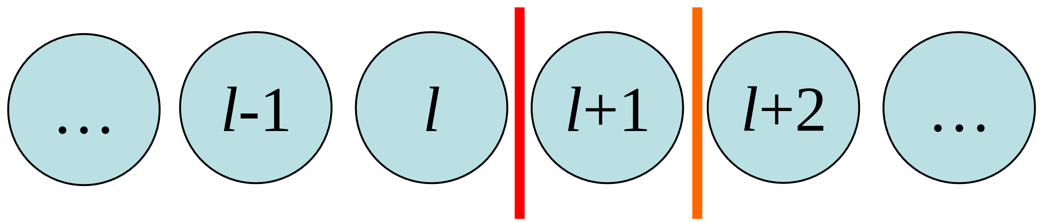

Draw T_l samples from $(1 \ldots l)$
Draw T_{l+1} samples from $(1 \ldots l+1)$

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 $l+1$

Epipolar geometry estimation

Executed on all TC

Background $N = 783, \varepsilon = 79\%$

	I	k	time [sec]
PROSAC	617	1.0	0.33
RANSAC	617	15	1.10

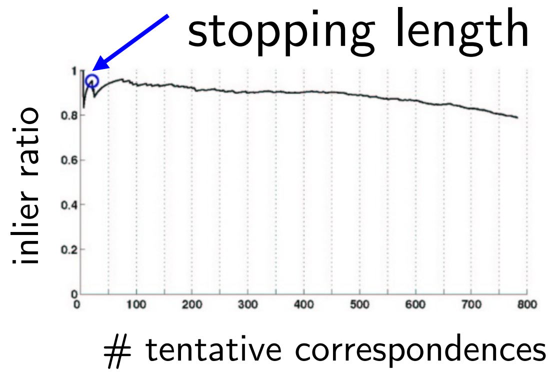

Executed on outliers to the background model

Mug $N = 166, \varepsilon = 31\%$

	I	k	time [sec]
PROSAC	51.6	18	0.12
RANSAC	52.3	10,551	0.96

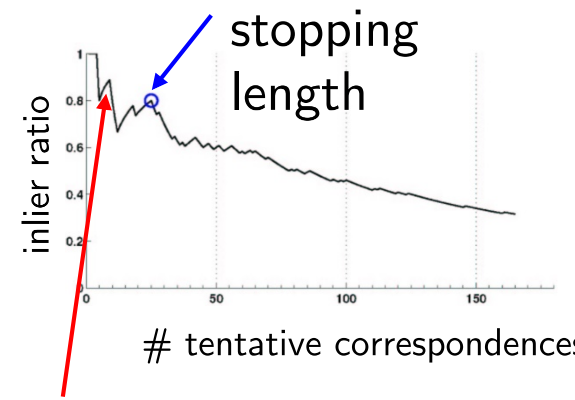

Too small set of TC – could be random

Progressive NAPSAC

A sampler for RANSAC-like robust estimators.

Idea:

- Start sampling locally and progressively blend to global sampling.
- A hypersphere is assigned to each point. Its radius is increased gradually.
- The 1st point $\boldsymbol{p}$ is selected by PROSAC.
- The rest of the sample is from the hypersphere around $\boldsymbol{p}$.

Barath et al., MAGSAC++, a fast, reliable and accurate robust estimator, CVPR 2020

Progressive NAPSAC

	Pro	Con
Small circle	+ higher inlier ratio	- poorly conditioned
Big circle	+ well conditioned	- lower inlier ratio

Comparison of neighborhood sizes.

Example. A selected point (red dot); the assigned neighborhood (black circle); the inliers of the sought 2D line (green dots) and outliers (blue crosses).

Growth function

1. Given the first selected point $\boldsymbol{p}_i$.

2. Let $\{\mathcal{M}_{i,j}\}_{j=1}^{T(i)} = \{\boldsymbol{p}_i, \boldsymbol{p}_{x_{i,j,1}}, \boldsymbol{p}_{x_{i,j,2}}, \dots \boldsymbol{p}_{x_{i,j,m-1}}\}_{j=1}^{T(i)}$ be a sequence of samples containing point $\boldsymbol{p}_i$ where indices $x_{i,j,1}, \dots, x_{i,j,m-1}$ refers to points and m is the sample size.

3. For all indices, the points are ordered w.r.t. to their distance to $\boldsymbol{p}_i$.
   Thus, if $\mathrm{k} \leq \mathrm{l}$, $|p_k - p_i| \leq |p_l - p_i|$.

4. Given a sphere containing the k closest neighbors $\mathcal{S}_{i,k}$ of point $\boldsymbol{p}_i$.

5. The number of samples containing points from $\mathcal{S}_{i,k}$ and $\boldsymbol{p}_i$ is $T_k(i)$.

6. Expected number of $T_k(i)$ is $\mathrm{E}\big(T_k(i) \mid T(i)\big) = T(i)\binom{k}{m-1} / \binom{n-1}{m-1}$.

7. After $T_k(i)$ iterations, the sphere radius is increased to contain the closest $k+1$ neighbors.

Noise Level
Inlier – Outlier Threshold

Robust Model Estimation (Standard Setting)

Input:

- A set of data points.
- *Inlier-outlier threshold.*

Output:

- Model parameters, e.g., 2D line.
- Set of inliers.

Note: the threshold should correspond to the noise level to find the sought inliers.

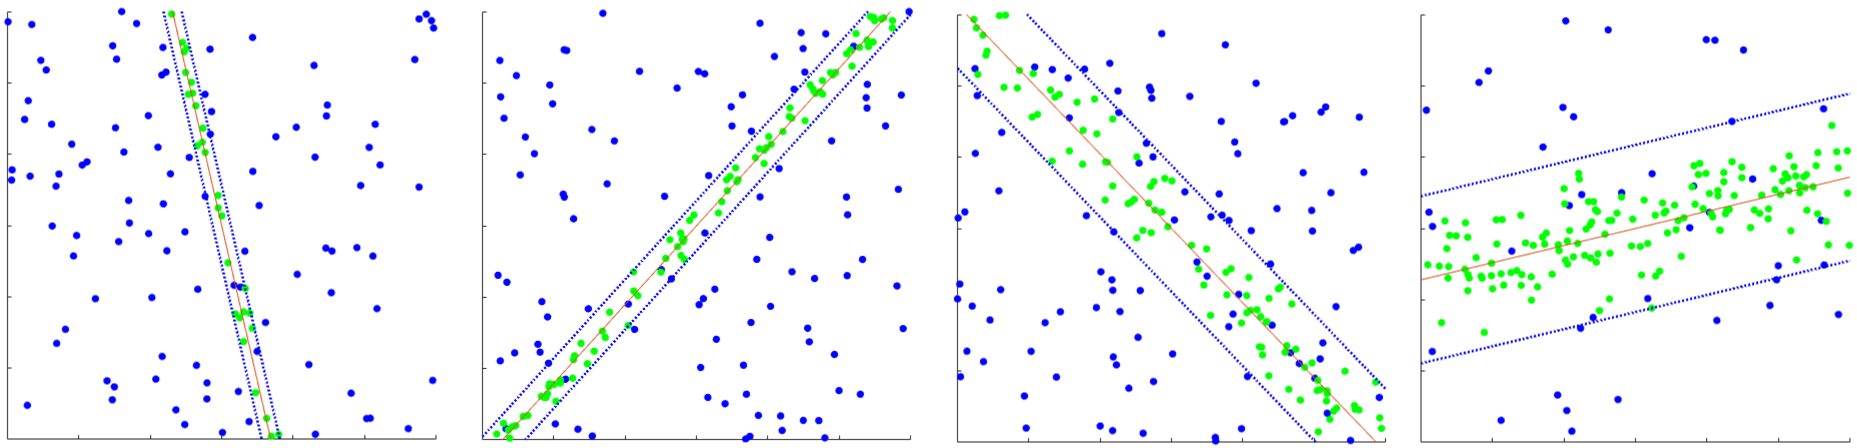

Randomly generated points (green) on a 2D line (red) and outliers (blue dots). The threshold which the synthetic noise added to the inlier point coordinates' implies is shown by blue dotted lines.

The issue: no single inlier-outlier threshold suits all problems.

