Recent Advances on RANSAC

Presenter: Dániel Baráth

RANSAC in 2020: A CVPR tutorial

Organizers: Daniel Barath, Jiri Matas, Ondra Chum,
Tat-Jun Chin, Rene Ranftl, Dmytro Mishkin

Main topics of the talk

Baseline: USAC, Raguram et al. TPAMI 2012

1. **Robust estimators without inlier-outlier threshold.**

   i) Finding the threshold adaptively (MINPRAN, a contrario RANSAC);

   ii) Averaging minimal sample models (RANSAAC);

   iii) Averaging over a range of noise scales (MAGSAC, MAGSAC++)

2. **Exploiting the spatial coherence of geometric data.**

   i) Local optimization on spatially coherent structures (Graph-Cut RANSAC)

   ii) Sampling minimal samples in a local-to-global manner (NAPSAC, Progressive NAPSAC)

3. **Multi-model fitting with RANSAC-like methods.**

   i) Algorithms applying RANSAC directly (sequential RANSAC, multiRANSAC).

   ii) State-of-the-art multi-model fitting.

Robust estimators without inlier-outlier threshold (V1)

Objective is traditionally

- Return the sought model parameters and the inliers.
- The two problems considered eqivalent, since
  - Inliers -> model easy, i.e., the standard estimation problem
  - Model -> inliers easy, *if you have the inlier outlier threshold.*

Note: the inlier - outlier decision is crisp.

Approaches discussed:

- Find the threshold adaptively from the data
  (MINPRAN, a contrario RANSAC).

Robust estimators without inlier-outlier threshold (V2)

Objective is

- Return the sought model parameters,
- and (*optionally*) the inliers
- without requiring a manually set inlier-outlier threshold.

Approaches discussed

- Generate minimal sample models and average them (RANSAAC).
- Marginalize over the threshold, thus, not making strict inlier-outlier decisions (MAGSAC, MAGSAC++).

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.

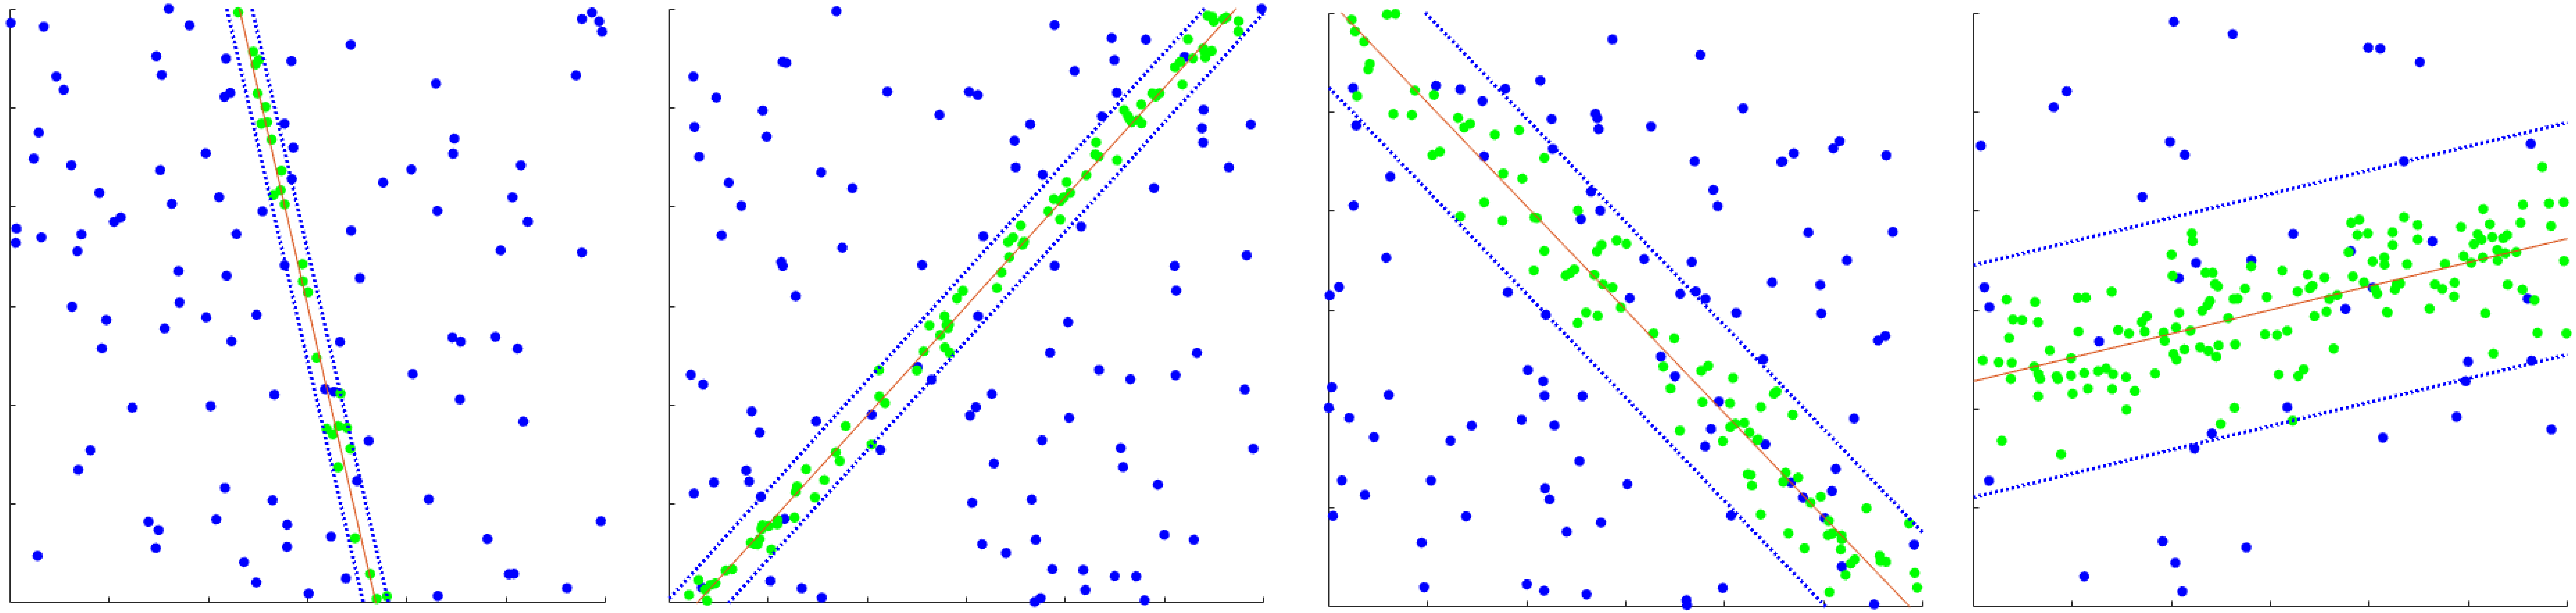

Randomly generated points (green) on a 2D line (red) and outliers (blue dots). The threshold which the synthetic noise added to the inlier point coordinates' implies is shown by blue dotted lines.

Robust Model Estimation (Standard Setting)

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

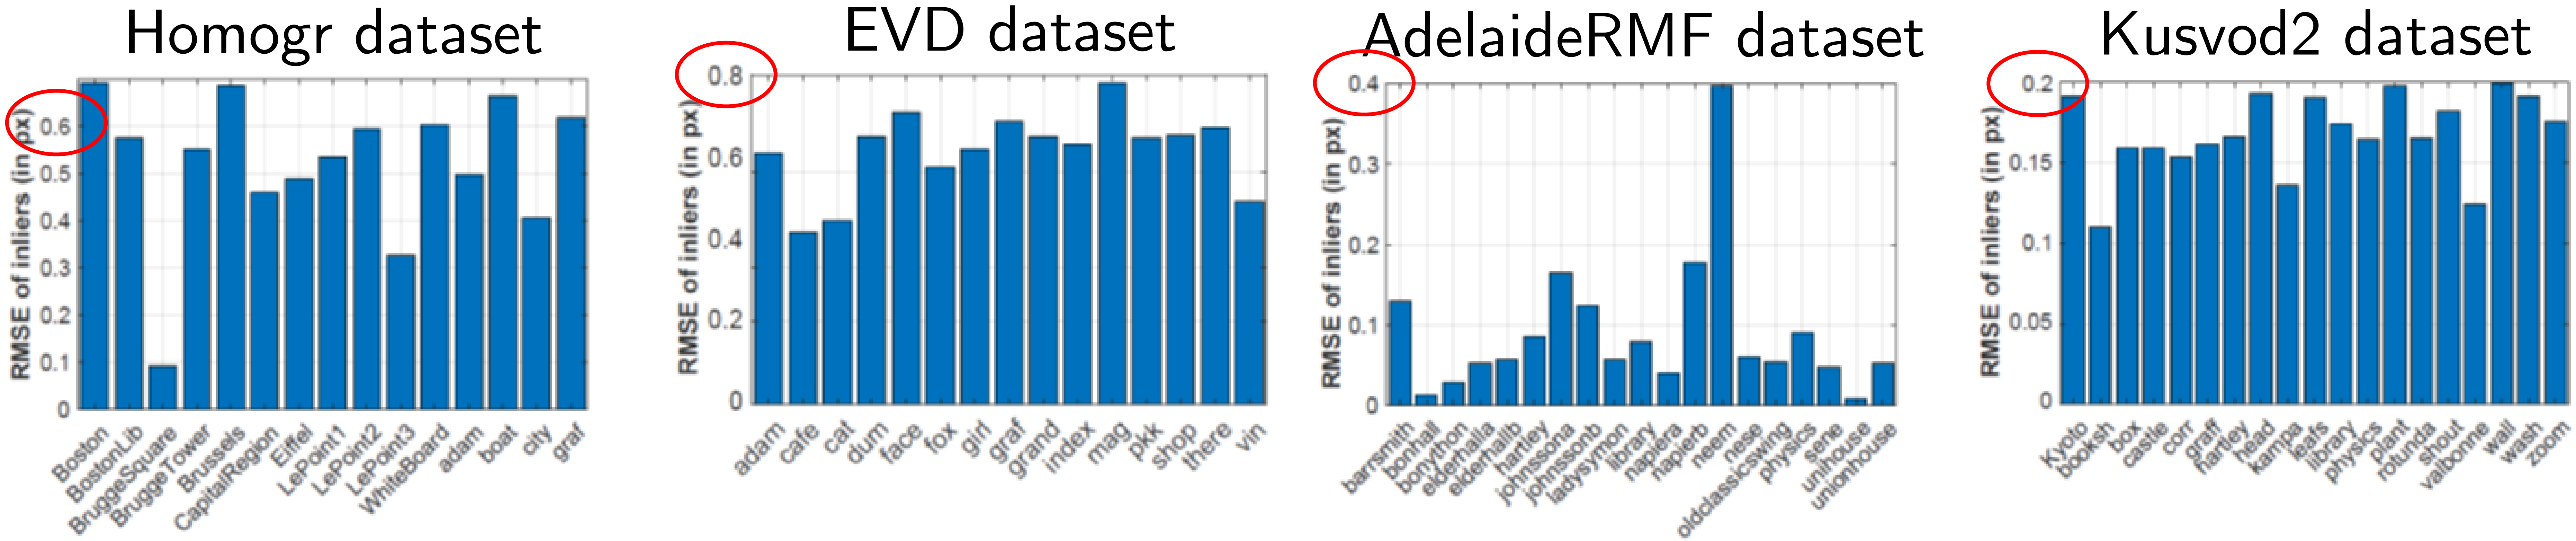
Homogr dataset

EVD dataset

AdelaideRMF dataset

Kusvod2 dataset

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. What happens if:

  (i) there is rolling shutter,

  (ii) there is radial distortion,

  (iii) the scene is not completely rigid,

Example of rolling shutter.

Example of radial distortion.

Example Sources of Model Inaccuracy

- **Imperfection of the residual function**

  E.g., zero symmetric epipolar distance does not necessarily mean a good correspondences.

- **Other things making the threshold unclear**

  (i) image size is different

  (ii) far/close object...

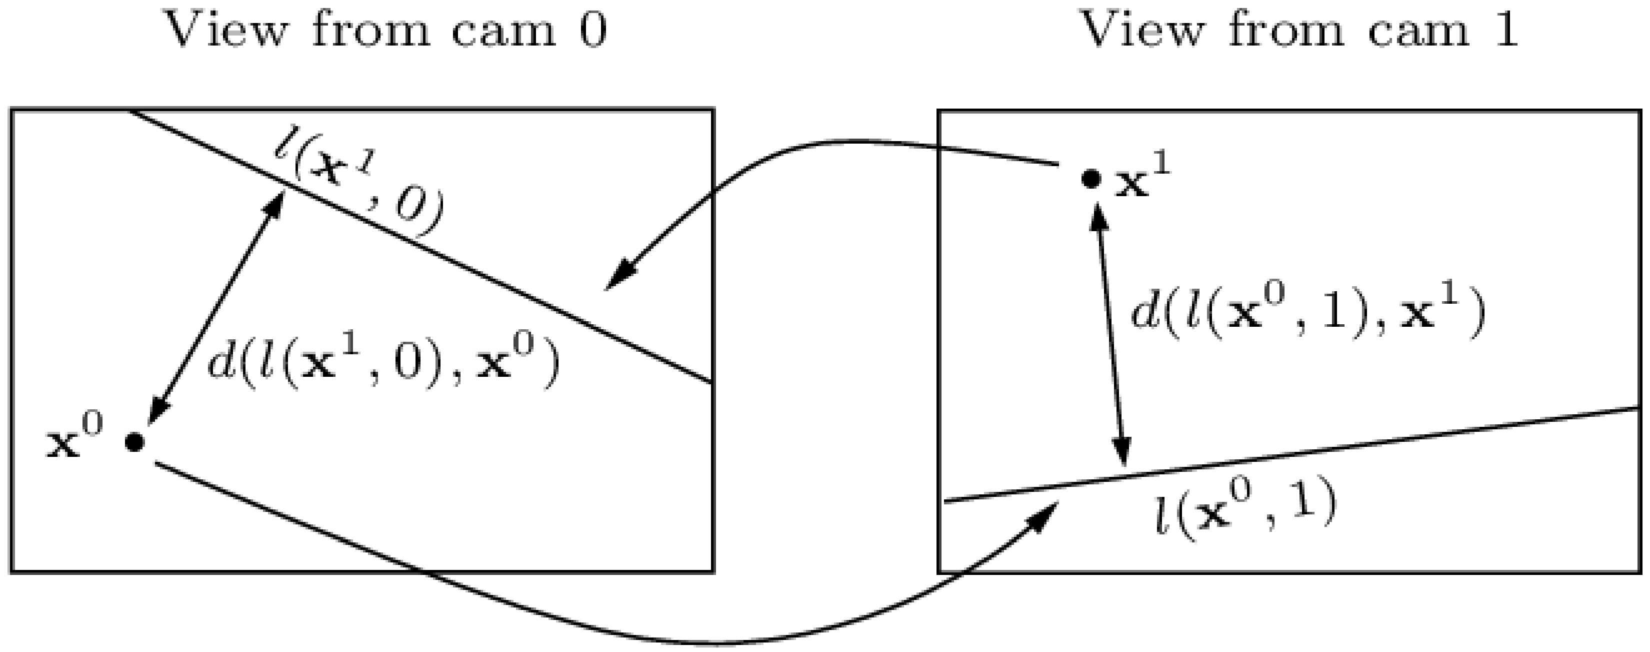

Symmetric epipolar distance.
The point can move on the line arbitrarily with 0 residual.

Conclusion: setting the threshold is really hard without knowing anything about the scene.

Close and far objects.