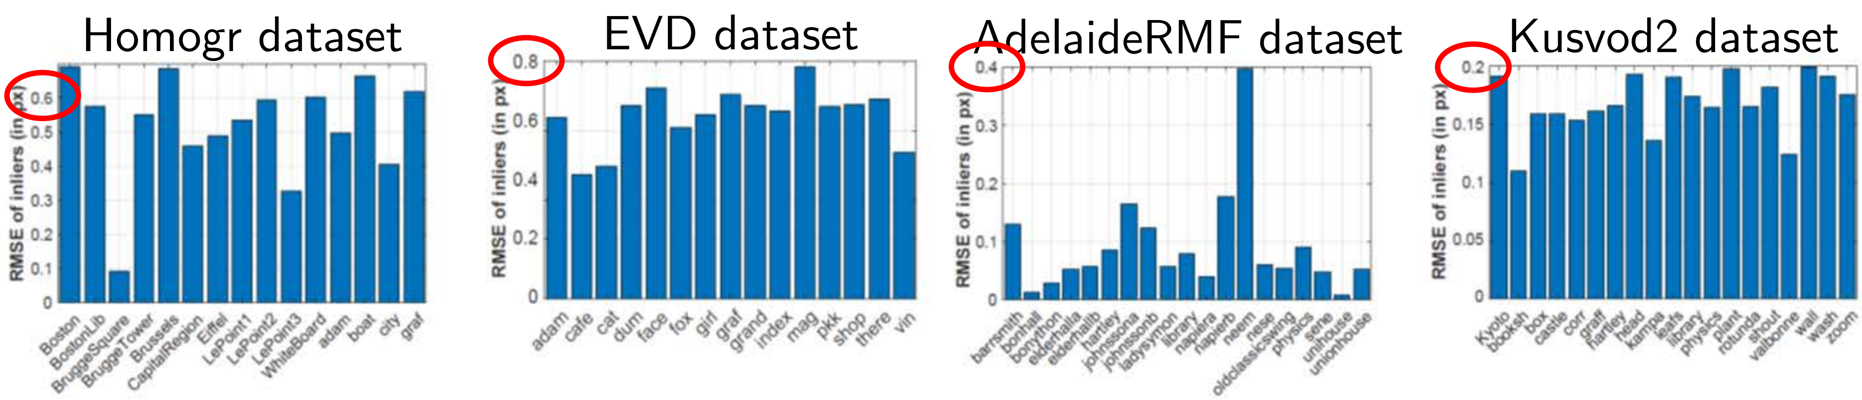

Model inaccuracy (RMSE) measured on carefully selected inlier correspondences in image pairs (bars) from standard datasets.

Note: the Y-ranges of the plots are different.

Example Sources of Model Inaccuracy

- **Noise in the measured data.**
  E.g., noise in the coordinates of points due to compression artifacts or resizing....

- **Imperfection of the model.**
  We are estimating, e.g., standard homography or fundamental matrix.

This happens if:

(i) there is rolling shutter,

(ii) there is radial distortion,

(iii) the scene is not completely rigid,

Example of rolling shutter effects

Example of radial distortion.

MINPRAN

Stewart, Charles V. MINPRAN: A new robust estimator for computer vision. TPAMI 1995

Idea:

- Assume that the outliers are uniformly distributed within the sensor range (e.g., image size).
- Find the model and threshold where the points closer than the threshold (i.e., inliers) are the least likely to have occurred uniformly randomly.

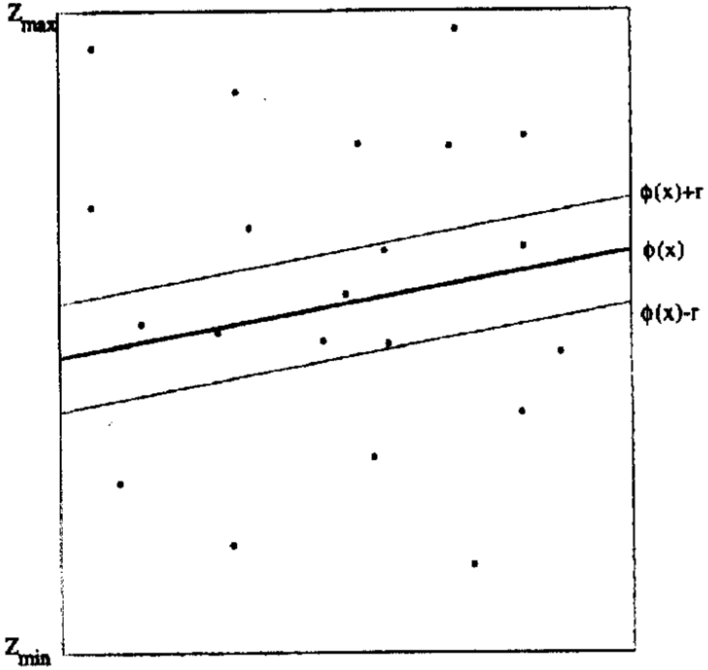

Problem:

- Given model ϕ, n points and probability $\mathcal{F}(r,k,n)$ that at least k points falls closer than threshold r.
- The problem to solve is $\mathcal{H}(\phi) = \arg_r \min \mathcal{F}(r,k,n)$ to select r, minimizing the randomness criterion.
- The final model will be where $\mathcal{H}(\phi)$ is minimal.

The MAGSAC Approach

D. Barath, J. Noskova, J. Matas, MAGSAC: Marginalizing Sample Consensus, CVPR 2019

Idea: Eliminate the threshold by marginalizing over it.

- The model quality does not depend on a single threshold,
- and the final model parameters are obtained *without* a strict inlier/outlier decision.

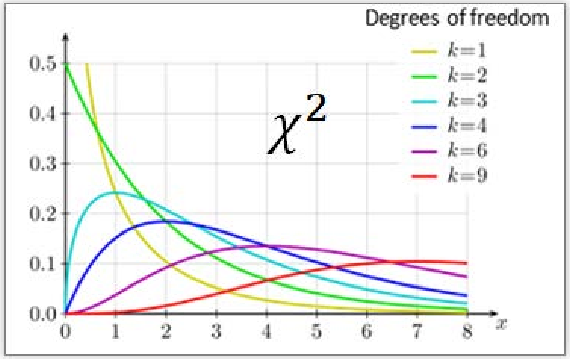

Design choices and data interpretation:

- Outlier are uniformly distributed.
- The squared inlier residuals have the χ^2 distribution.

→ Holds for residuals that are them sum of squared values with Gaussian distribution, e.g., re-projection error.

Notes:
- MAGSAC is capable of assuming other distributions.
- The χ^2 distribution should be used even in vanilla RANSAC.

Model Quality and its Implied „Problems"

New quality function:

$$Q^*(\theta) = \frac{1}{\sigma_{max}} \int_0^{\sigma_{max}} \ln L(\theta \mid \sigma)\, d\sigma$$

- Function $Q^*(\theta)$ does not depend on the noise scale σ and, thus, on a threshold.
- No inlier/outlier decision is made.

Implied „Problems":

- No knowledge about the inliers $\rightarrow$ Final least-squares fitting is not applicable.
- No knowledge about the inlier number $\rightarrow$ We don't know when to stop.

Note: Parameter σ_{max} is a loose upper bound for the noise scale.

The σ-consensus Model Fitting

Problem: the inliers are not known to use them in an LSQ model polishing step.

A solution:

1. Given an initial model θ, e.g., from a minimal sample.

2. Calculate the inlier probabilities of each point x as

$$L(x \mid \theta) \approx \frac{2C(x)}{\sigma_{max}} \sum_{i=1}^{K} (\sigma_i - \sigma_{i-1}) \sigma_i^{-p} D^{p-1}(\theta_{\sigma_i}, x) \exp \frac{-D^2(\theta_{\sigma_i}, x)}{2\sigma_i^2}.$$

3. Apply weighted least-squares fitting using the probabilities as weights.

Note: The (piecewise constant) integral in $L(x \mid \theta)$ is replaced by a weighted sum.

Termination Criterion

Problem: the inlier number is not known, so cannot be used to terminate.

Original RANSAC criterion:

Manually set confidence

Model parameters   Point set

RANSAC iteration number

$$k(\theta, \ \sigma, \ \mathcal{X}) = \frac{\ln(1-\mu)}{\ln\left(1 - \left(\frac{|I(\theta,\sigma)|}{|\mathcal{X}|}\right)^m\right)}$$

Sample size

Noise scale

Inlier ratio

Criterion via marginalization:

MAGSAC iteration number

$$k^*(\theta, \mathcal{X}) = \frac{1}{\sigma_{max}} \int_0^{\sigma_{max}} \underbrace{k(\theta, \sigma, \mathcal{X})}_{\text{RANSAC iteration number}} \, d\sigma$$

Noise scale upper bound

The MAGSAC Algorithm

Input:

- A set of data points.
- An *upper bound for the threshold*.

Output: model parameters.

Algorithm:

1. Generate a minimal sample model and calculate its quality.
2. If it is a new best model: Estimate the inlier weight of each point using a (data-dependent) number of different noise scales.
3. Estimate the final model parameters by weighted least-squares.
4. Terminate or go to 1.