The same threshold would not work for both.

Example Sources of Model Inaccuracy

- **Imperfection of the residual function**

  E.g., zero symmetric epipolar distance does not necessarily mean a good correspondences.

- **Other things making the threshold unclear**

  (i) image size is different

  (ii) far/close object...

Conclusion: setting the threshold is really hard without knowing anything about the scene.

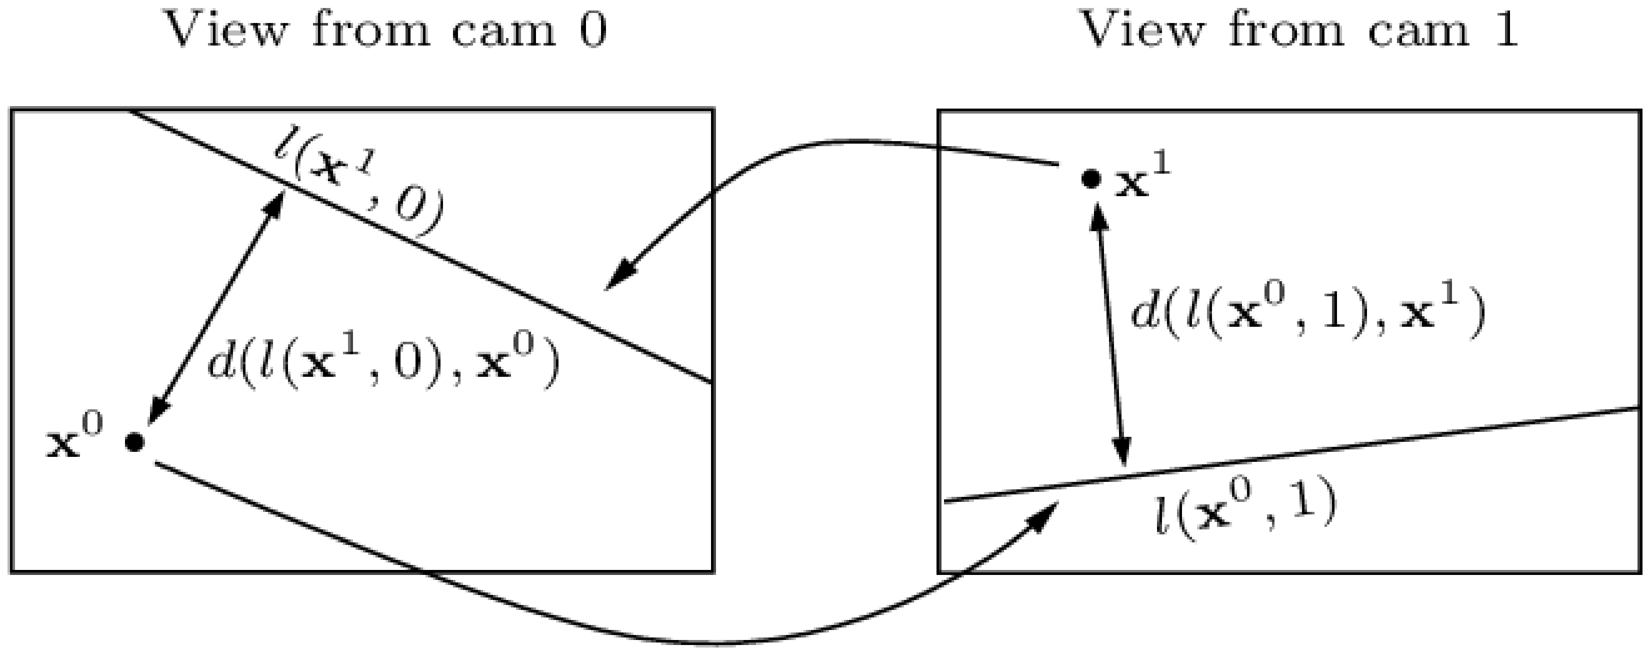

Symmetric epipolar distance.
The point can move on the line arbitrarily with 0 residual.

Close and far objects.

The same threshold would not work for both.

Example Sources of Model Inaccuracy

- **Imperfection of the residual function**

  E.g., zero symmetric epipolar distance does not necessarily mean a good correspondences.

- **Other things making the threshold unclear**

  (i) image size is different

  (ii) far/close object...

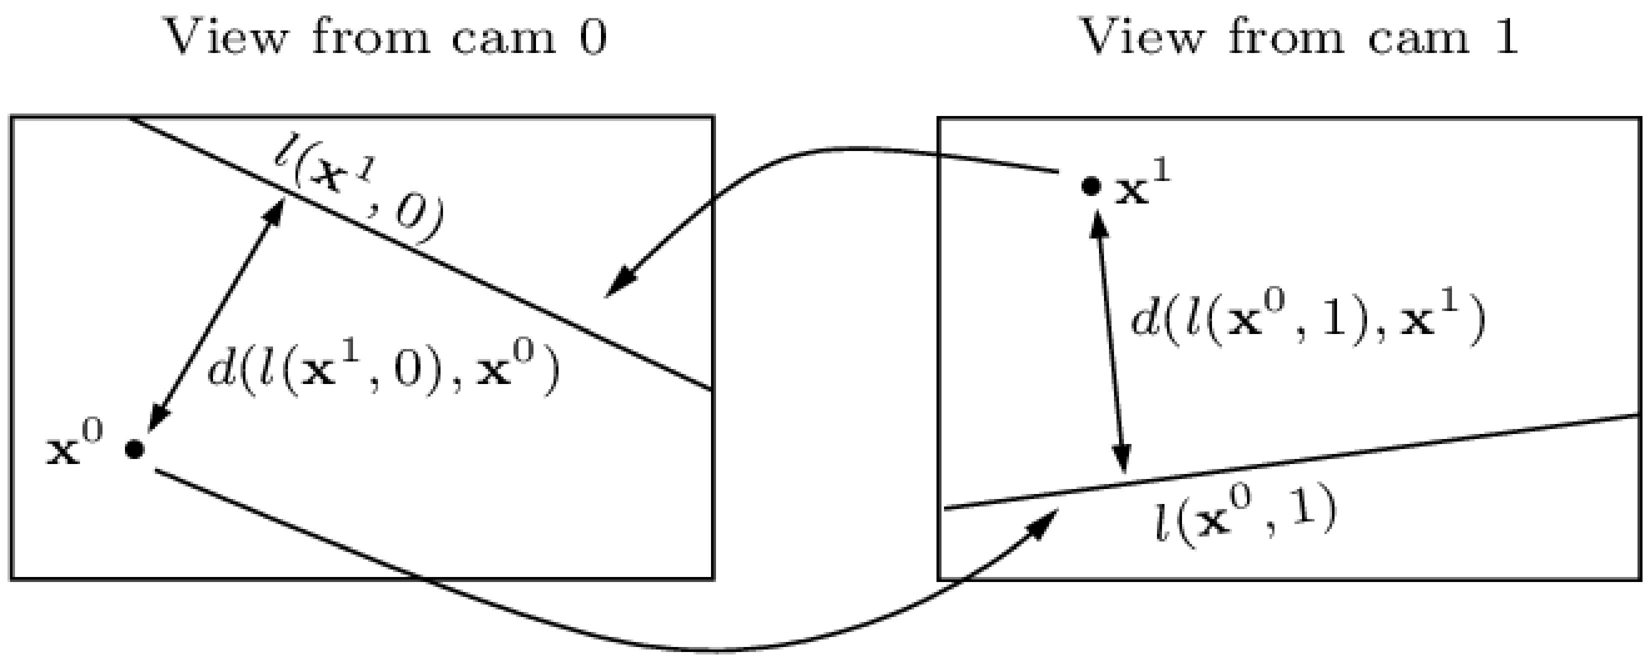

Symmetric epipolar distance.
The point can move on the line arbitrarily with 0 residual.

Conclusion: setting the threshold is really hard without knowing anything about the scene.

Close and far objects.

The same threshold would not work for both.

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 occured 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.