Example.
The minimal sample (red dots), the line which it initializes (red line), the current threshold (blue dotted lines),
the implied inliers (green dots), and the model fit to the implied inliers (red dashed line).

Pros and Cons

Pros	Cons
Very accurate on the tests	Due to the number of LS fits, MAGSAC can be slow on some problems.
Smaller failure ratio than the competitors	
The setting of σ_{max} is much easier than setting σ.	

The MAGSAC++ Approach

D. Barath, J. Noskova, M. Ivashechkin, J. Matas, MAGSAC++, a fast, reliable and accurate robust estimator, CVPR 2020

Idea:

Eliminate the threshold similarly as in MAGSAC, but

- assume nothing about the outliers,
- and use an efficient iteratively re-weighted least-squares approach.

Design choices and data interpretation:

- No assumption on the outliers.
- The squared inlier residuals have the χ^2 distribution.

Note: the algorithms is capable of handling other distributions

Input:

- A set of data points.
- Upper bound for the threshold. E.g., 10 pixels for fundamental matrix fitting.

Output:

- Model parameters.

Point-to-model residual

$$\theta_{i+1} = \arg\min_{\theta} \sum_{p \in P} w(r(\theta_i, p)) r^2(\theta, p)$$

Model in the (*i*+1)th iteration

Model in the *i*th iteration

Algorithm:

Iteratively re-weighted LS fitting, where *model parameters in the (i+1)th iteration* are calculated from points *weighted* via marginalizing over the noise σ.

Prior on $\sigma \sim \mathcal{U}(0, \sigma_{max})$

$$w(r(\theta_i, p)) = \int_0^{\sigma_{max}} \mathrm{P}(p \mid \theta_i, \sigma)\, f(\sigma)\, d\sigma$$

Conditional probability of point *p* given model θ_i and σ

Multiple Instances

Multi-model fitting with RANSAC-like methods

Multi-model fitting problem:

- We are looking for a set of models (e.g., 3 planes) interpeting the scene,
- and a point-to-model assignment.

Note: there is an outlier model.

Connection of RANSAC-like methods and multi-model fitting:

- Early methods were using RANSAC directly.
- State-of-the-art methods:
  - Use RANSAC-like initialization and, then, some optimization to select the model interpeting the scene.
  - Use RANSAC inside the optimization.

Early Methods (Sequential RANSAC)

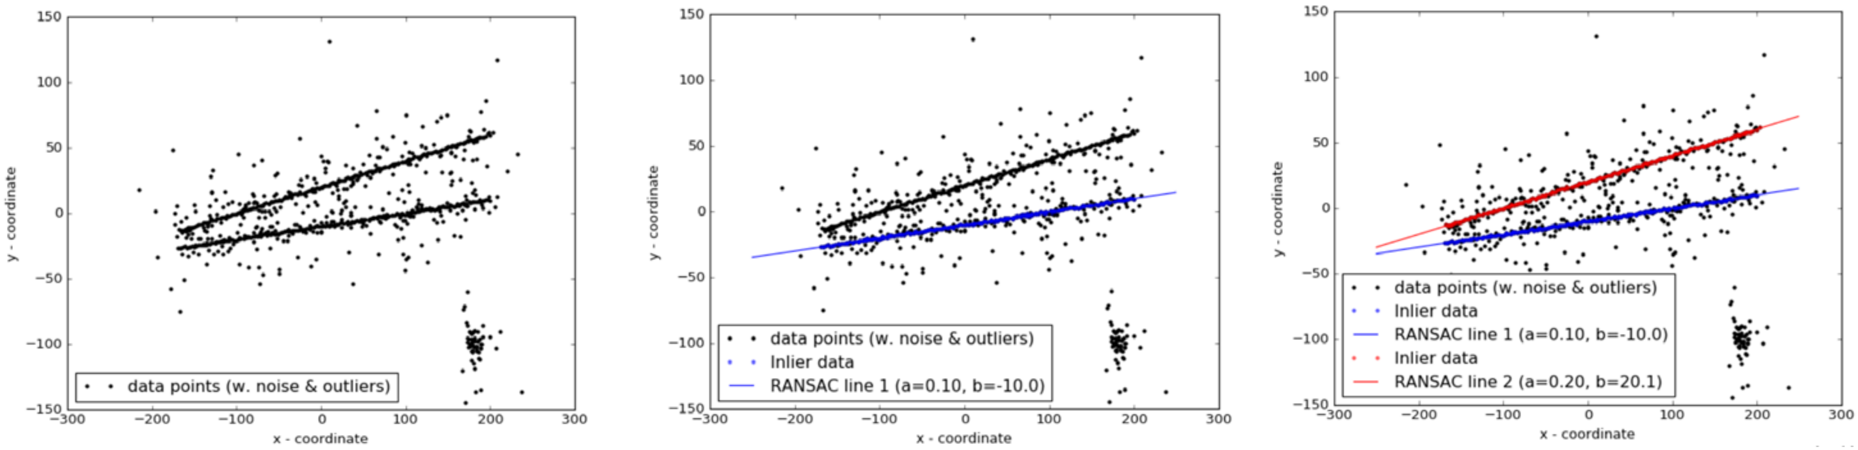

Example fitting by Sequential RANSAC.

Notes:

- Greedy algorithm, but works reasonably well for finding the most dominant models in the data.
- Very easy to implement.
- Scalable, in contrast to most state-of-the-art techniques.

Early Methods (MultiRANSAC)

- The number of sought models k is a parameter.
- In each RANSAC iteration, it selects k minimal samples and fits k models.
- Due to the increased sample size, it requires too many iterations.

Outlier ratio / Sample size	2	$3*2 = 6$
0.5	16	292
0.75	71	$1.8 * 10^4$
0.9	458	$4.6 * 10^6$

Theoretical number of iterations to achieve 0.99 confidence in the results.

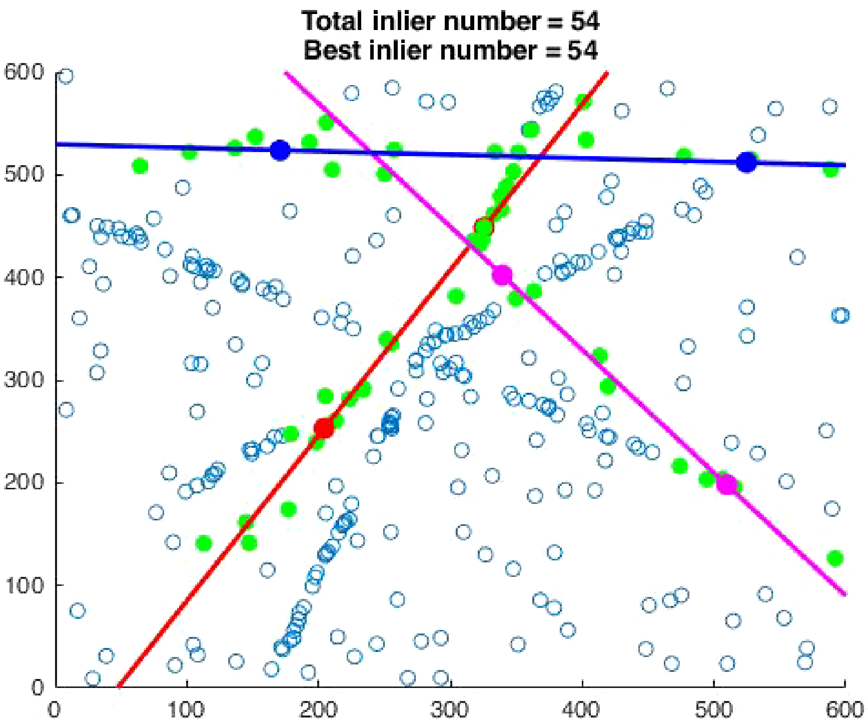

Example fitting by MultiRANSAC.