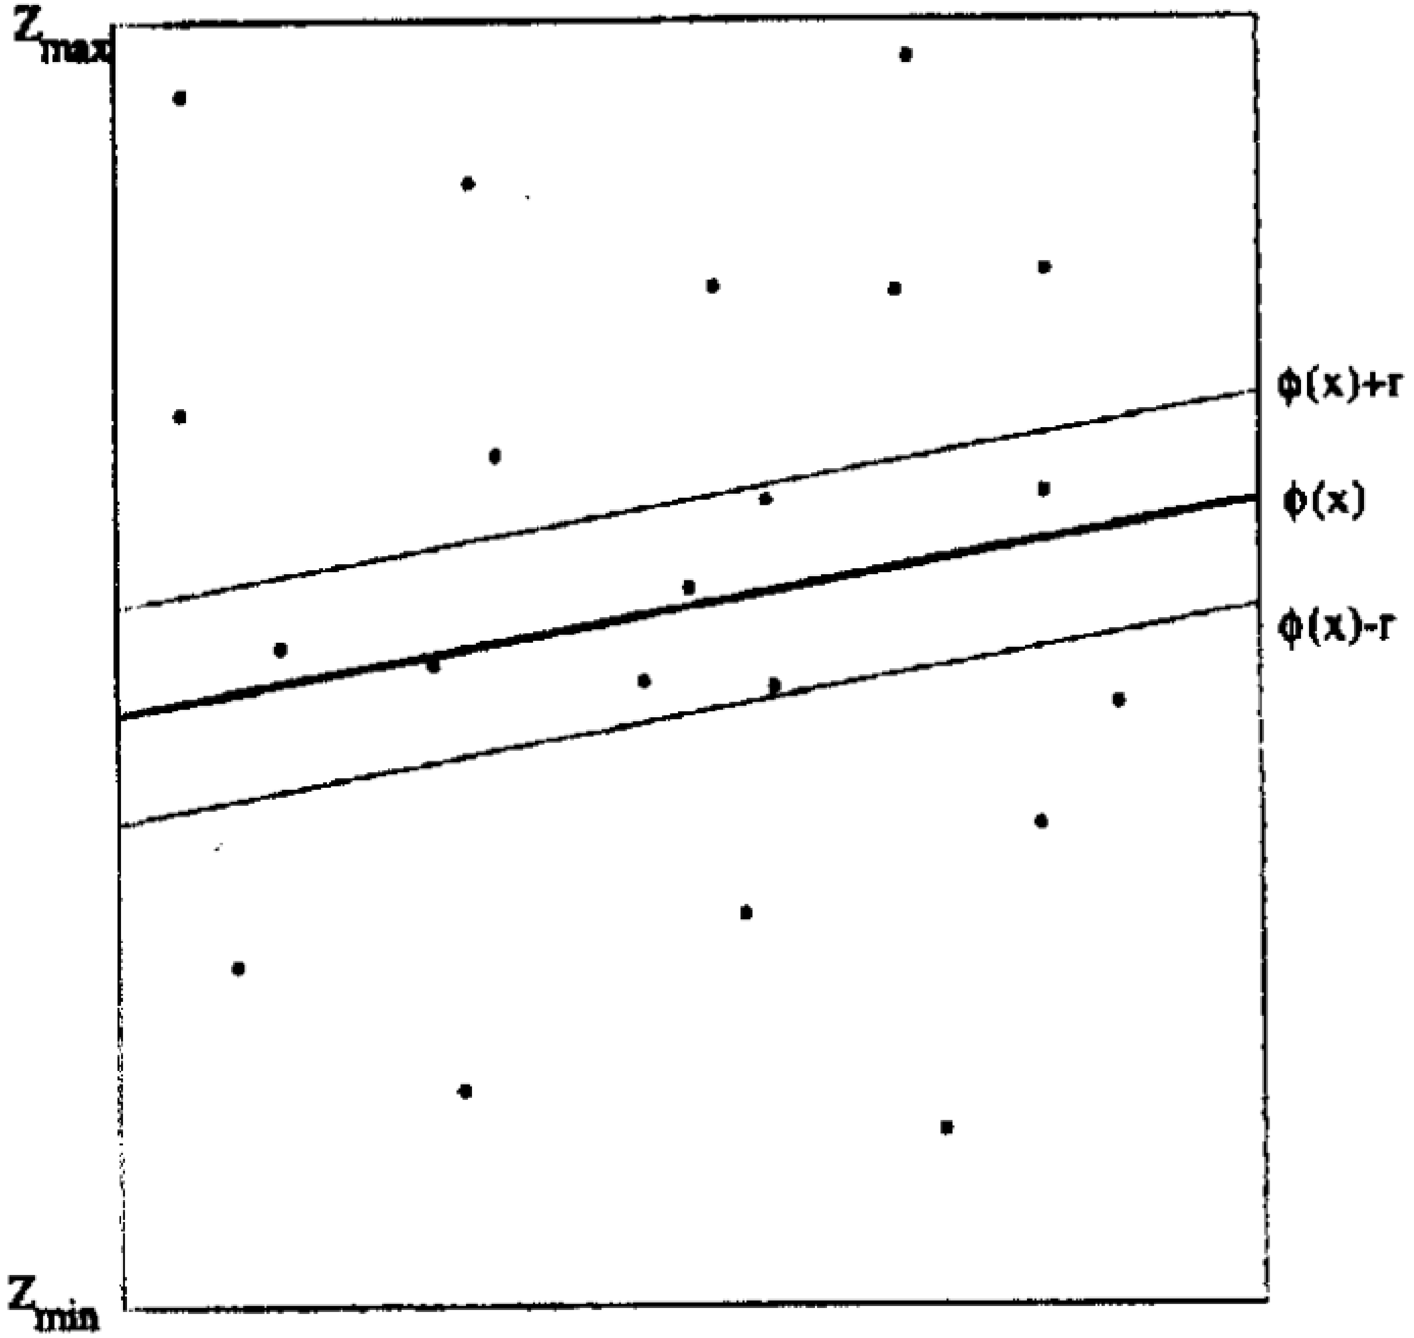

MINPRAN

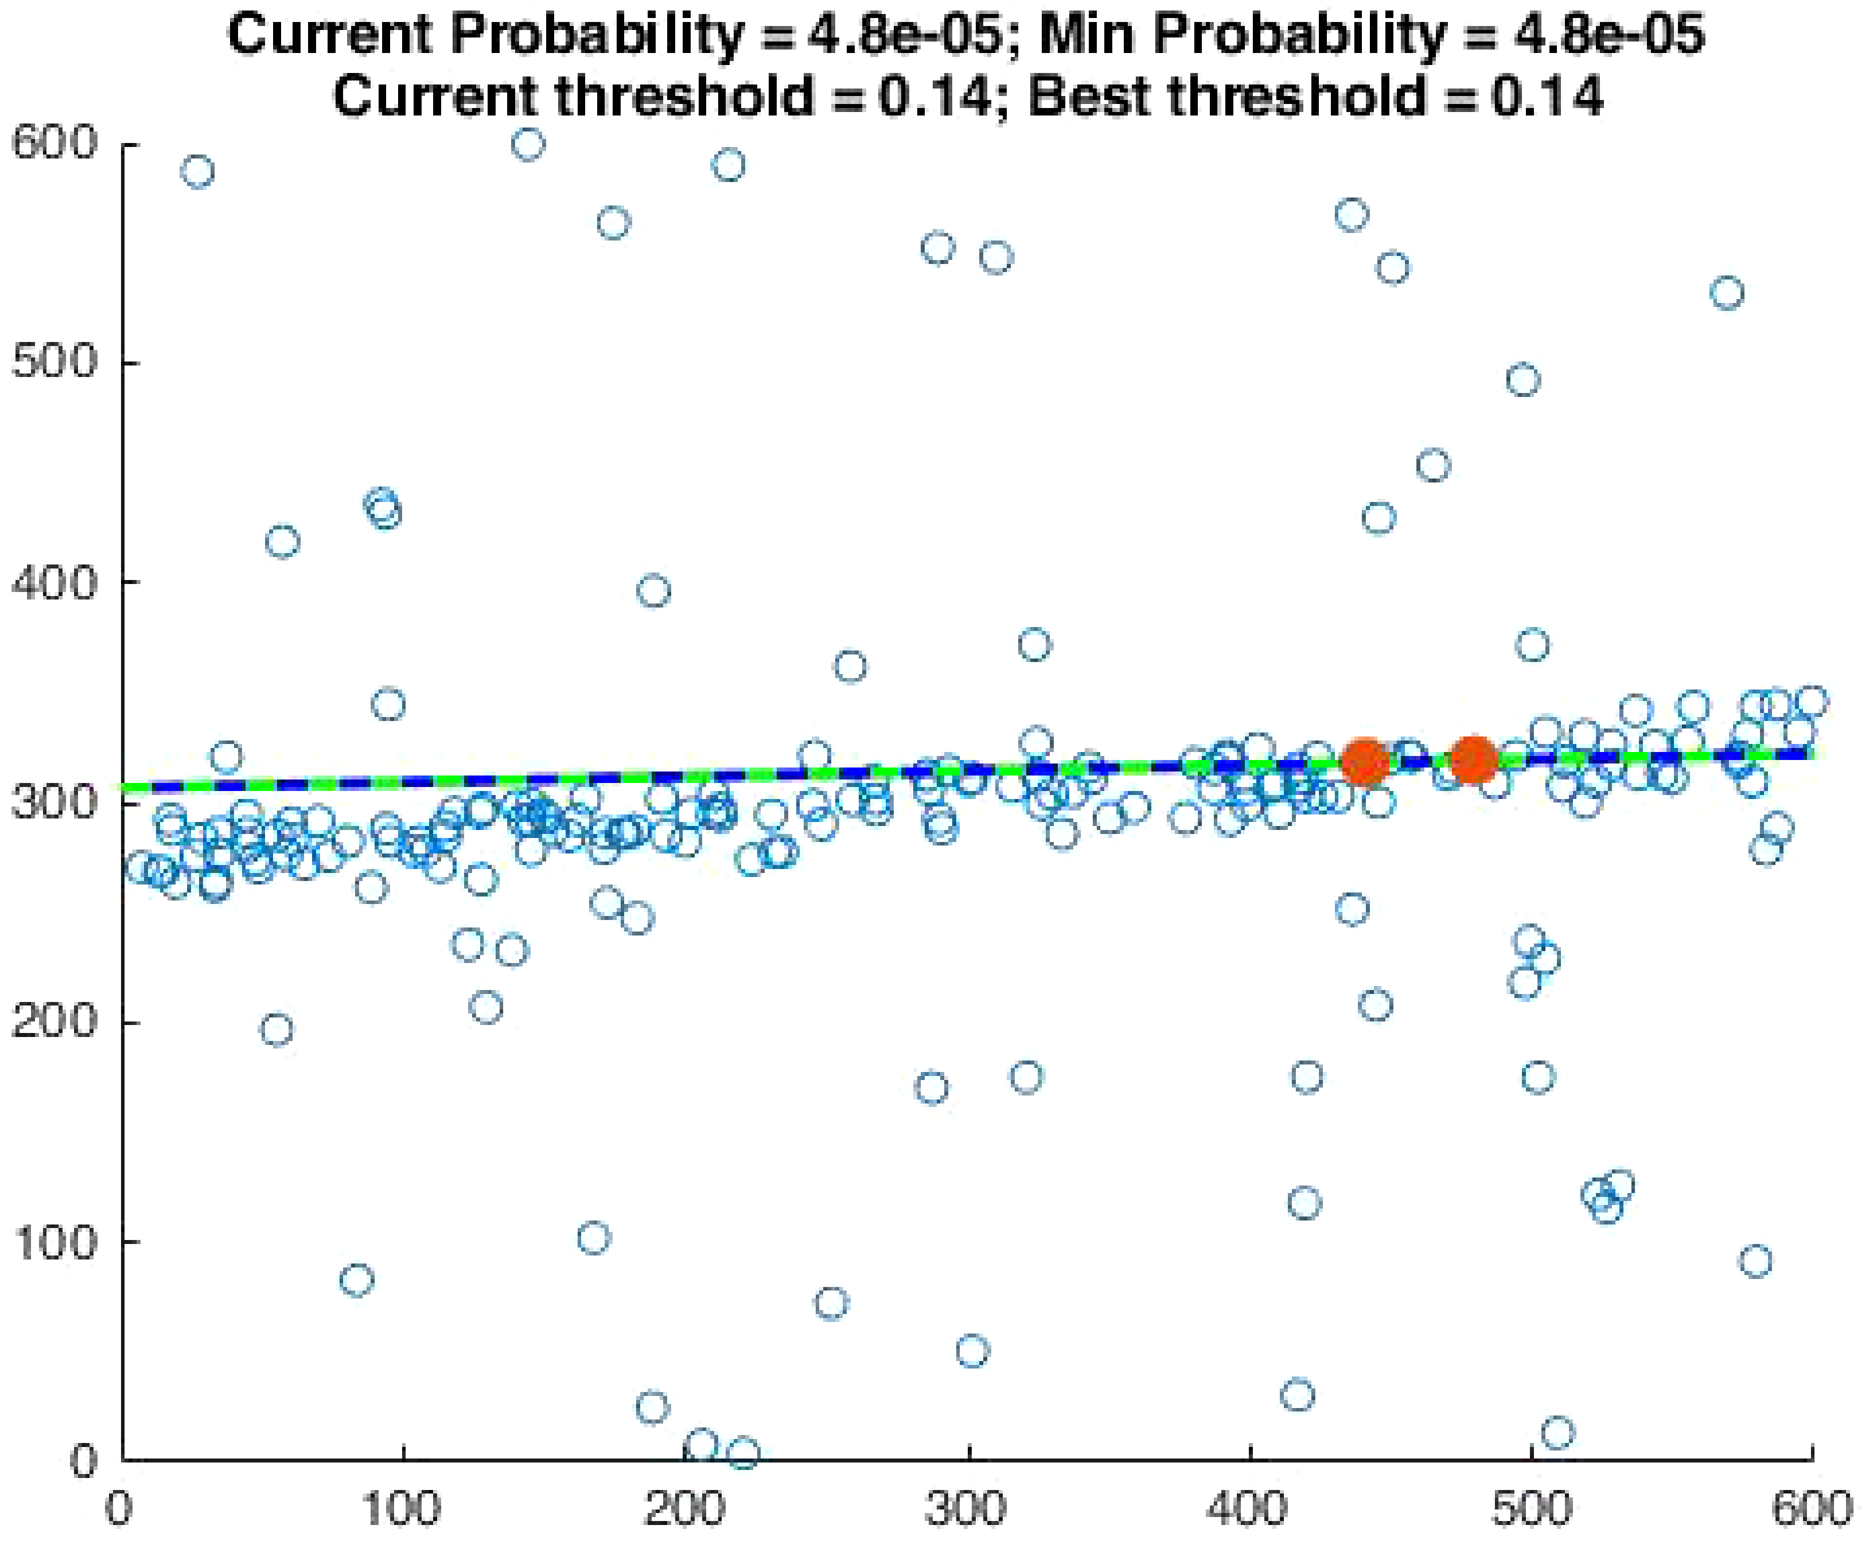

Example of determining the threshold by MINPRAN given a set of 2D points, a minimal sample (red dots) and the 2D line it implies (red line). The currently tested threshold (blue) and the best one (green) are shown.

The best threshold is the one minimizing the probability of randomness.