M. Zuliani et al. ***The multiransac algorithm and its application to detect planar homographies****. ICIP 2005.*

Recent Methods (RANSAC-like initialization)

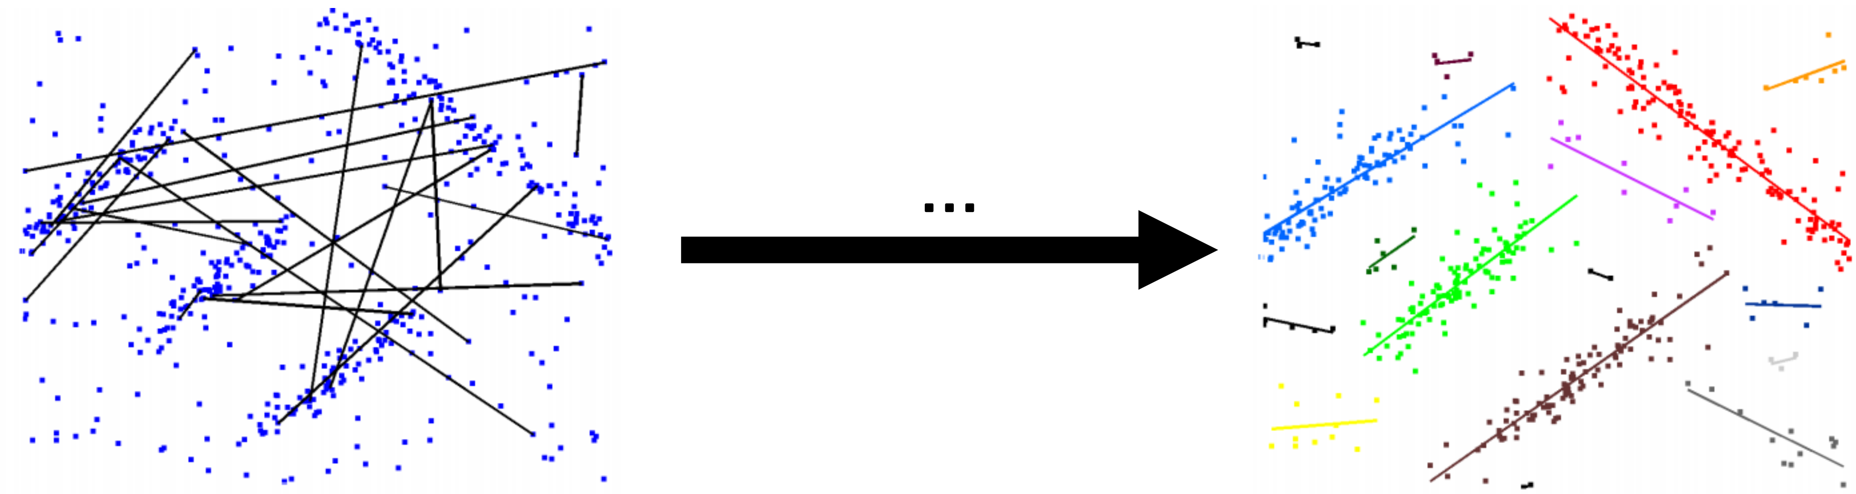

Pipeline:

- Generate an initial set of models by a RANSAC-like procedure.
- Do model selection – optimization - point assignment to find the dominant models.

P. Amayo et al. ***Geometric multi-model fitting with a convex relaxation algorithm****. CVPR 2018.*
D. Barath et al. ***Multi-class model fitting by energy minimization and mode-seeking****. ECCV 2018.*
A. Delong et al. ***Minimizing energies with hierarchical costs****. IJCV 2012*
L. Magri et al. T-Linkage: ***A continuous relaxation of J-Linkage for multi-model fitting****. CVPR 2014.*
L. Magri and A. Fusiello. ***Robust multiple model fitting with preference analysis and low-rank approximation****. BMVC 2015*
L. Magri and A. Fusiello. ***Multiple model fitting as a set coverage problem****. CVPR 2016*
T. T. Pham, T.-J. Chin, K. Schindler, and D. Suter. ***Interacting geometric priors for robust multi-model fitting****. TIP 2014*
H. Wang, G. Xiao, Y. Yan, and D. Suter. ***Mode-seeking on hypergraphs for robust geometric model fitting.*** *ICCV 2015.*
H. Wang, G. Xiao, Y. Yan, and D. Suter. ***Searching for representative modes on hypergraphs for robust geometric model fitting****. PAMI 2018.*
...

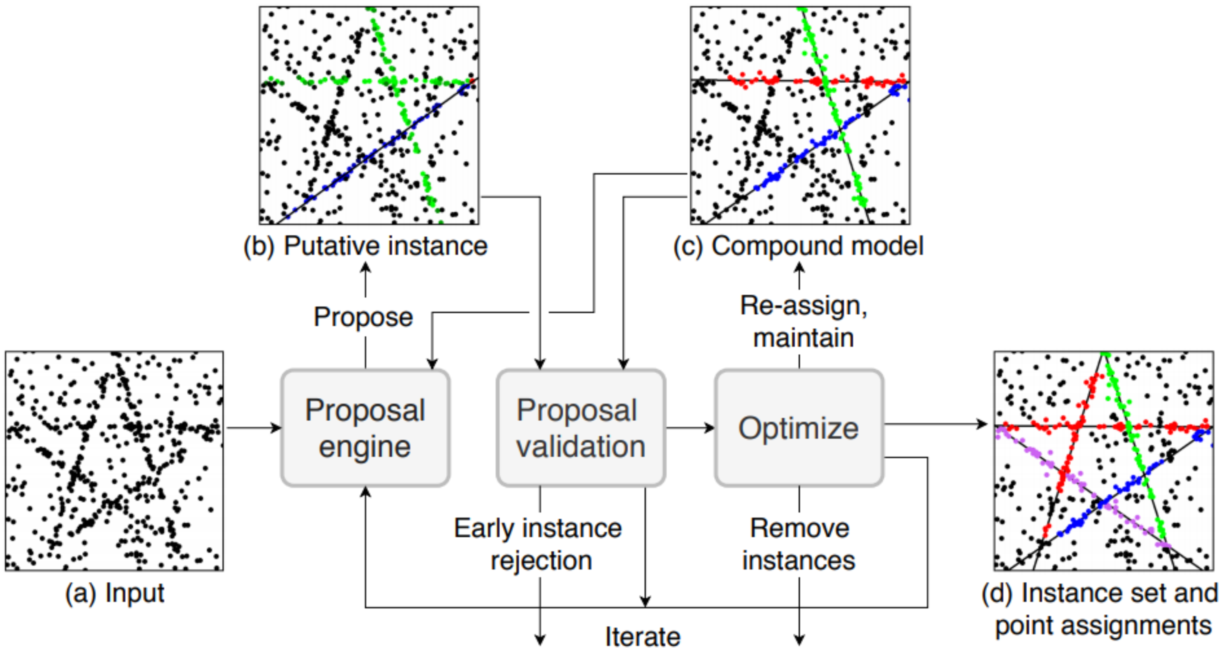

Pipeline of Progressive-X.

Note: also, there is CONSAC at CVPR 2020 following a similar strategy.

D. Barath et al. ***Progressive-X: Efficient, Anytime, Multi-Model Fitting Algorithm***. *ICCV 2019.*
F. Kluger et al. ***CONSAC: Robust Multi-Model Fitting by Conditional Sample Consensus***. *CVPR 2020*

Recent methods (RANSAC Inside the Optimization)

Example homography fitting by Progressive-X.

D. Barath et al. ***Progressive-X: Efficient, Anytime, Multi-Model Fitting Algorithm***. *ICCV 2019.*

Speed

RANSAC time complexity depends on the inner loop

Repeat k times (k depends on sample size m, inlier number 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

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

t_M - the time to estimate model parameters from a sample,
number of models generated from a sample
$\overline{m}_s$ - number of models tested per sample

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

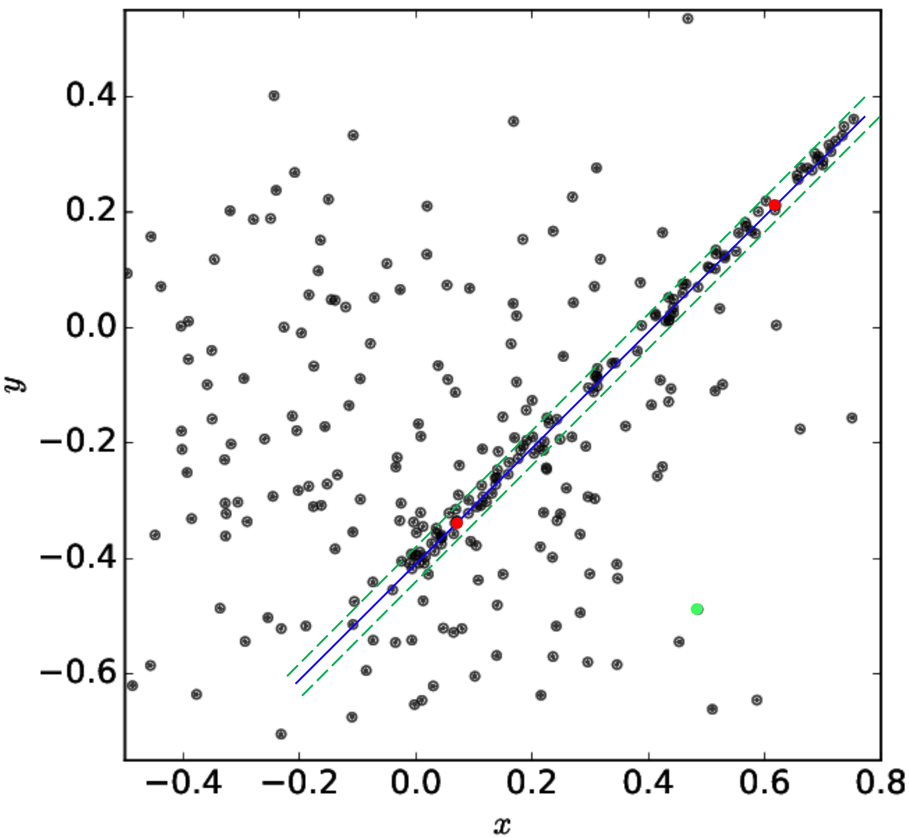

Example (d=1)

1. Generate a model (sample 2 points)

2a. Sample another point ●

Does it fall within threshold?

No. Go to 1.

Optimal Randomised Strategy

Model Verification employing Sequential Decision Making

H_g: $P(x_i = 1|H_g) \geq \varepsilon$
H_b: $P(x_i = 1|H_b) = \delta$
$x_i = 1$ $\quad x_i$ 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]

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*.

SPRT

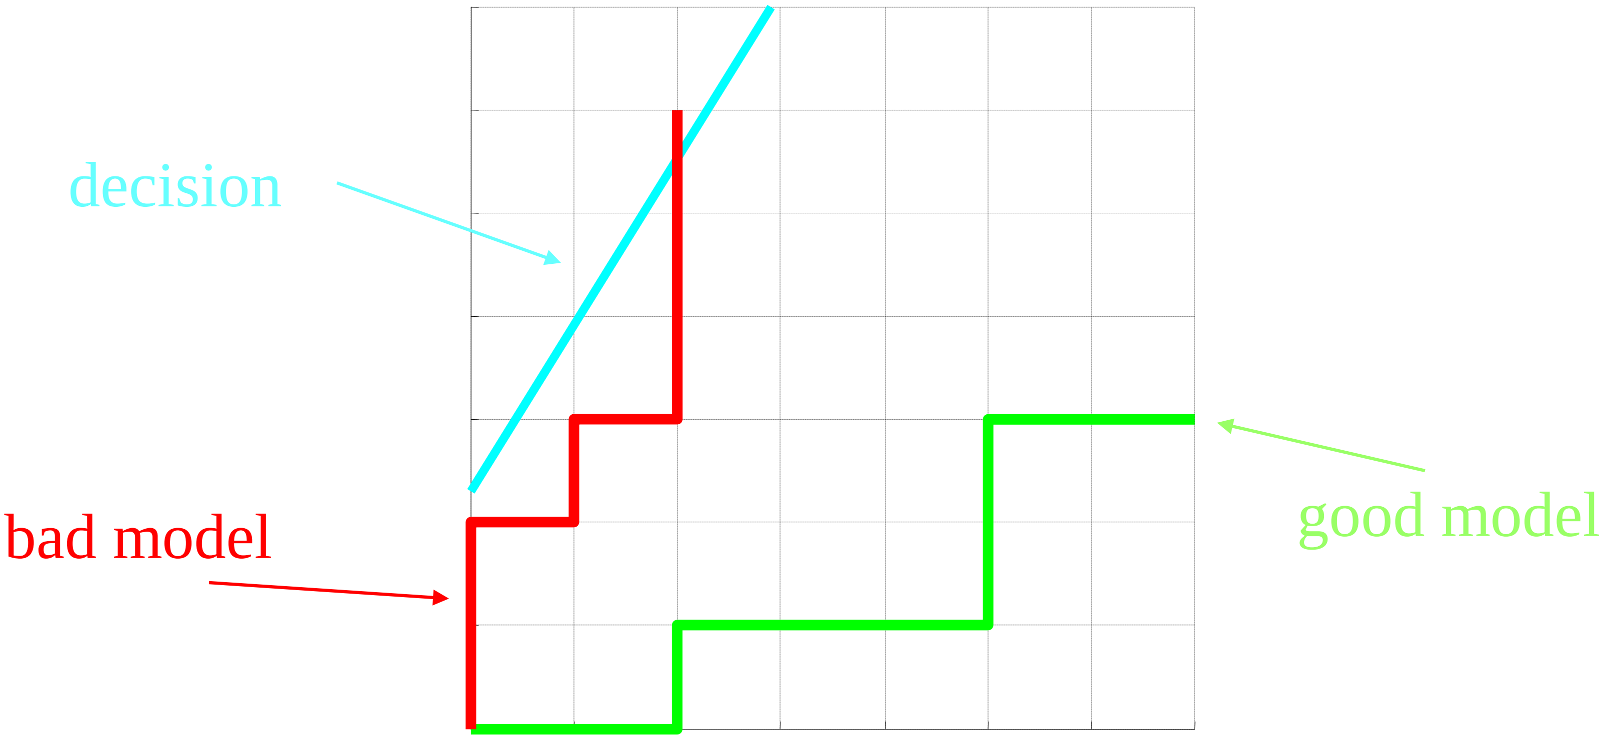

Note: the Wald's test is equivalent to series of $T(d, c)$, where $c = \lceil (\log A - d \log \lambda_1) / \log \lambda_0 \rceil$