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.*

A contrario RANSAC

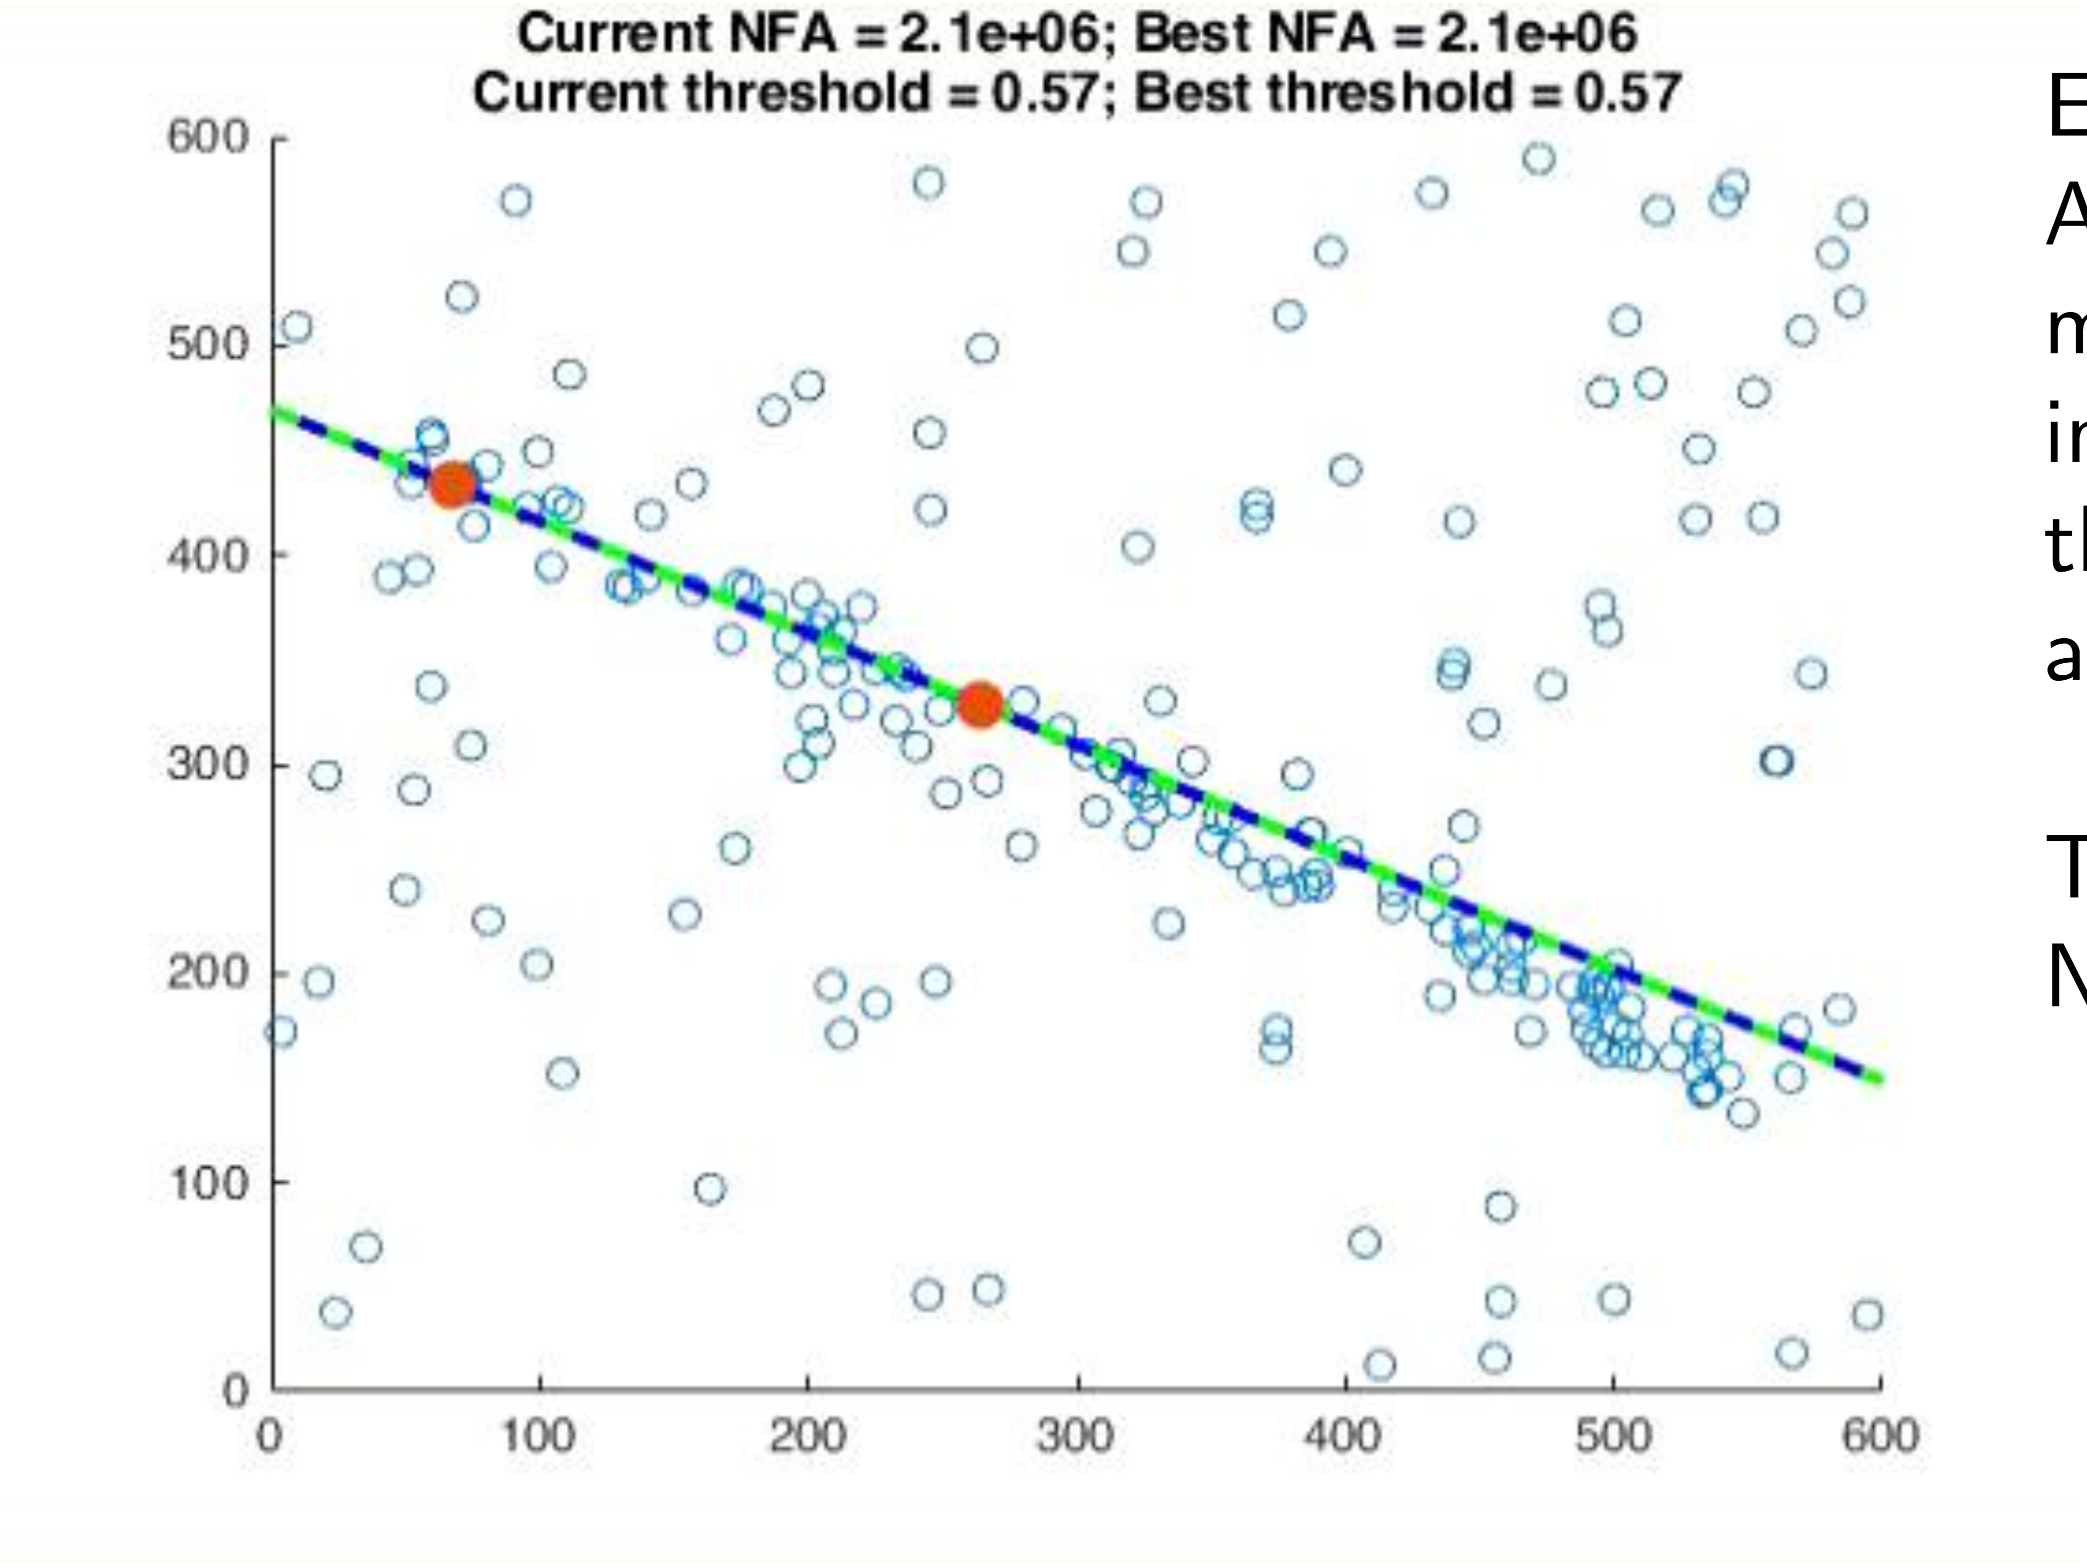

Example of determining the threshold by AC-RANSAC given a set of 2D points, a minimal sample (red dots) and the 2D line it implies (red line). The currently tested threshold (blue) and the best one (green) are shown.

The best threshold is the one minimizing the NFA value.

Random Sample Aggregated Consensus (RANSAAC)

Rais et al., Accurate motion estimation through random sample aggregated consensus. 2017

Observations motivating RANSAAC:

- Inlier-outlier threshold is hard to set (as shown earlier).
- RANSAC's results come from the best hypothesis, while everything else is ignored including useful „all-inlier'' samples.

Idea:

- Combine all generated hypotheses instead of using only the best one.
- Done by weighted averaging of the model parameters.

Issues/Questions:

- What is an average model?
- How weighted averaging can be robust? What weights are robust?

Random Sample Aggregated Consensus (RANSAAC)

Requirement: averaging has to be done in the model parameter space.

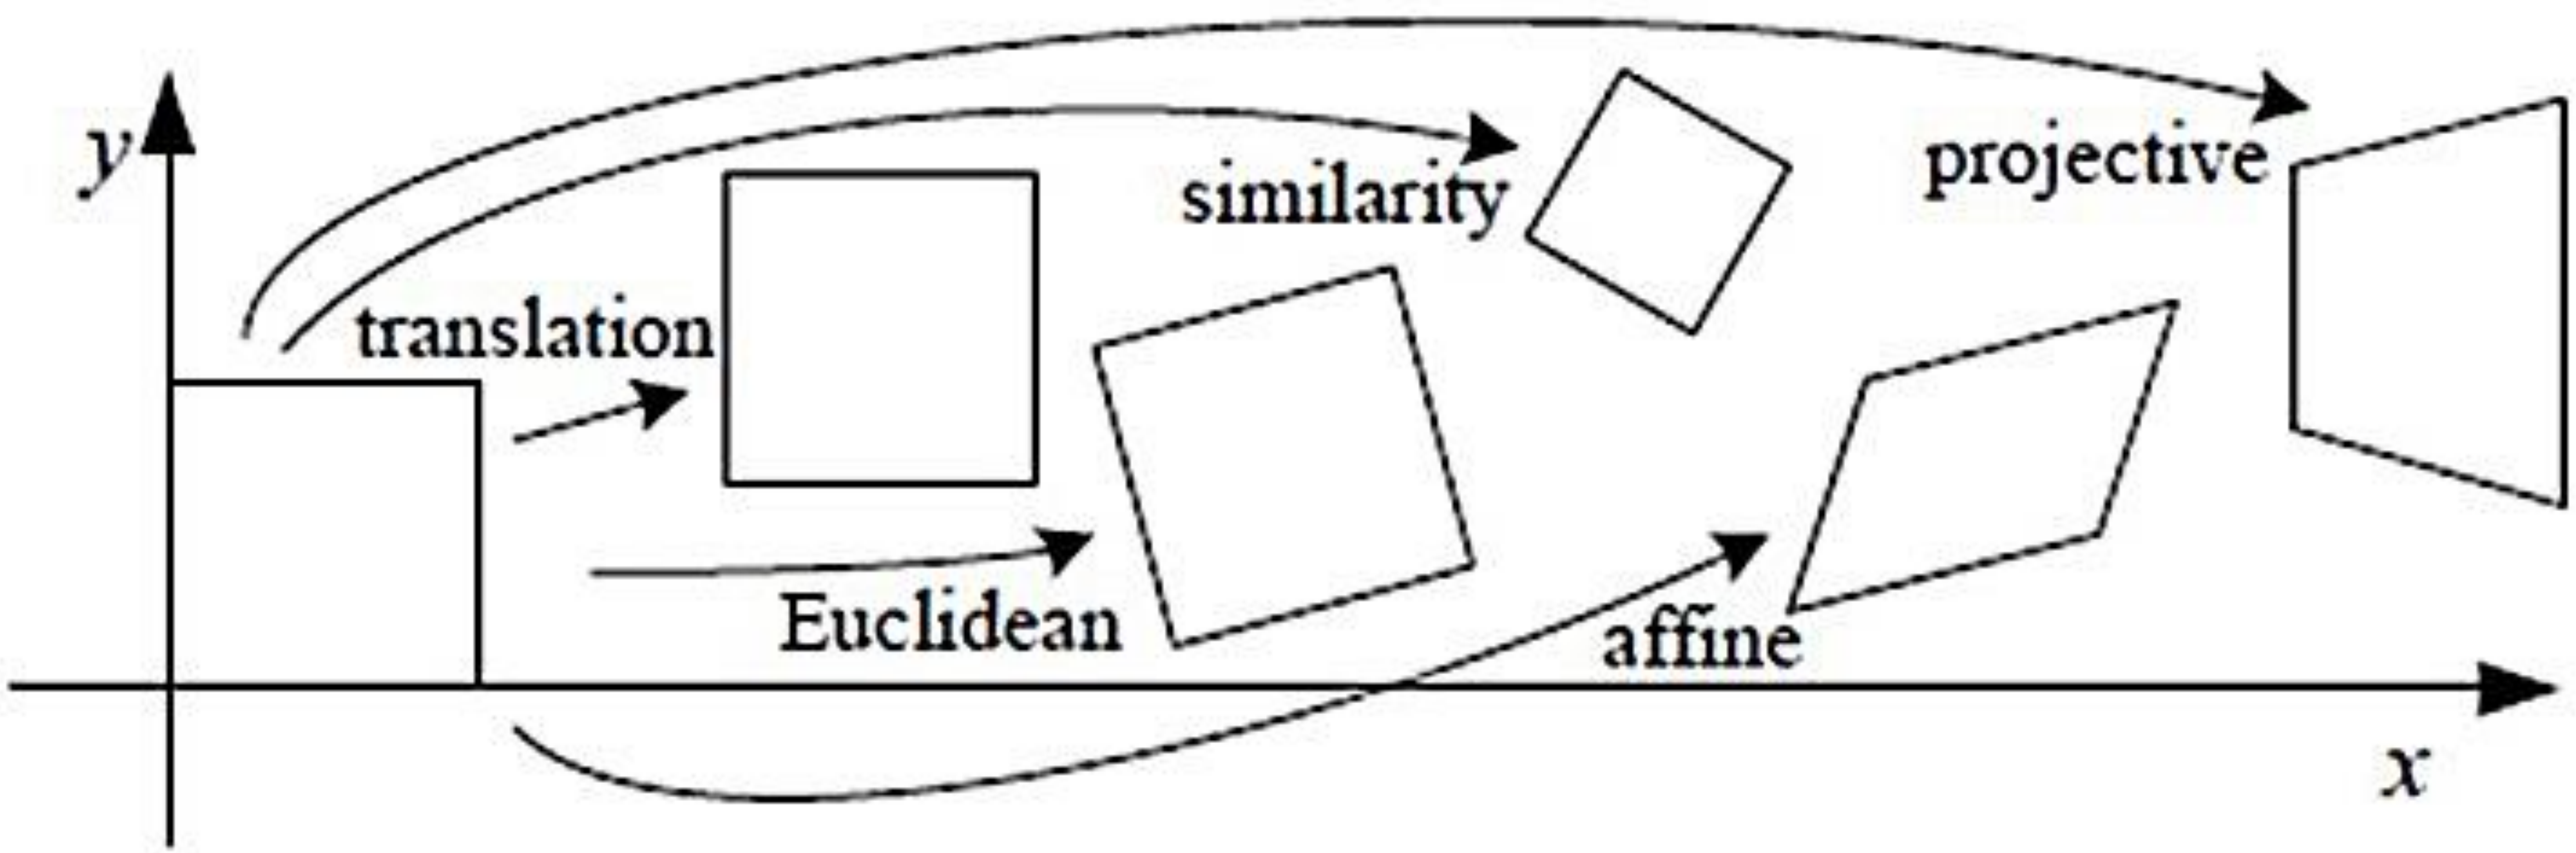

Example transformations to be estimate. We need to know how to average them.
E.g., what is the average of 2 affine transformations?

Averaging Models

A solution:

1. Select anchor points, e.g., corners of the image.
2. Transform the selected points by the estimated transformation.
3. Keep the transformed points and the model quality.
4. The averaging is done on the transformed points.

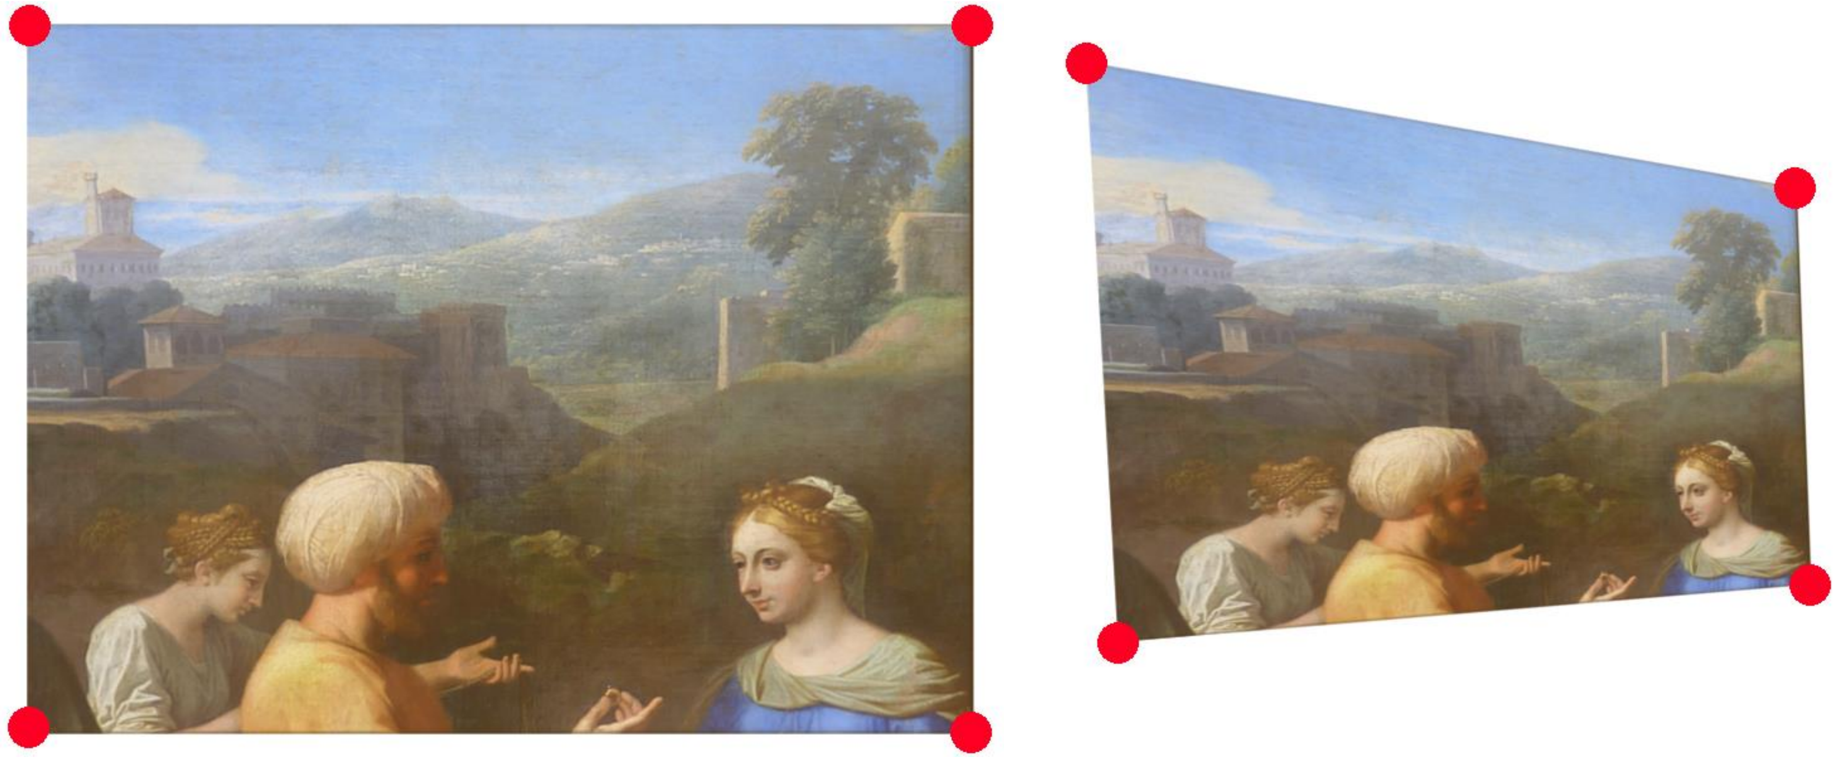

Anchor and transformed points of a homography

Anchor and transformed points of a radial distortion homography

Example

Point correspondences detected in an image pair.