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

Evaluation

Evaluation at CVPR 2020 : Classical F methods, 1k iterations

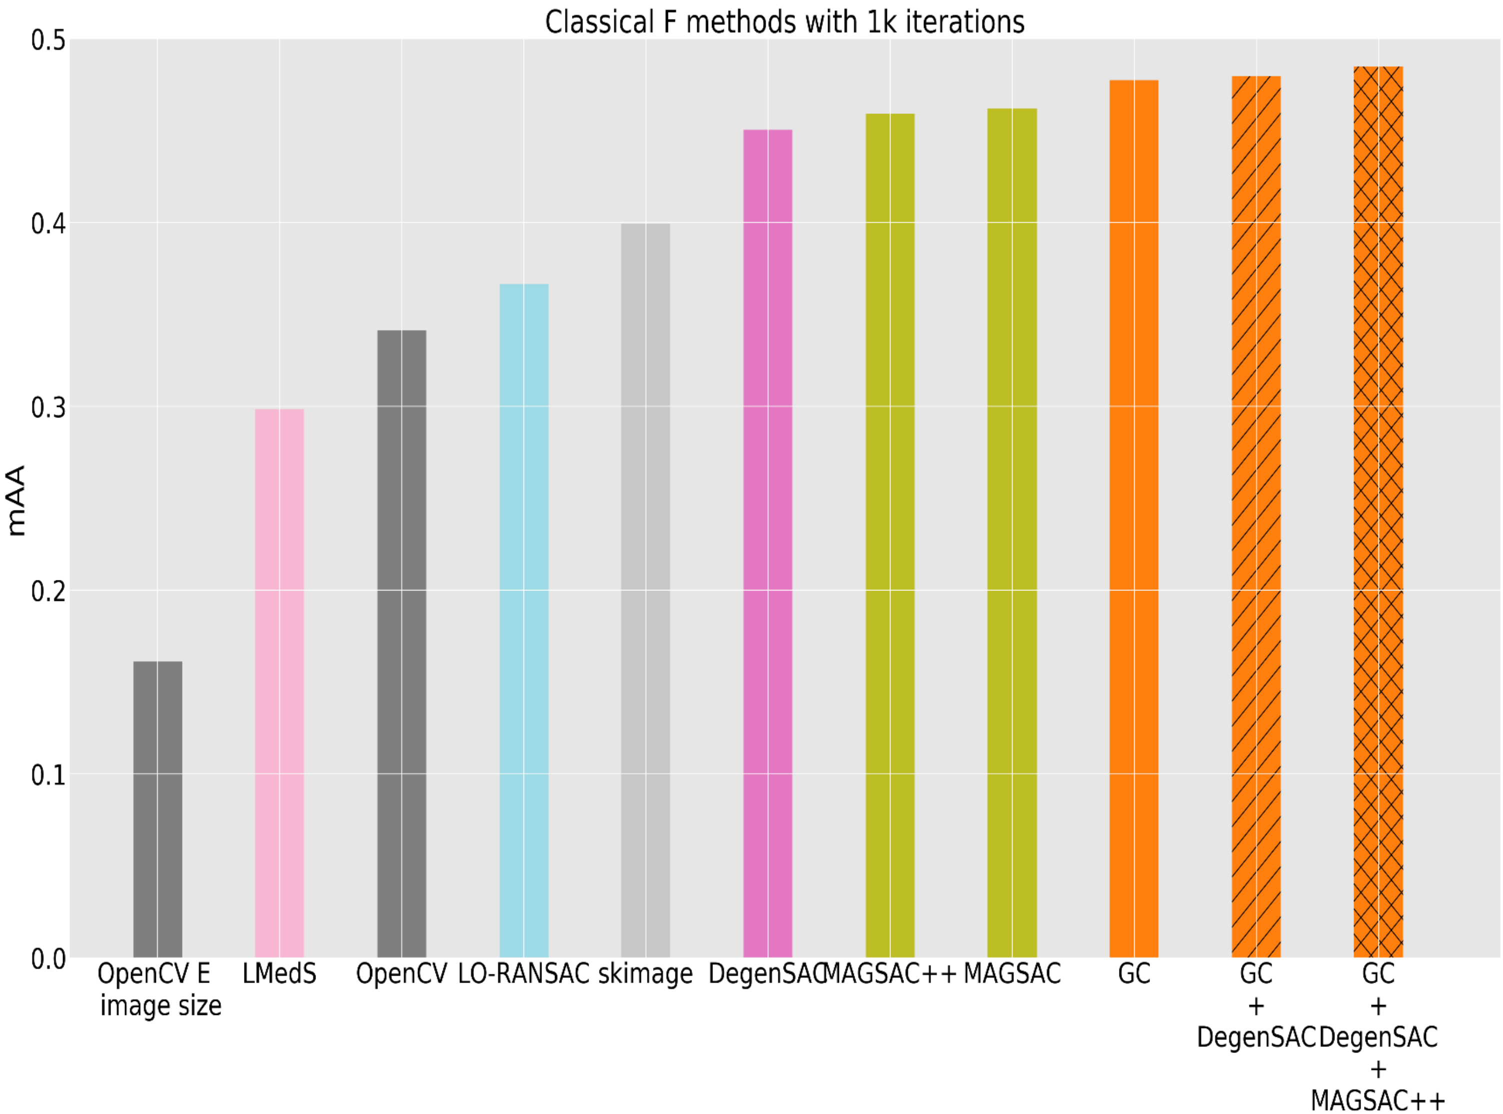

- Methods are sorted by accuracy
- sk-image RANSAC is orders of magnitude slower than the rest
- OpenCV is the least precise RANSAC

Evaluation at CVPR 2020: Learned methods F

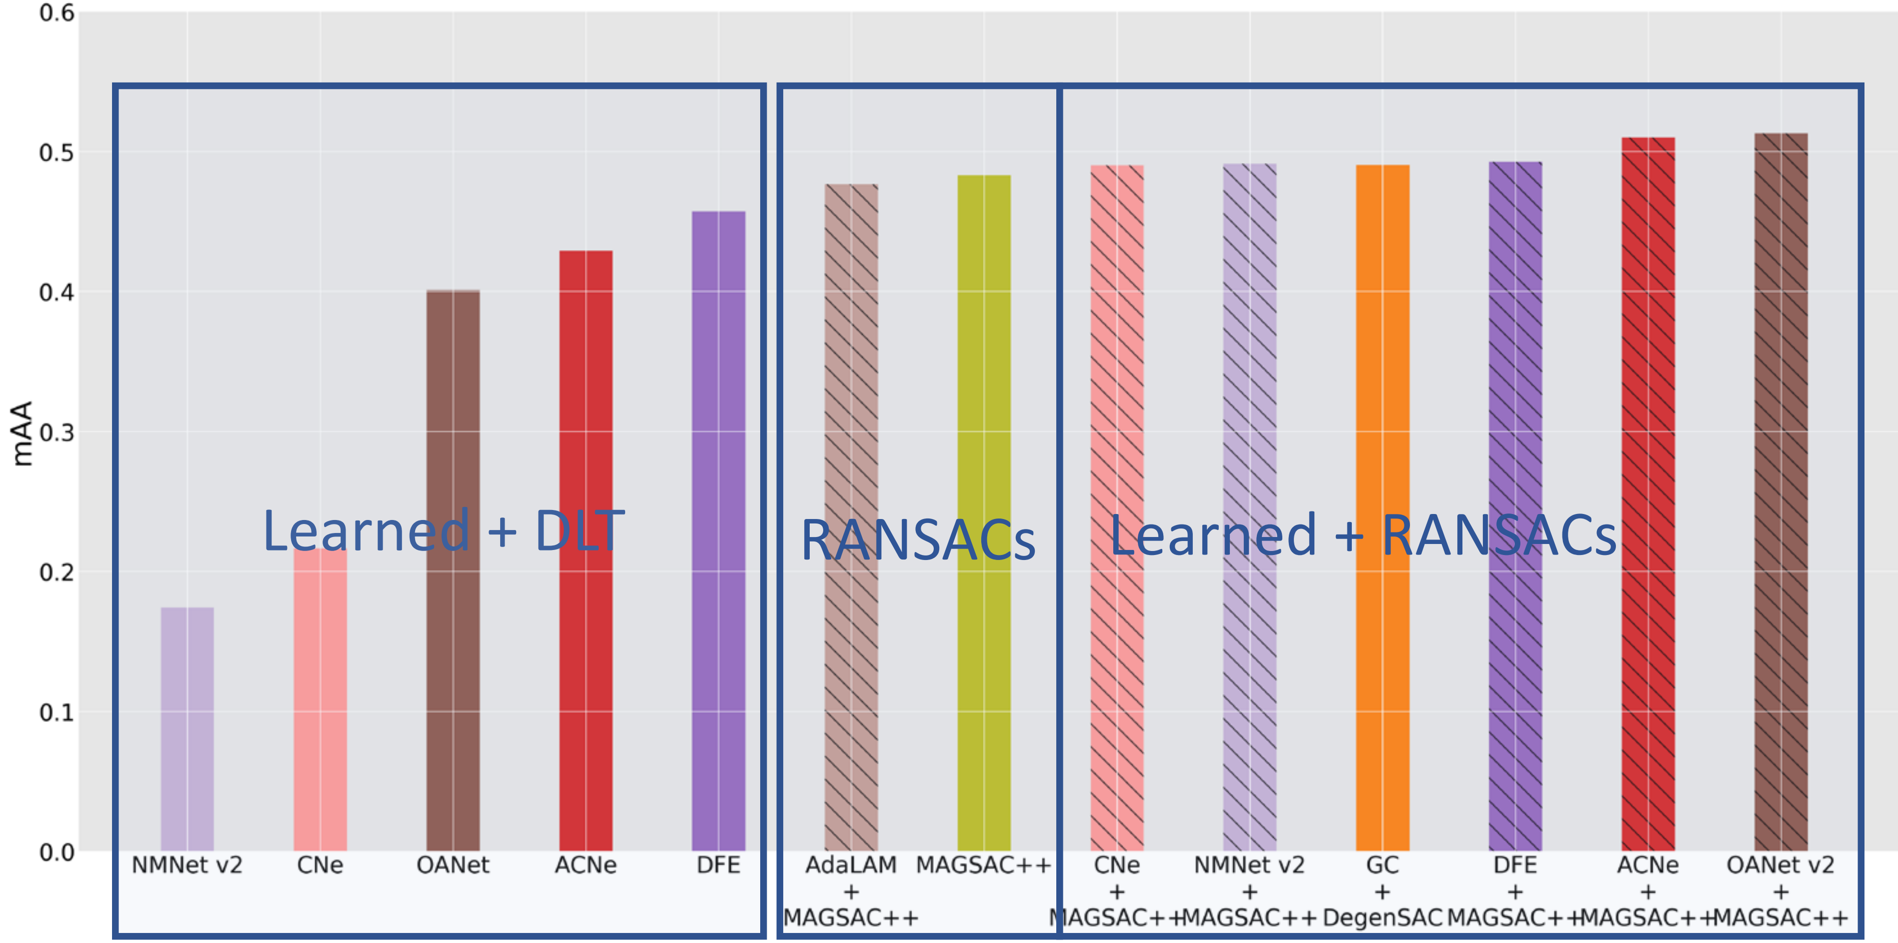

Classical methods, H, 10k iterations

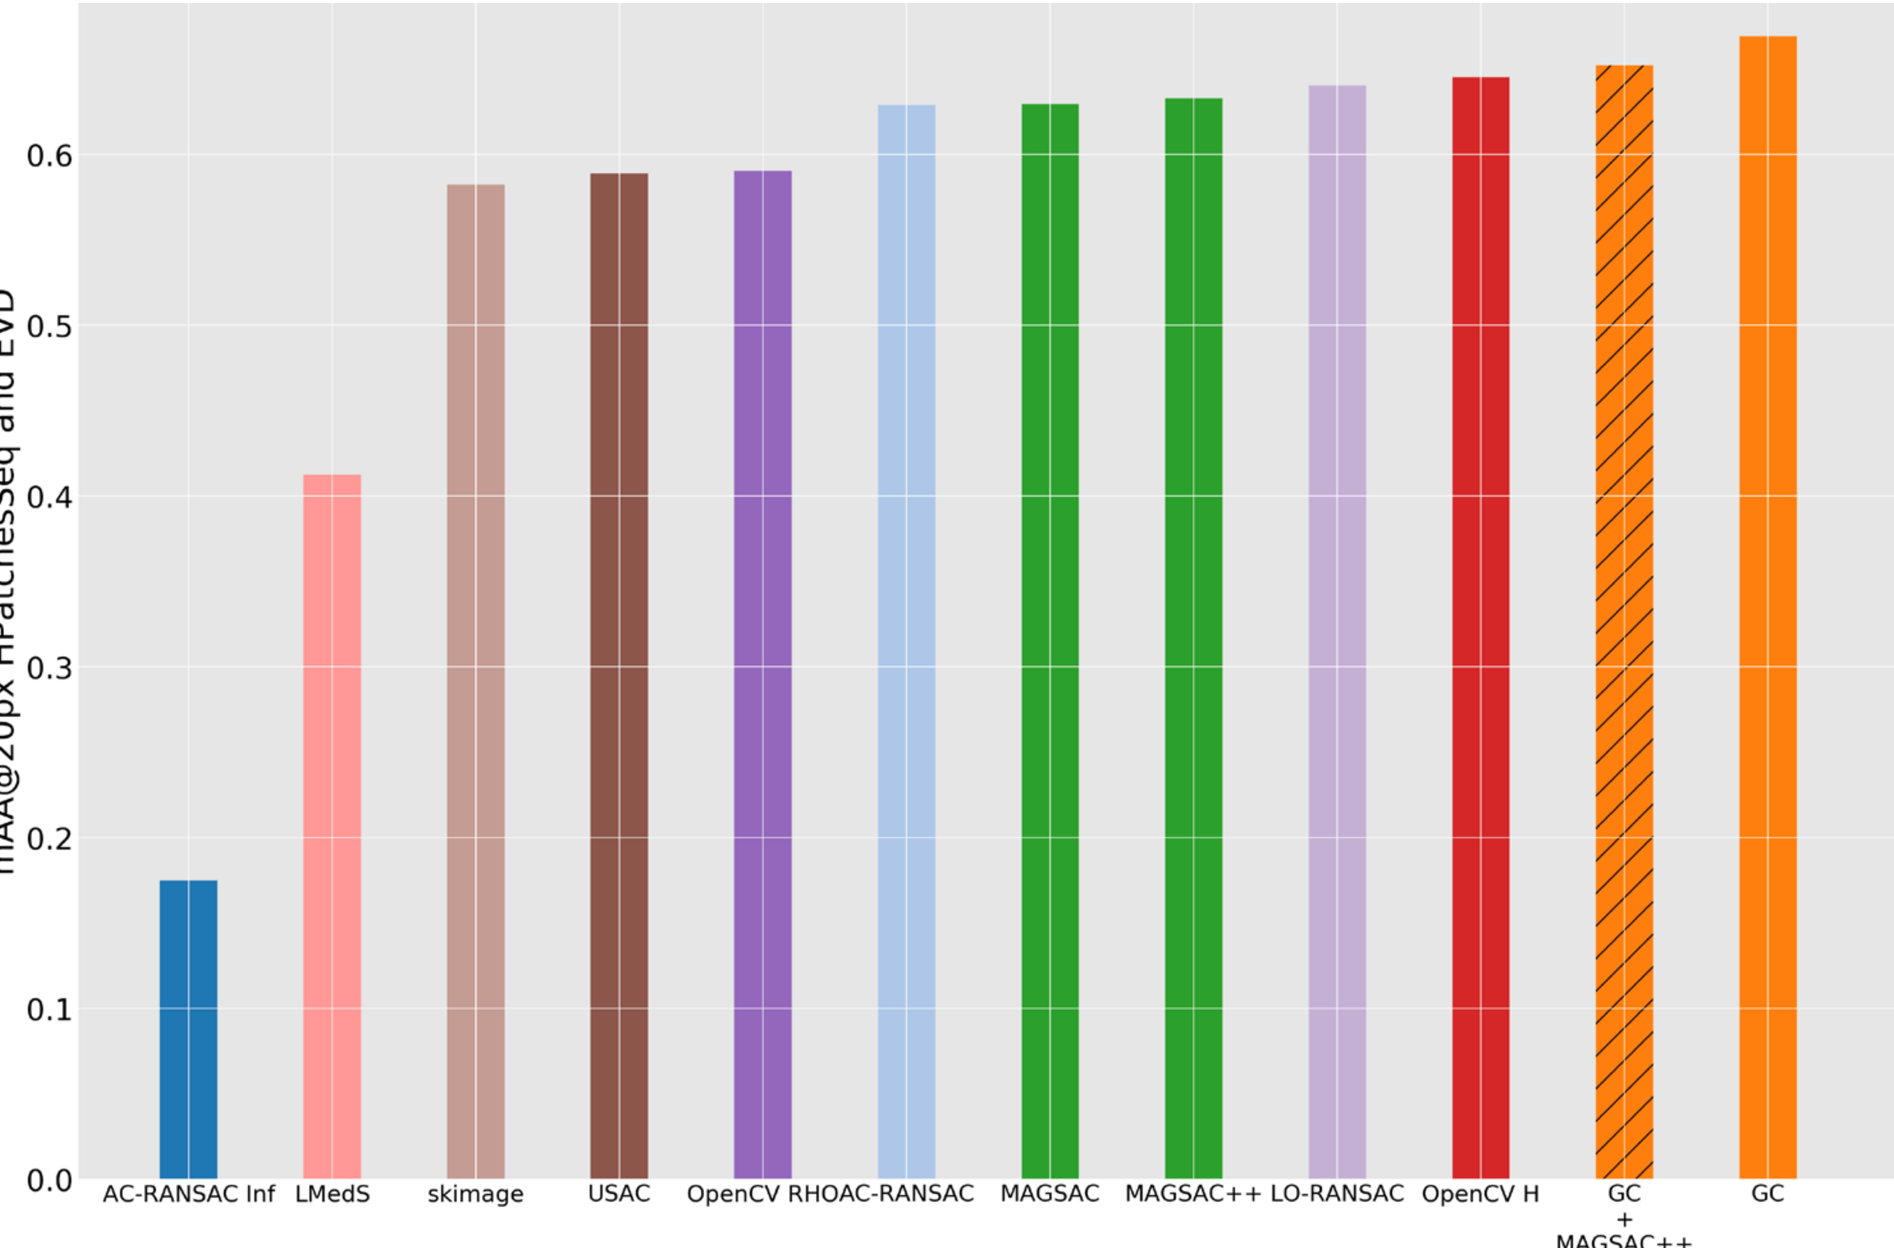

Classical methods, H, 1M iterations

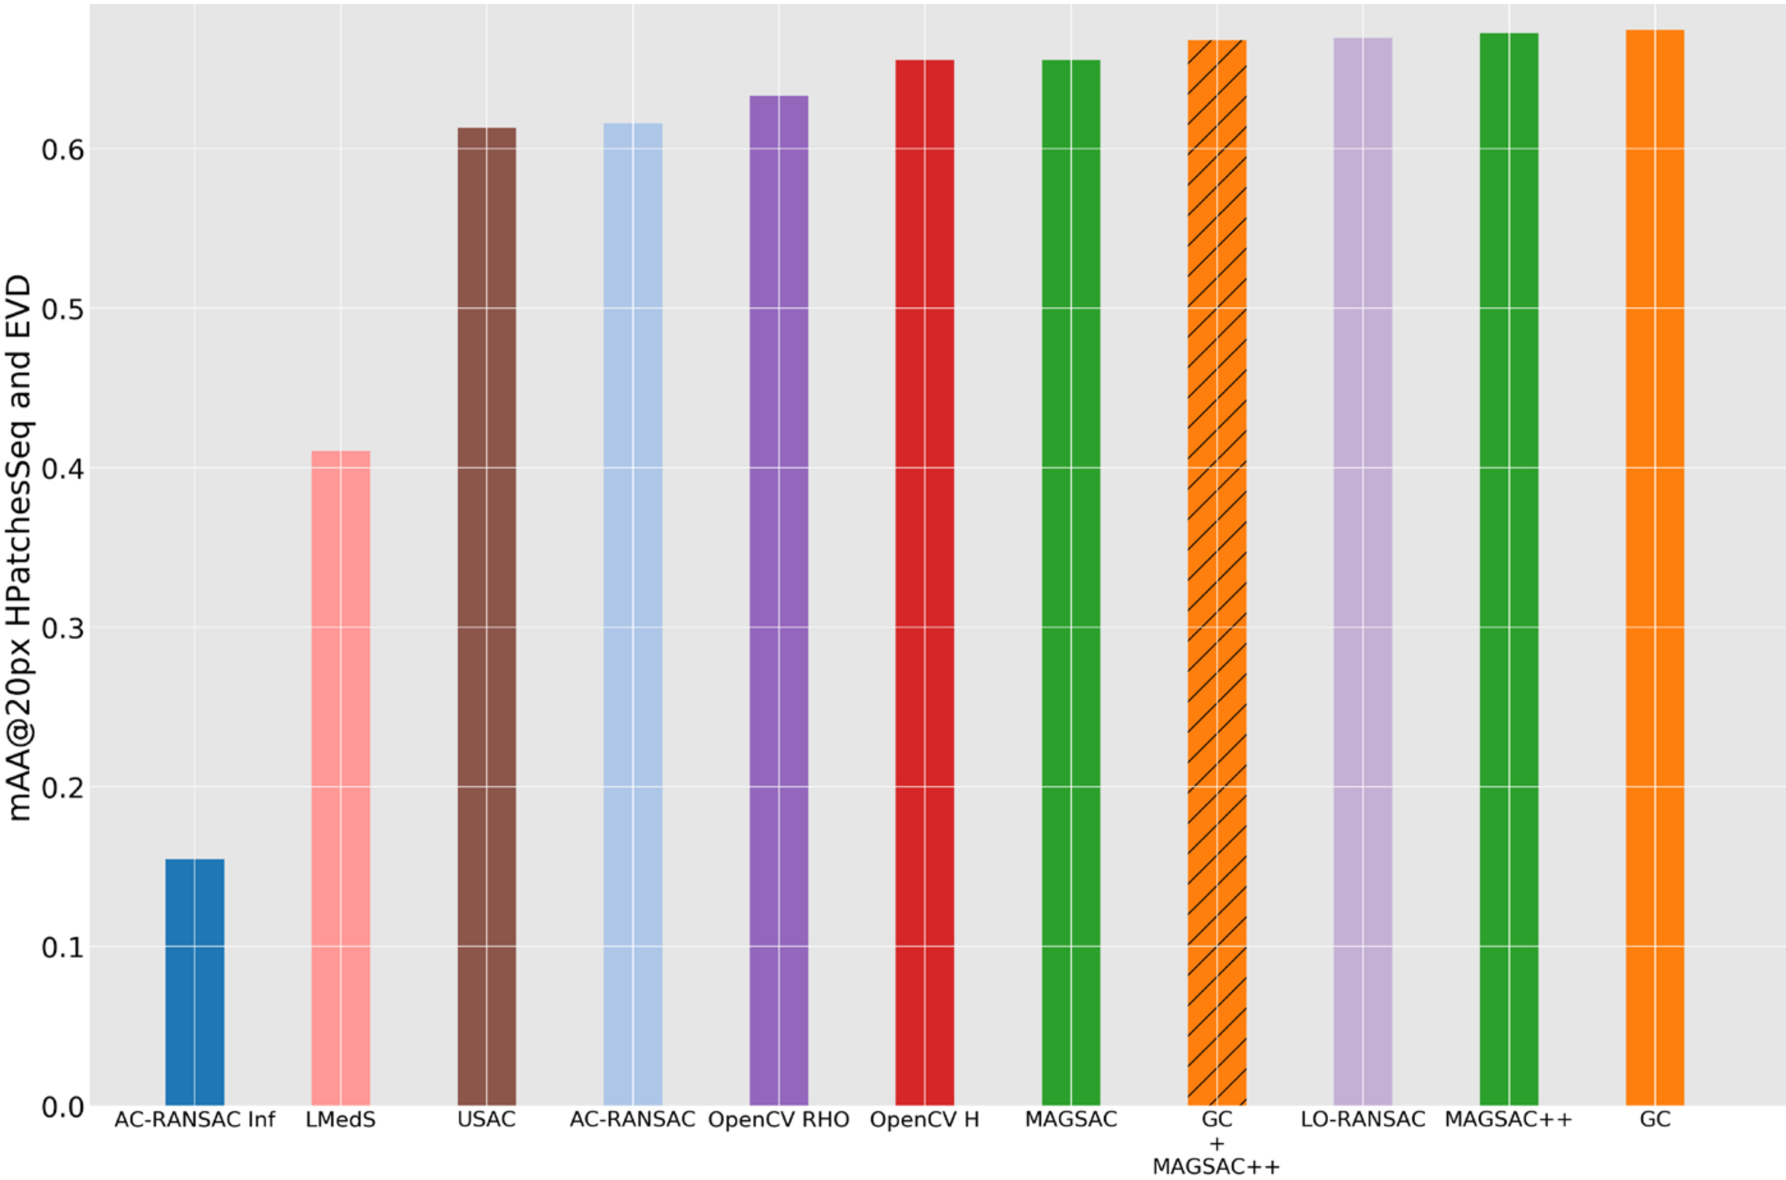

Conclusions

1. Fitting of models is an old and yet open problem.
   .RANSAC is the most popular robust fitting method.
2. Its simplicity is a huge plus. On the other hand, its assumptions allow for multiple improvements.
3. It achieves state of the art with local optimization treated as a two-class labelling problem solved by graph cut
4. Marginalization over the outlier threshold improves precision and removes the need for an input parameter.
5. The labelling formulation extends to multiple instance problem.

Thank you for your attention.

A contrario RANSAC

Moisan, L., Moulon, P. and Monasse, P., Fundamental matrix of a stereo pair, with a contrario elimination of outliers, Image Processing On Line 6 2016

Objective: Revisit the MINPRAN idea.

- Find the best model parameters together
- with the best inlier-outlier threshold (i.e., the noise scale).

Idea:

- For each minimal sample model, test multiple thresholds.
- Return the (model M, threshold ϵ_k) pair with the fewest inconsistent points closer than the threshold.
- Inconsistent points are called *false alarms.*

Problem: find model M where

The point-to-model residuals ordered increasingly

$$M = \arg\min_{M} \min_{k=m+1\,..\,n} \mathrm{NFA}(\{\epsilon_i(M)\}, k),$$

Sample size

Expected number of false alarms

Number of False Alarms (NFA)

Idea in brief:

- The background model is assumed to have uniform distribution.
- We are looking for a set of data points, where the probability of being uniformly and independently distributed is low.
- False alarm: a data point inconsistent with the randomly distributed background model.