Example

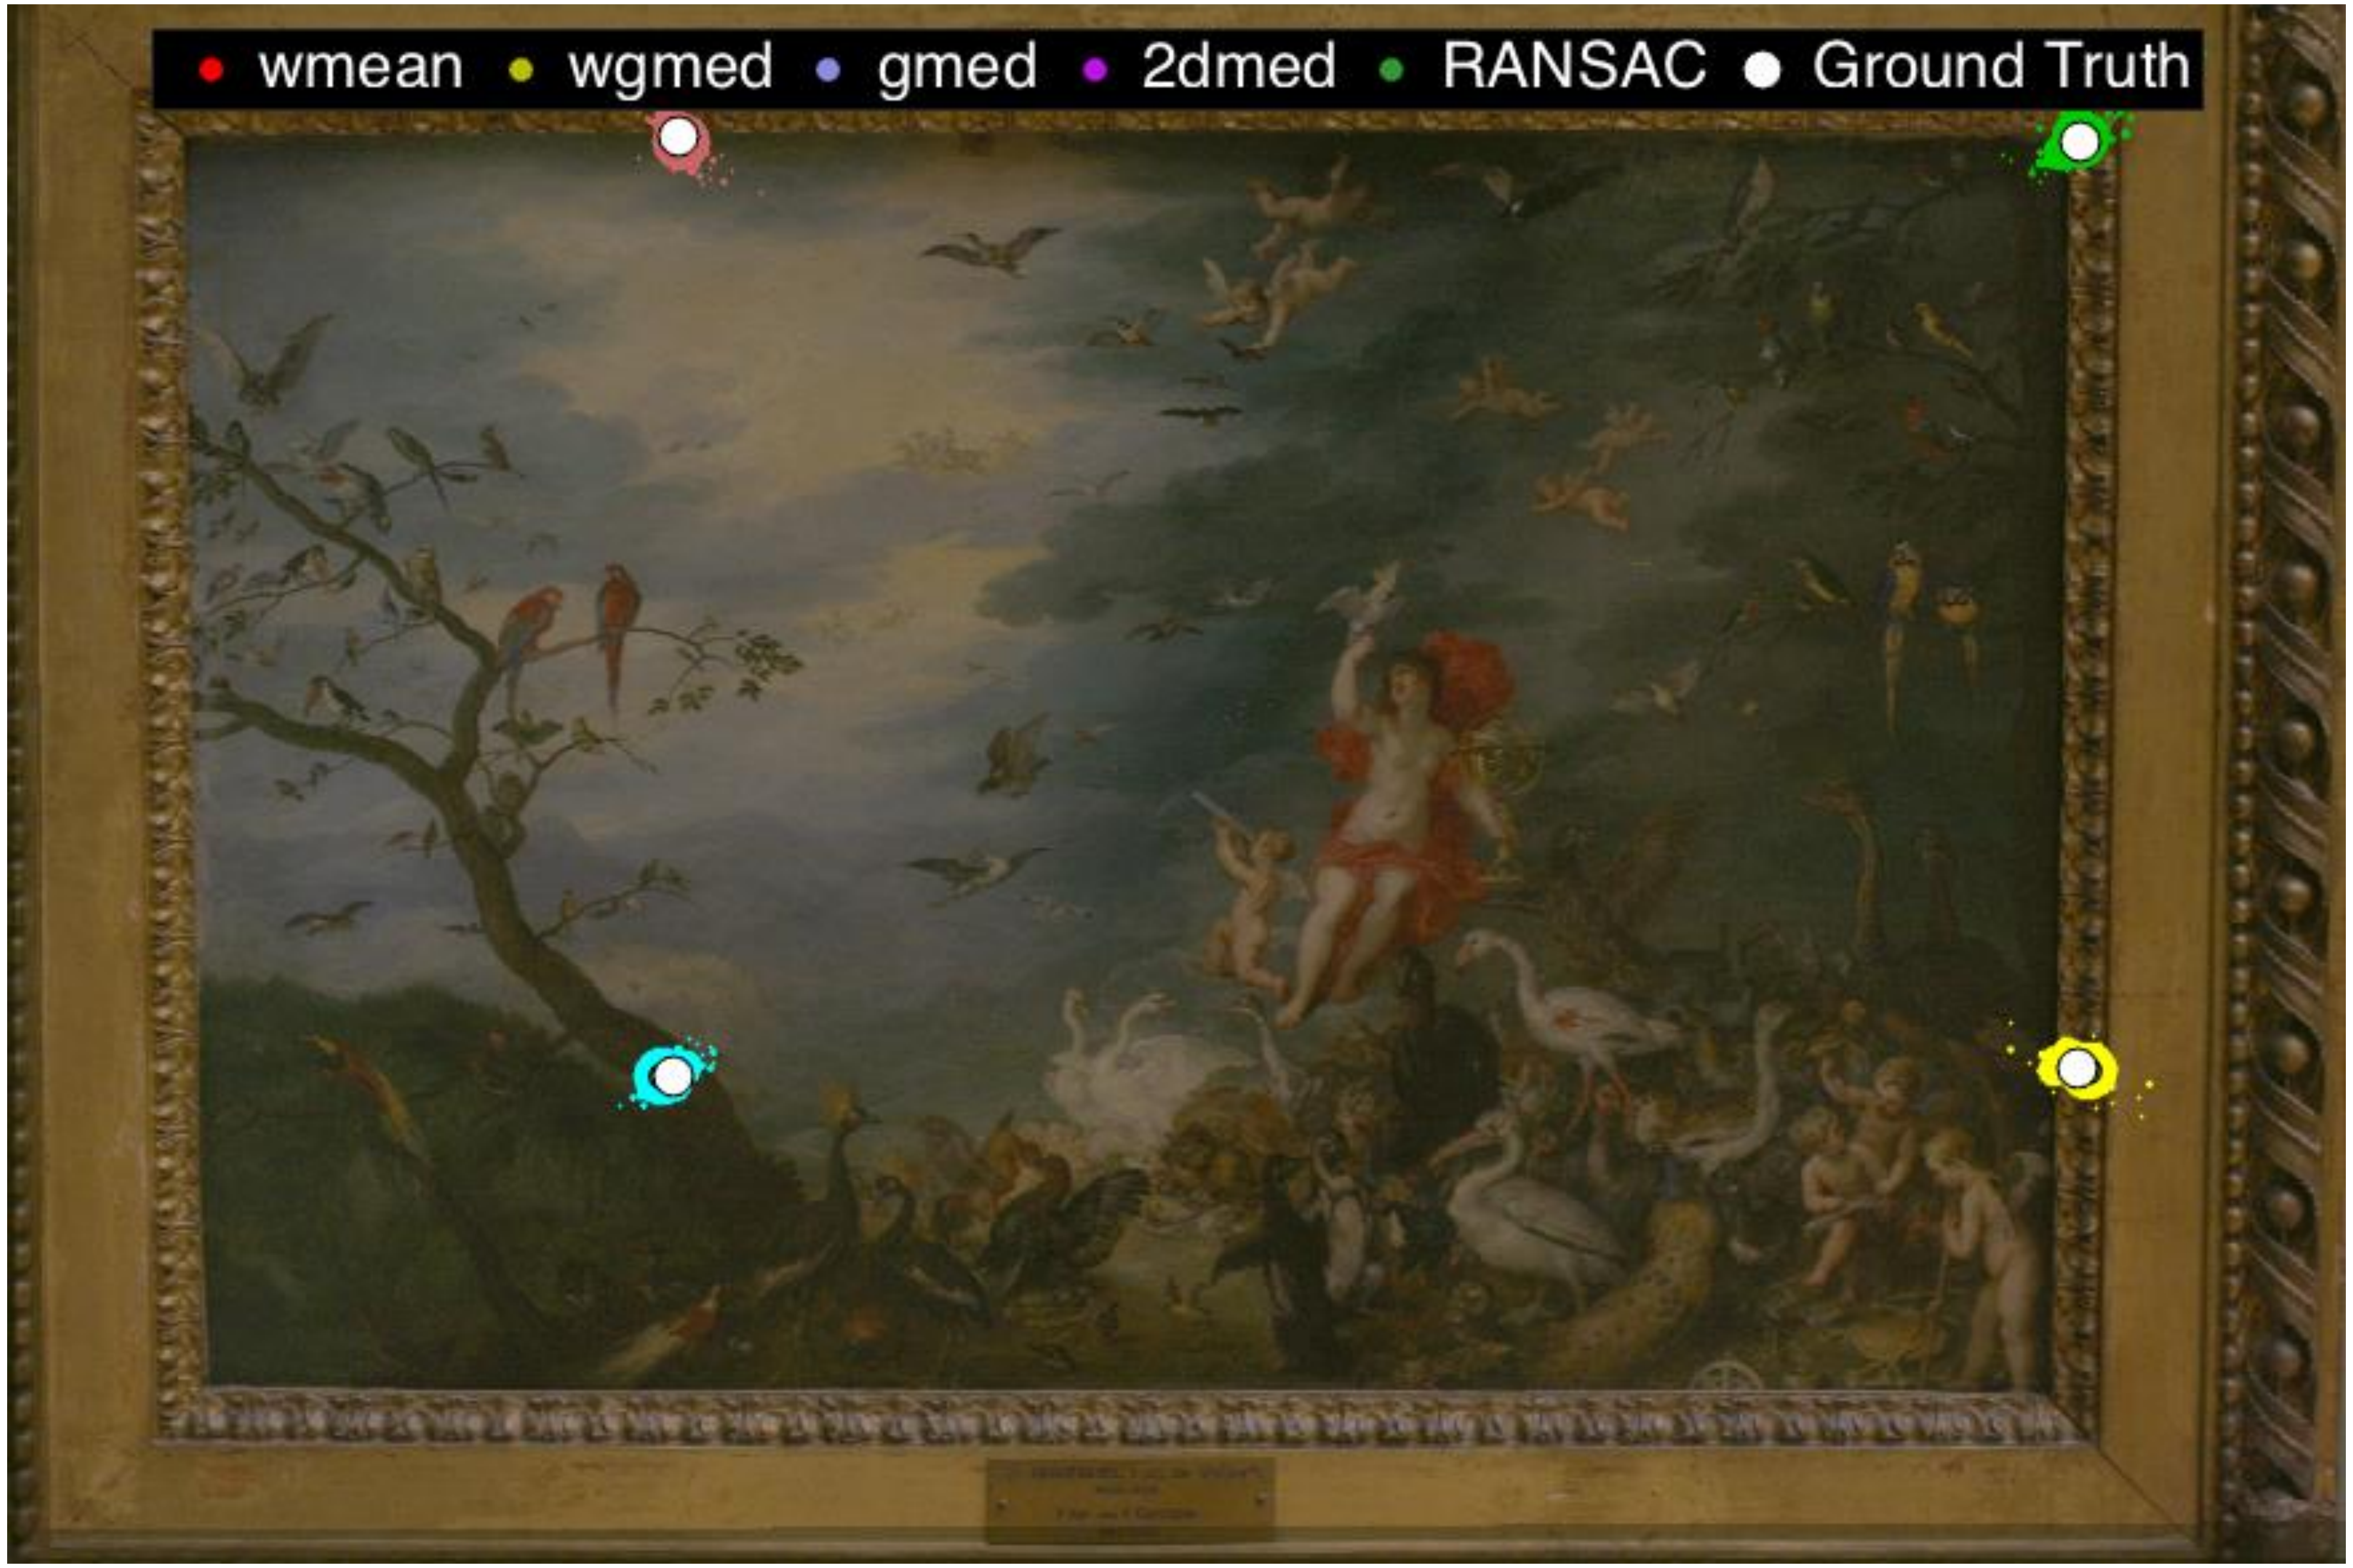

The corners (red / green / blue / yellow) of the right image transformed to the left one by homographies estimated from minimal samples.