- NFA is the expected number of false alarms given a set of points.

Moisan, Lionel, and Bérenger Stival. ***A probabilistic criterion to detect rigid point matches between two images and estimate the fundamental matrix.*** *International Journal of Computer Vision 57.3 (2004): 201-218.*

Example of determining the threshold by AC-RANSAC given a set of 2D points, a minimal sample (red dots) and the 2D line it implies (red line). The currently tested threshold (blue) and the best one (green) are shown.

The best threshold is the one minimizing the NFA value.

Conclusions of this Section and Everything

Even if RANSAC solves a different problem than multi-model fitting. It is a fundamental tool when approaching that problem as well.

Take home message of this presentation:

- The RANSAC inlier-outlier threshold is not trivial to set.
- Geometric data is spatially coherent. Use it.

Questions?

How to Ignore „Bad” Models?

Idea: models with good enough inlier supports are not distant from the true underlying model.

Algorithm (LO-RANSAAC): when a new best model is found,

1. Get its inliers and refit the model.
2. Store the polished model by converting it to points.
3. Start again from 1. with iteratively shrinking threshold.

4. The averaging is applied to these intermediate models.

Number of False Alarms (NFA)

Moisan, L., Moulon, P. and Monasse, P., Fundamental matrix of a stereo pair, with a contrario elimination of outliers, Image Processing On Line 6 2016

The probability of a random data term to have error at most 1

Index of the currently tested residual

Point number

Residual dimension

$$\mathrm{NFA}(\{\epsilon_i : i = 1,..,n\}, k) = N_{outcomes}(n - m)\binom{n}{k}\binom{k}{m}\left(\epsilon_k^d \alpha_0\right)^{k-m}$$

Number of models estimated from a minimal sample, e.g., 3 for fundamental matrix

Residual of the kth closest point