Example

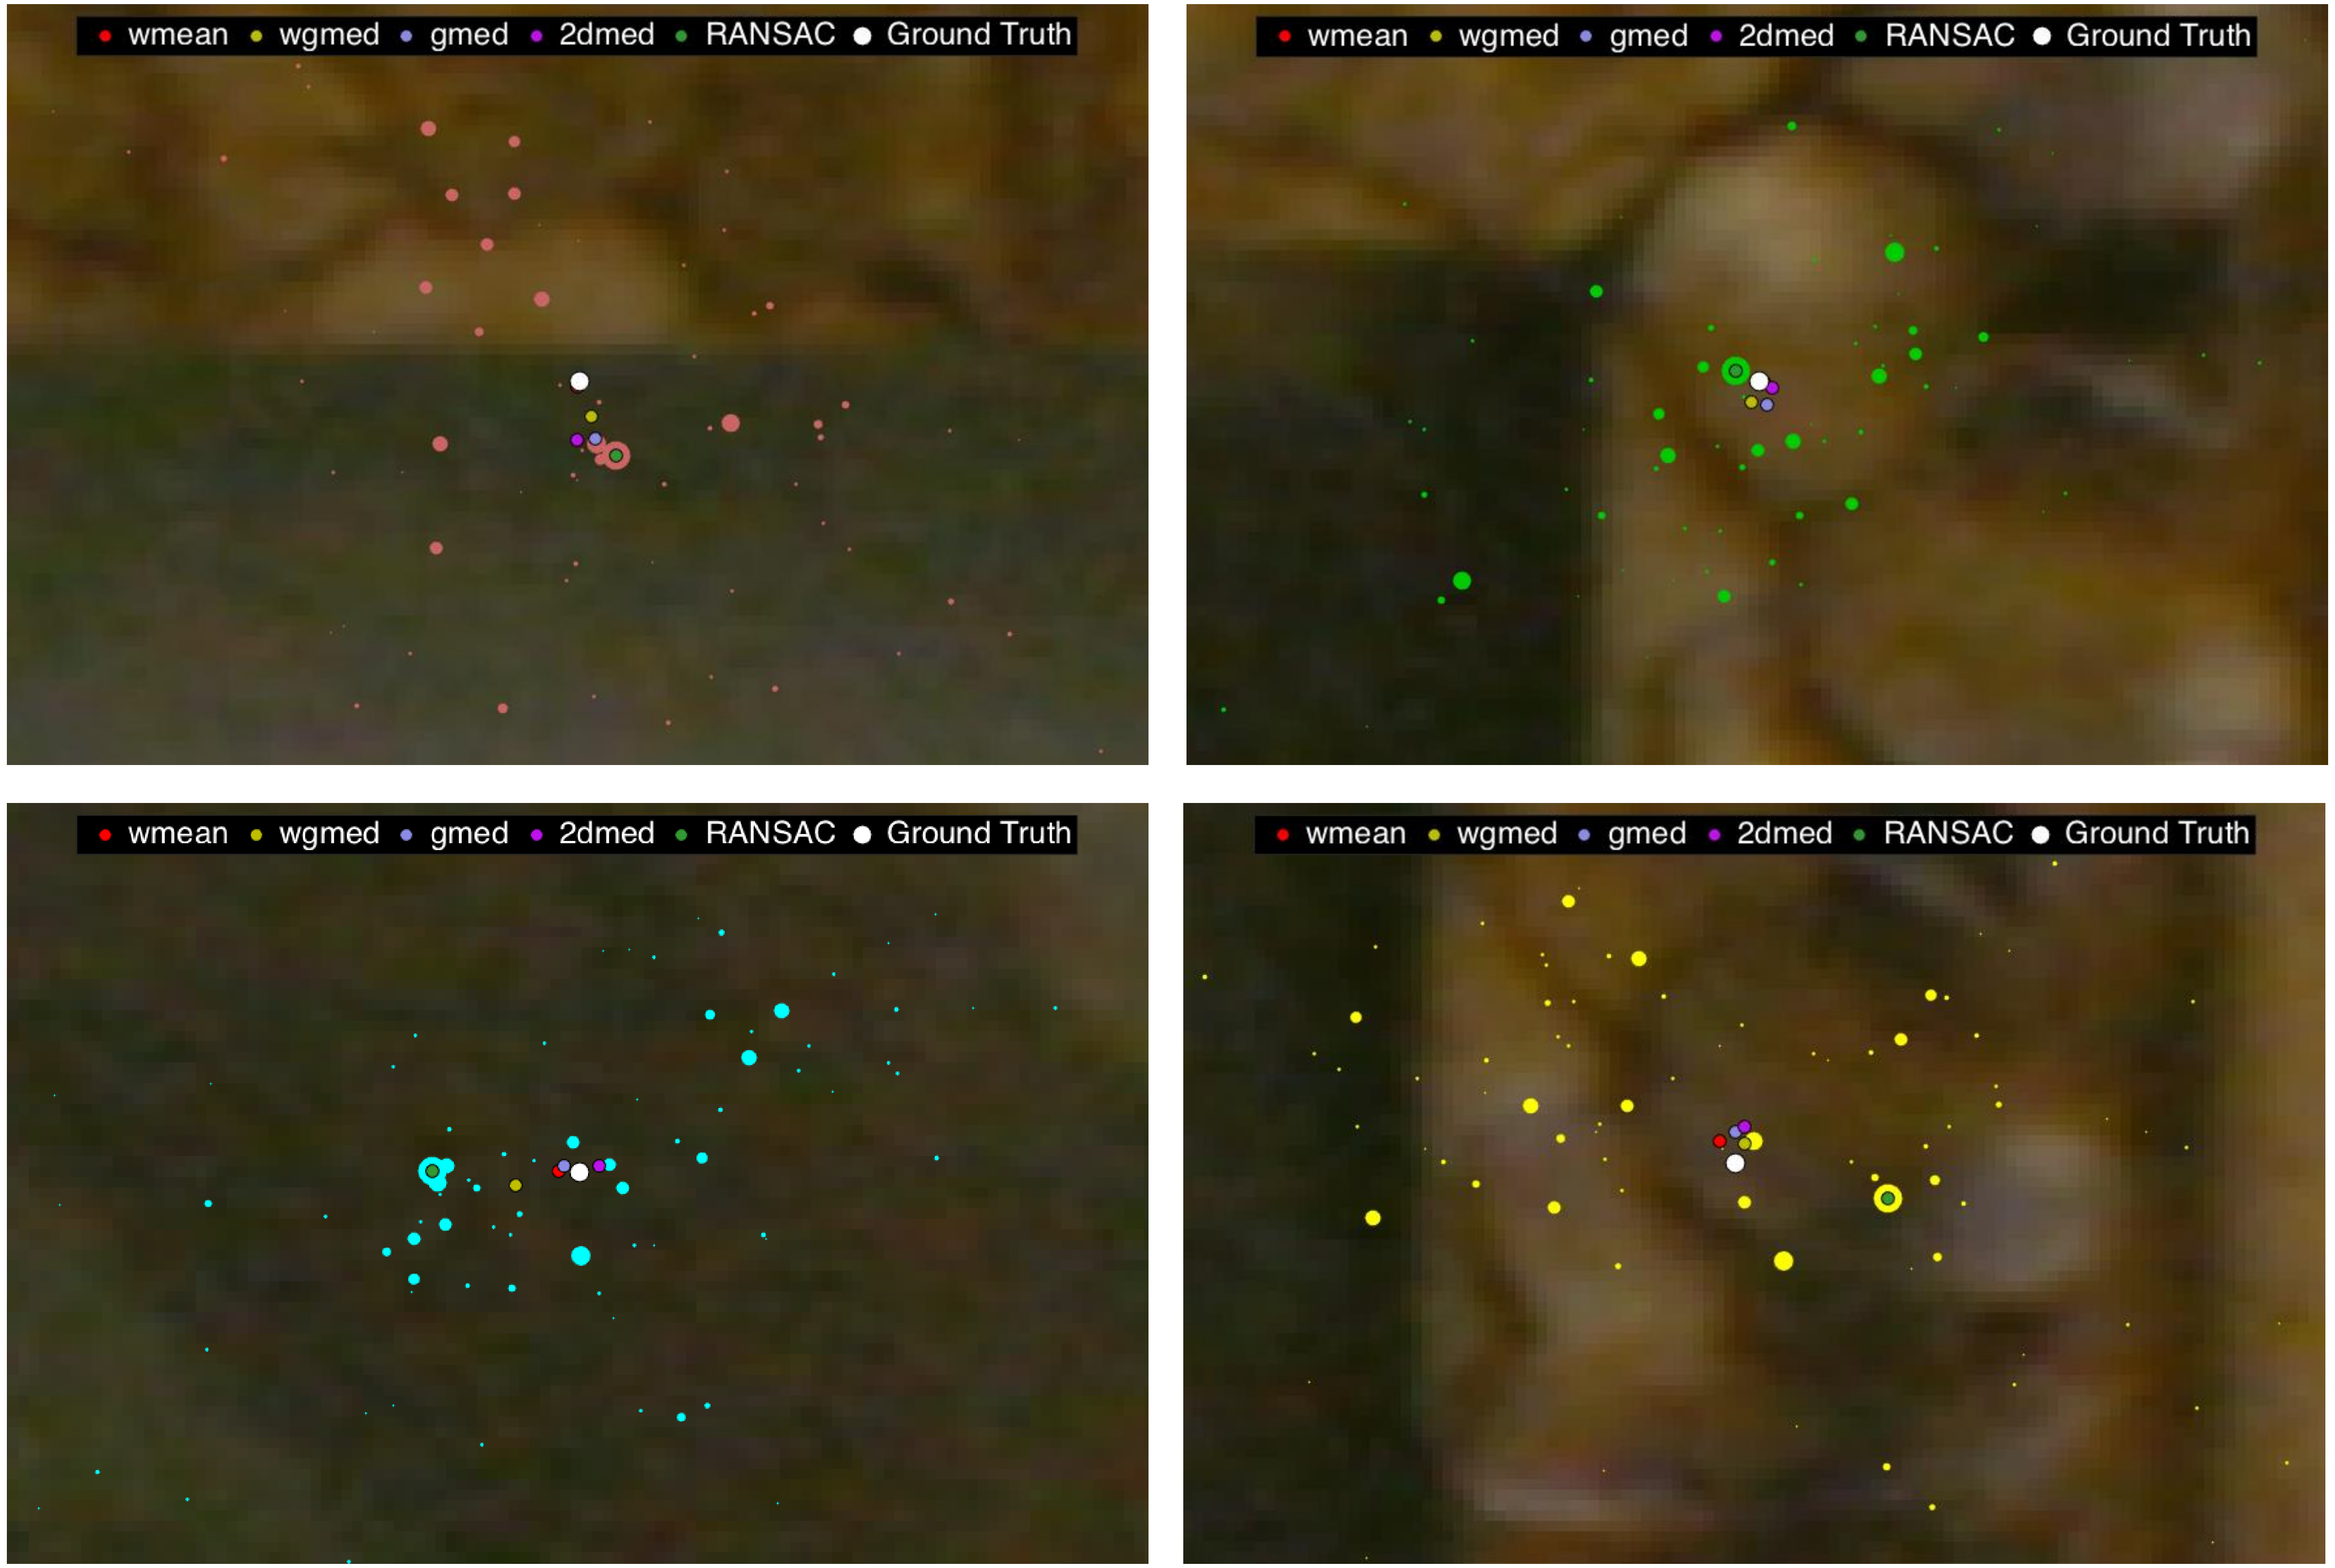

The projected corners after zooming into the image.

The results of different methods (e.g., weighted mean, RANSAC) for calculating the final projected corners are shown.

What happens if there is an outlier?

Example

The projected corners after zooming into the image.

The results of different methods (e.g., weighted mean, RANSAC) for calculating the final projected corners are shown.

What happens if there is an outlier?

Limitations

1. To estimate a model, we need to know
   - *How to convert a model to a unique set of points.*
     For a homography, the corners of the image might suit.
     For a fundamental matrix…?
   - *How to convert points to a model.*
     This is clear when estimating from point correspondences.
     What happens if we have corresponding lines?
2. In theory, *it requires more than a single „all-inlier sample"* to be drawn to find the sought model parameters.
3. Still, the method needs an inlier-outlier threshold to calculate the weights in the weighted averaging.

Limitations

4. Is the model-to-point conversion robust?

   Not really… it has a number degeneracies.

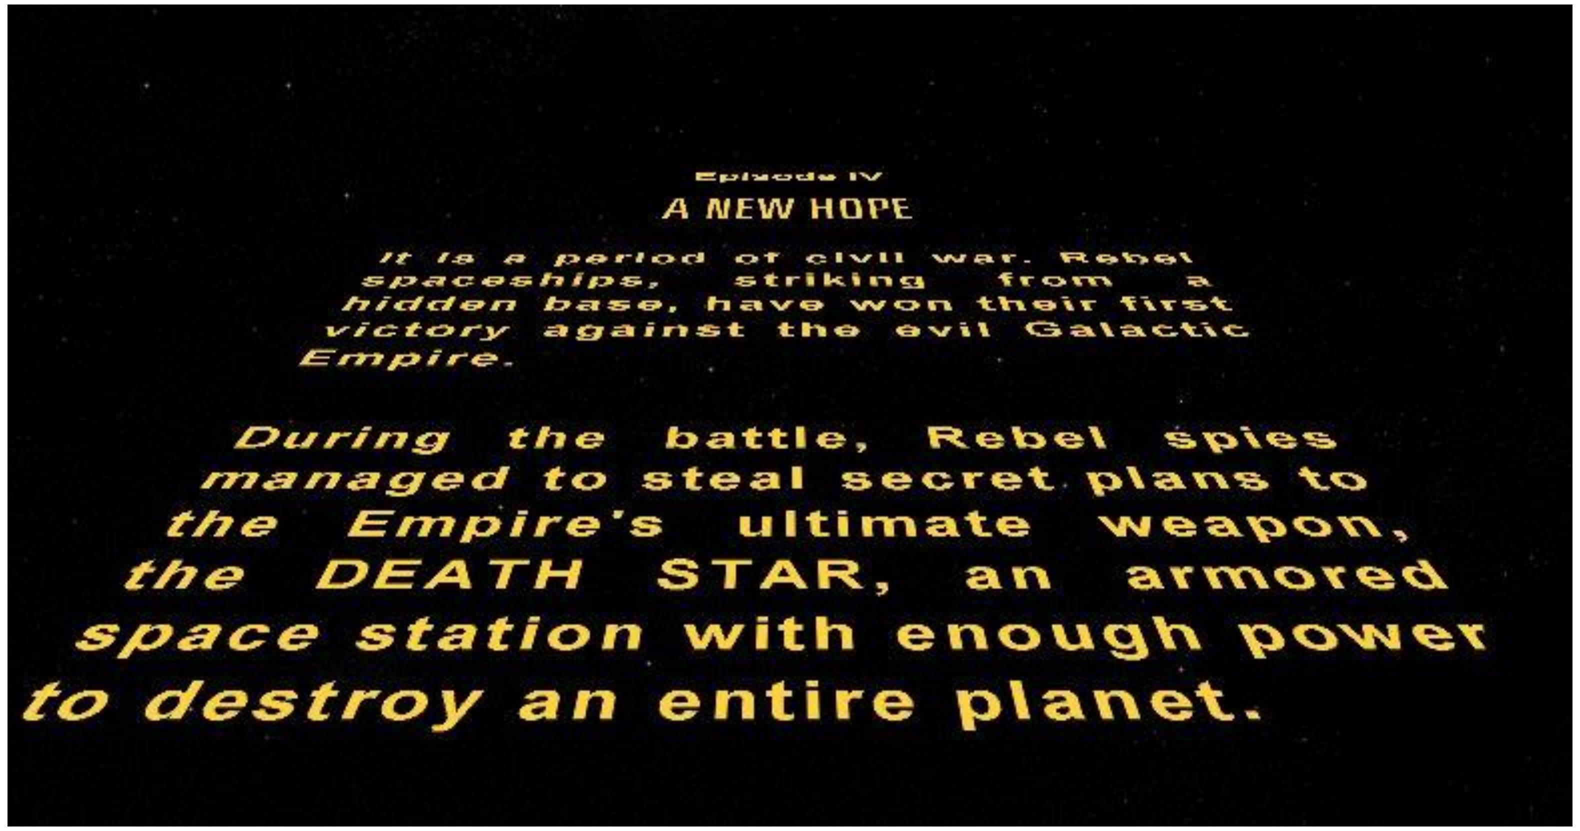

Episode IV

A NEW HOPE

It is a period of civil war. Rebel spaceships, striking from a hidden base, have won their first victory against the evil Galactic Empire.

During the battle, Rebel spies managed to steal secret plans to the Empire's ultimate weapon, the DEATH STAR, an armored space station with enough power to destroy an entire planet.

Example degeneracy when converting the homography to points and the corners are transformed to a single point at infinity.

Example degeneracy when converting the homography to points and the plane flips.

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.

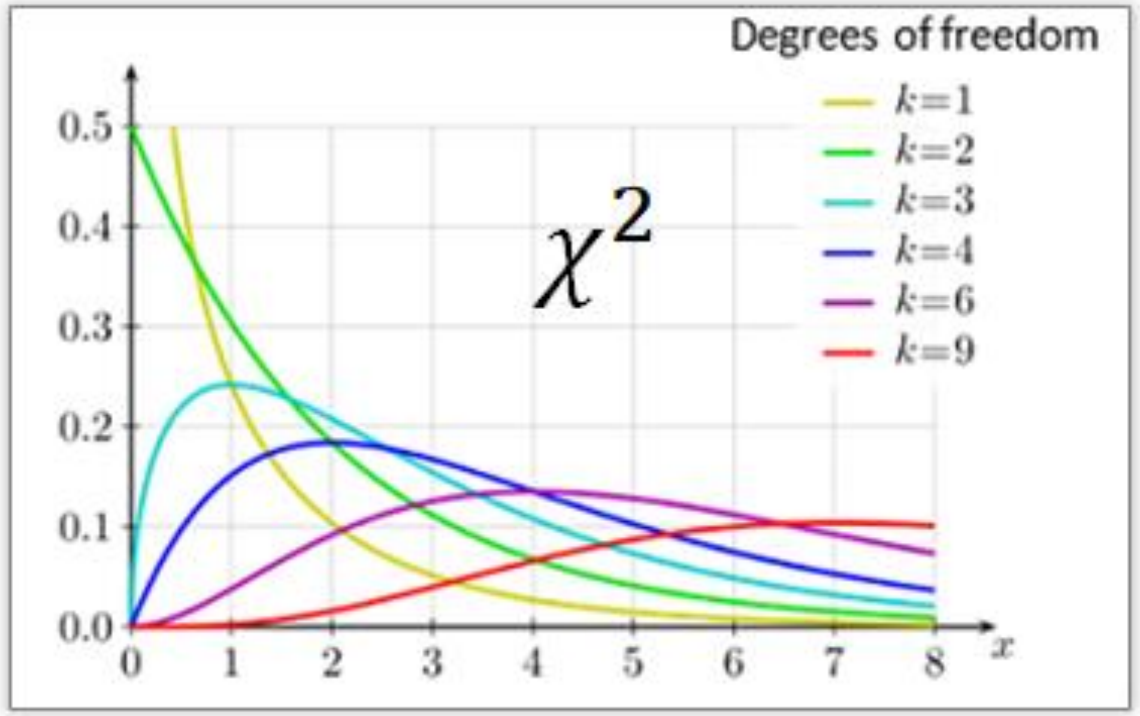

Design choices and data interpretation:

- 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 Likelihood Given a Noise Scale

$$L(\underset{\substack{\uparrow\\\text{Model}\\\text{parameters}}}{\theta} \mid \overset{\substack{\text{Noise}\\\text{level}\\\downarrow}}{\sigma}) = \underbrace{\frac{1}{l^{|\mathcal{X}|-|I(\sigma)|}}}_{\substack{\text{Comes from the}\\\text{outliers' distribution}}} \prod_{\underset{\substack{\uparrow\\\text{Set of inliers}\\\text{which }\sigma\text{ implies}}}{x\in I(\sigma)}} \underbrace{\left[2\overset{\substack{\text{A constant}\\\text{of the}\\\text{distribution}\\\downarrow}}{C}(p)\sigma^{-p} D^{\overset{\substack{\text{Problem DoF}\\\downarrow}}{p}-1}(\theta,x)\exp\frac{-\overset{\substack{\text{Distance function}\\\downarrow}}{D^2}(\theta,x)}{2\sigma^2}\right]}_{\substack{\text{Comes from the inliers'}\\\text{distribution}}}$$

Note: Given noise scale σ, the likelihood of model θ is $L(\theta \mid \sigma)$.

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. -> Final least-squares fitting is not applicable.
- No knowledge about the inlier number. -> 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

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

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:

$$\text{RANSAC iteration number} \longrightarrow k(\overset{\text{Model parameters}}{\theta}, \underset{\text{Noise scale}}{\sigma}, \overset{\text{Point set}}{\mathcal{X}}) = \frac{\ln(1-\overset{\text{Manually set confidence}}{\mu})}{\ln\left(1-\left(\underset{\text{Inlier ratio}}{\frac{|I(\theta,\sigma)|}{|\mathcal{X}|}}\right)^{m}\right)} \longleftarrow \text{Sample size}$$

Criterion via marginalization:

$$\text{MAGSAC iteration number} \longrightarrow k^*(\theta,\mathcal{X}) = \frac{1}{\underset{\text{Noise scale upper bound}}{\sigma_{max}}}\int_0^{\sigma_{max}} \underbrace{k(\theta,\sigma,\mathcal{X})}_{\text{RANSAC iteration number}}\ d\sigma$$

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.

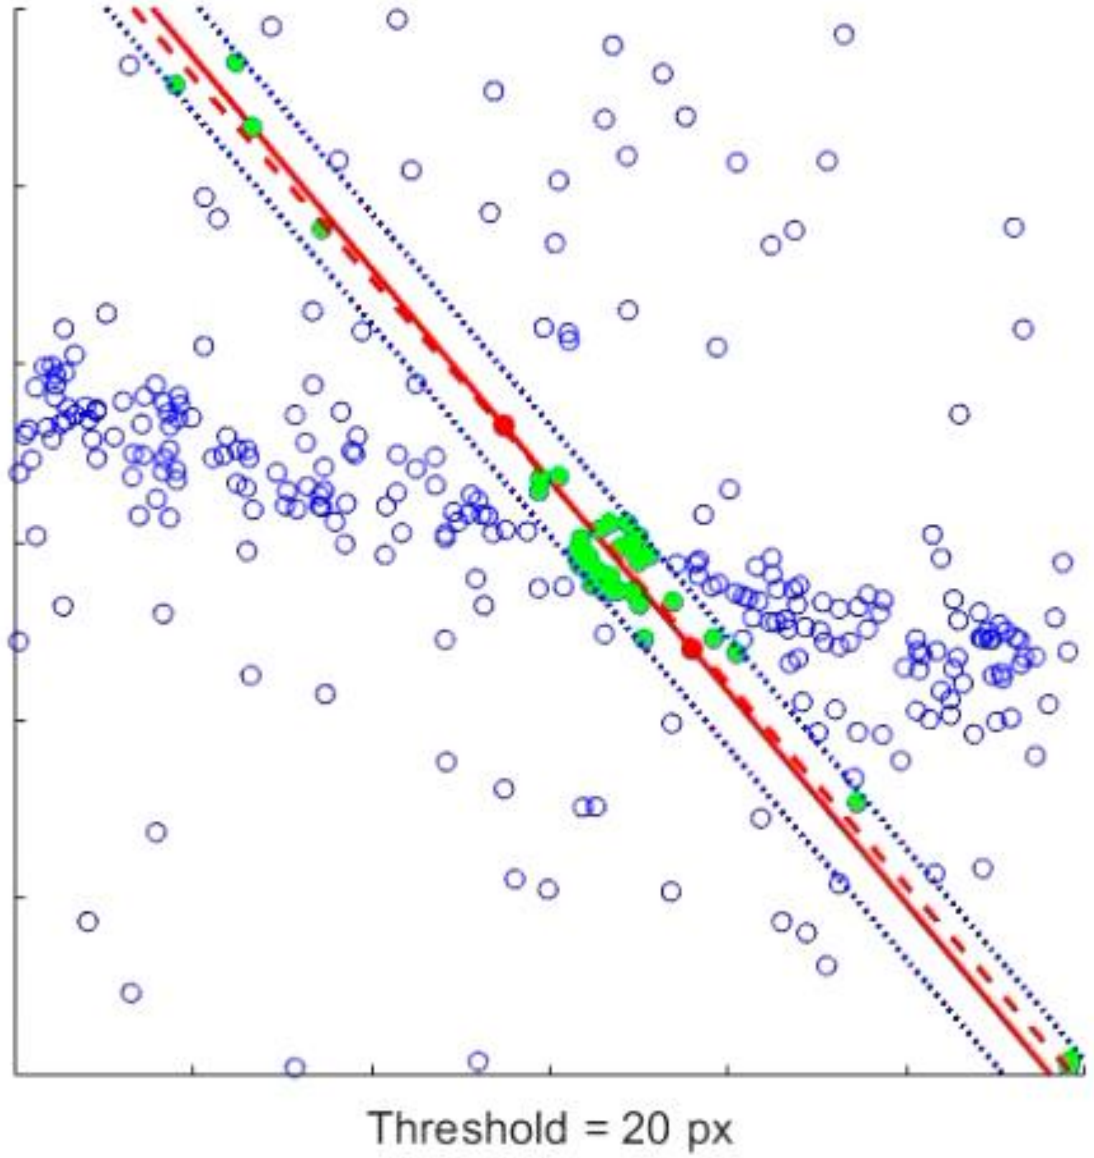

Example. The minimal sample (red dots), the line which it initializes (red line), the current threshold (blue dotted lines), the implied inliers (green dots), and the model fit to the implied inliers (red dashed line).

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 assuming other distributions as well.

The MAGSAC++ Approach

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

The MAGSAC++ Approach

Model quality function:

The quality is defined by an M-estimator built on the proposed weights $w(r)$ and is as follows:

$$Q^*_{++}(\theta, P) = \frac{1}{L(\theta, P)} = \frac{1}{\sum_{p \in P} \rho(r(\theta, p))}$$

where $L(\theta, P)$ is a loss function of the M-estimator and ρ is

$$\rho(r) = \int_0^r x w(x)\, dx \,.$$

$Q^*_{++}(\theta)$ can be calculated

precisely (MAGSAC approximated it), and **efficiently** using a look-up table.

Conclusions of this Section

- There seems to have an emerging trend of methods not requiring a single inlier-outlier threshold.
- There are a number of approaches to skip setting the threshold manually:
  (i) skipping entirely the inlier selection, by using only minimal sample models;
  (ii) selecting the threshold adaptively;
  (iii) marginalizing over a range of possible thresholds.

Questions before continuing?

Note: experiments will be shown in the last talk by Dmytro Mishkin.

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

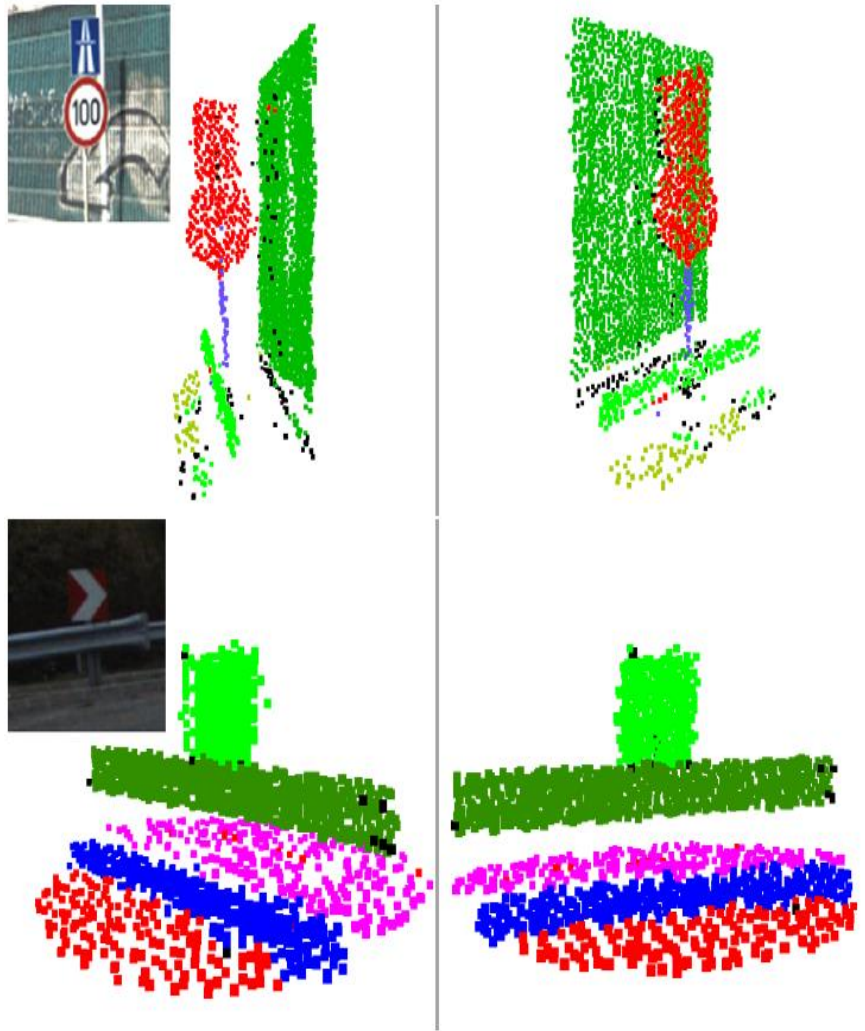

Planes in LiDAR data.

Vanishing points (similar line directions)

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

Daniel Barath and Jiri Matas, Graph-Cut RANSAC; CVPR 2018

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.

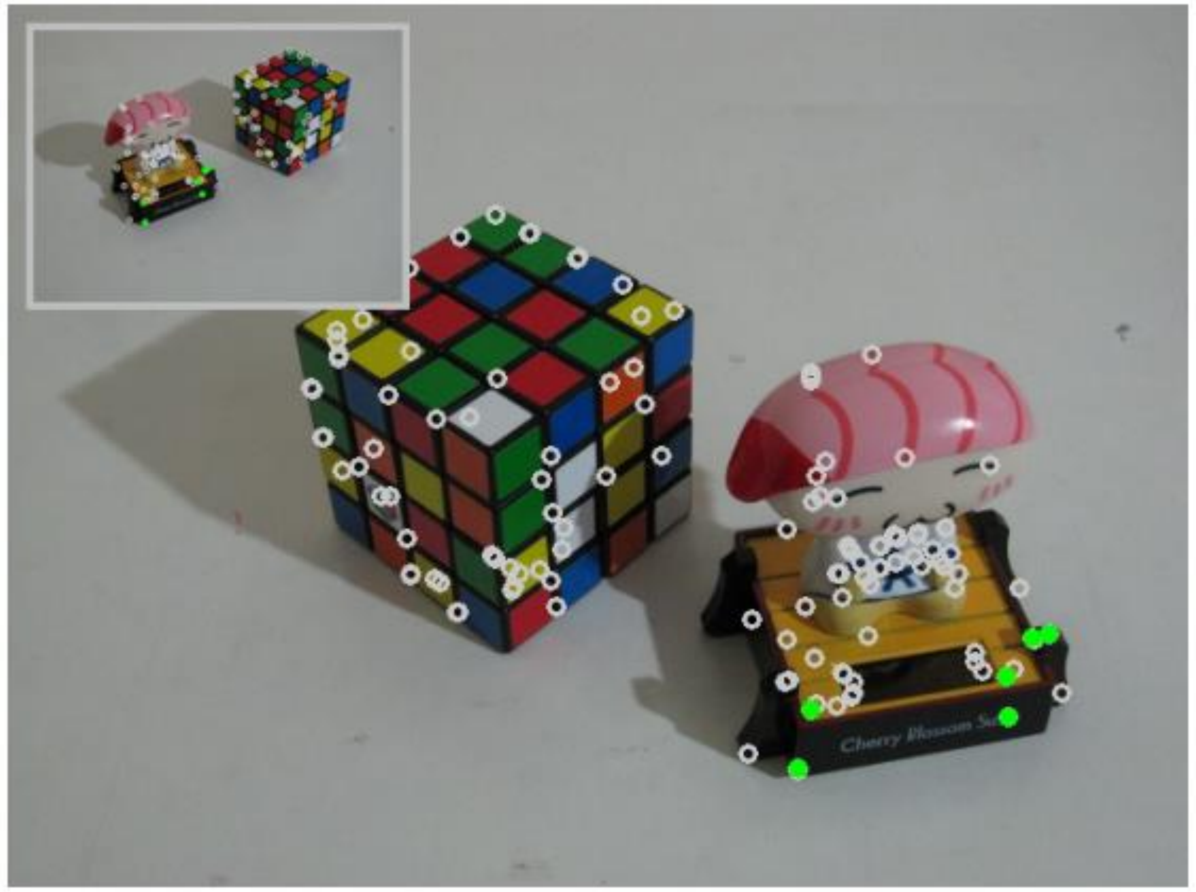

Minimal sample initializing a rigid motion.

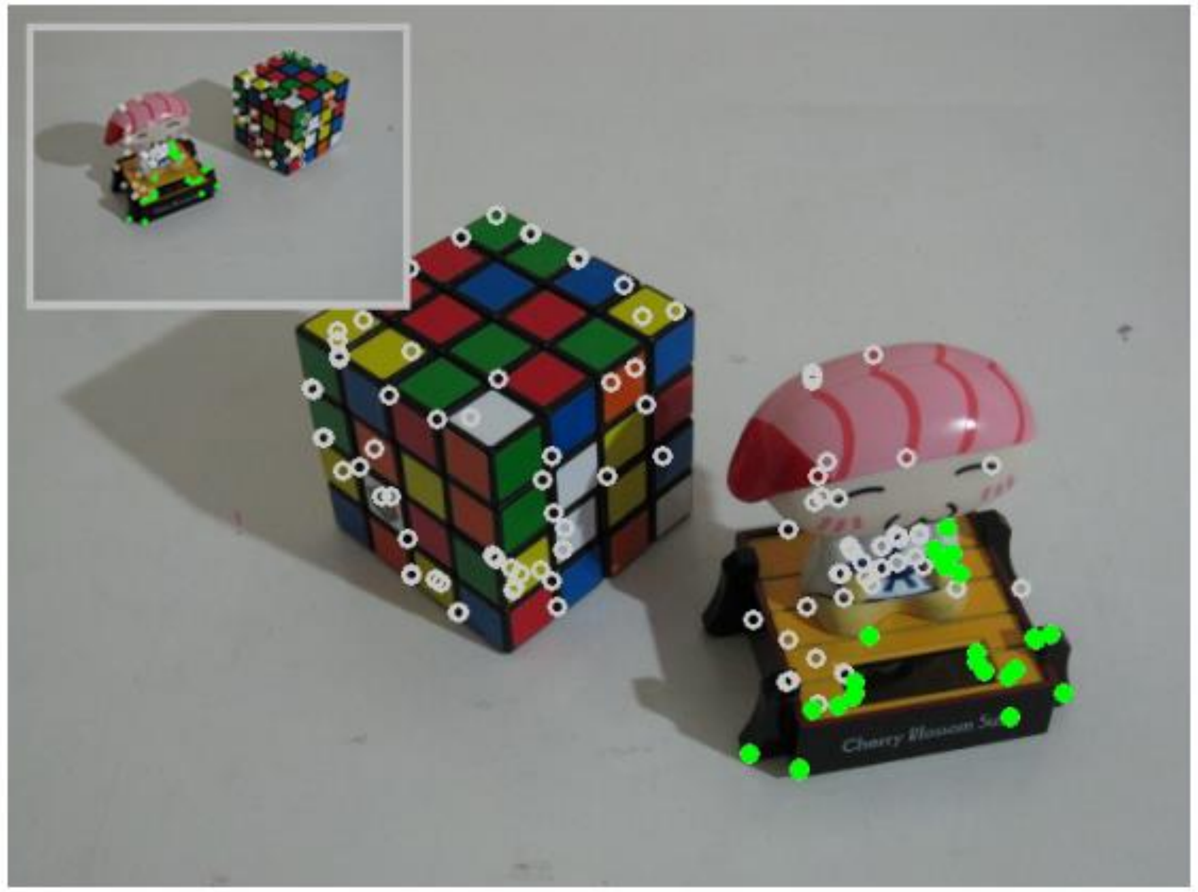

Inliers by standard RANSAC-like thresholding.

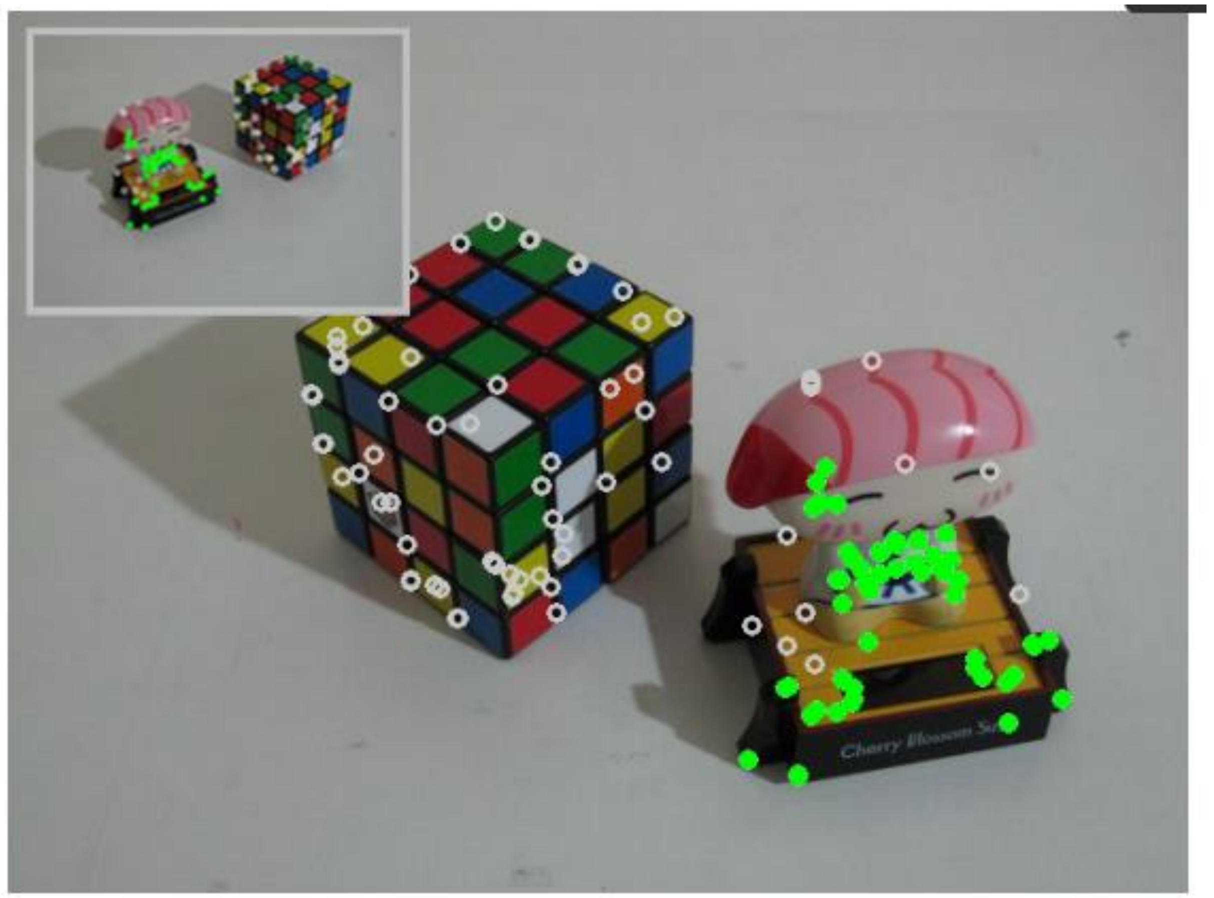

Inliers by considering spatial coherence.

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

```
$iter \leftarrow 0,\ J^* \leftarrow \infty$
repeat
    Select random $S \subseteq \mathcal{X}$ (sample size $m = |S|$)
    Estimate parameters $\theta = e(S)$
    Evaluate $J(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}, \theta)$
    If $J(\theta) < J^*$ then
        $\theta^* \leftarrow$ LocalOptimization$(\mathcal{X}_{in}, \theta)$, $J^* \leftarrow J(\theta)$
    $iter \leftarrow iter + 1$
until $P$(better solution exists) $= f(|\mathcal{X}|, J^*, iter) < \eta$
```

SAMPLING
MODEL ESTIMATION
VERIFICATION
SO-FAR-THE-BEST
LOCAL OPTIMIZATION

RANSAC as a Labeling Problem

Idea: formalize RANSAC 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*.

RANSAC 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. ⟶

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

„Close" and labeled inlier

„Far" and labeled outlier

RANSAC as a Labeling Problem

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.
- Advantage: new energy terms can be added.
- Using a different loss function than RANSAC's is easy and beneficial.

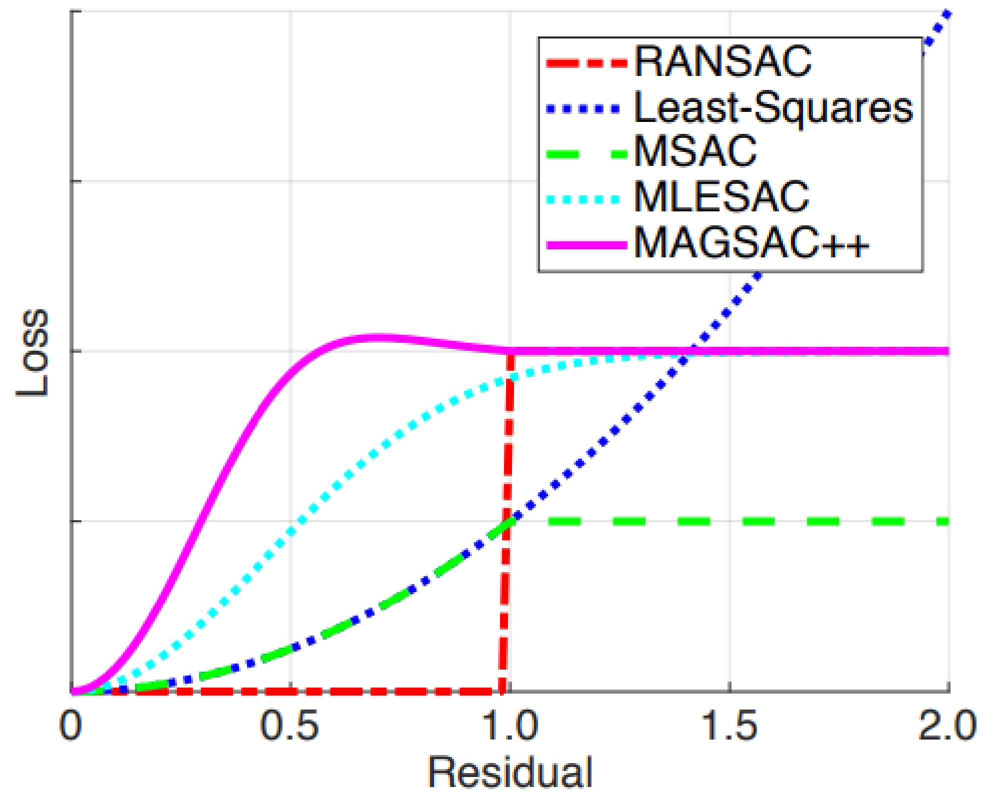

Spatial Coherence in RANSAC (Potts model)

Modelling the spatial coherence by the Potts model.

Binary energy →

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

p

q

Edge set of $\mathcal{G}$

Neighborhood graph $\mathcal{G}$

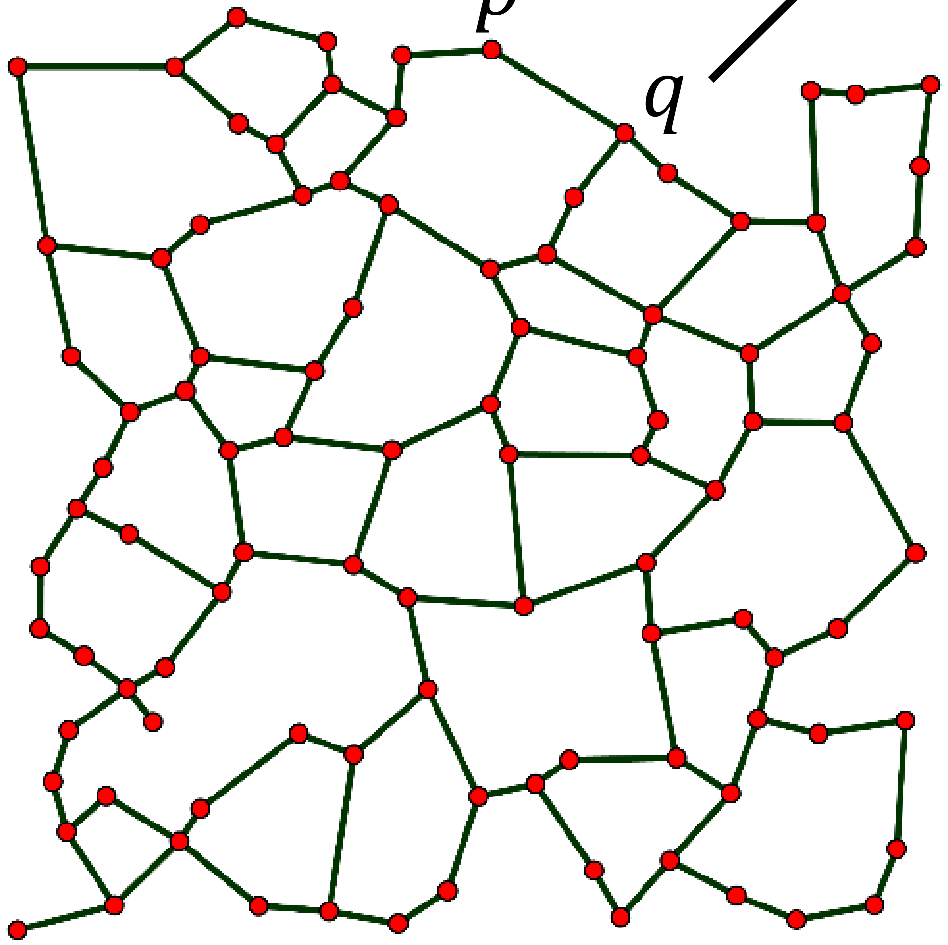

In short, if neighboring points

- 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 RANSAC (Potts model)

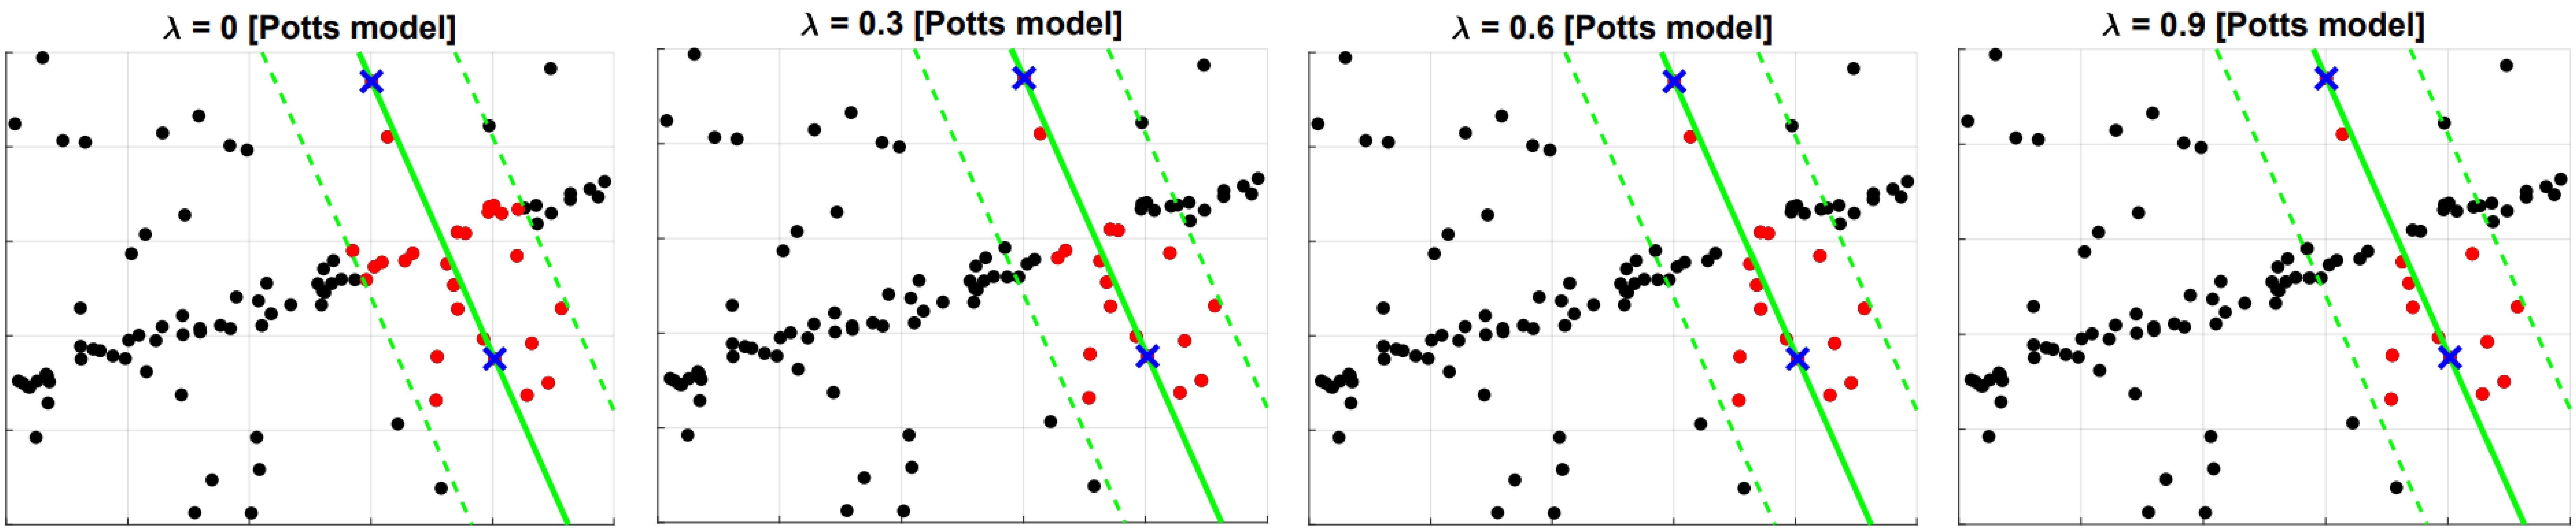

Example labeling using the Potts model to interpret spatial coherence.
Parameter $\lambda \in [0, 1]$ is the weight of the term.

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)

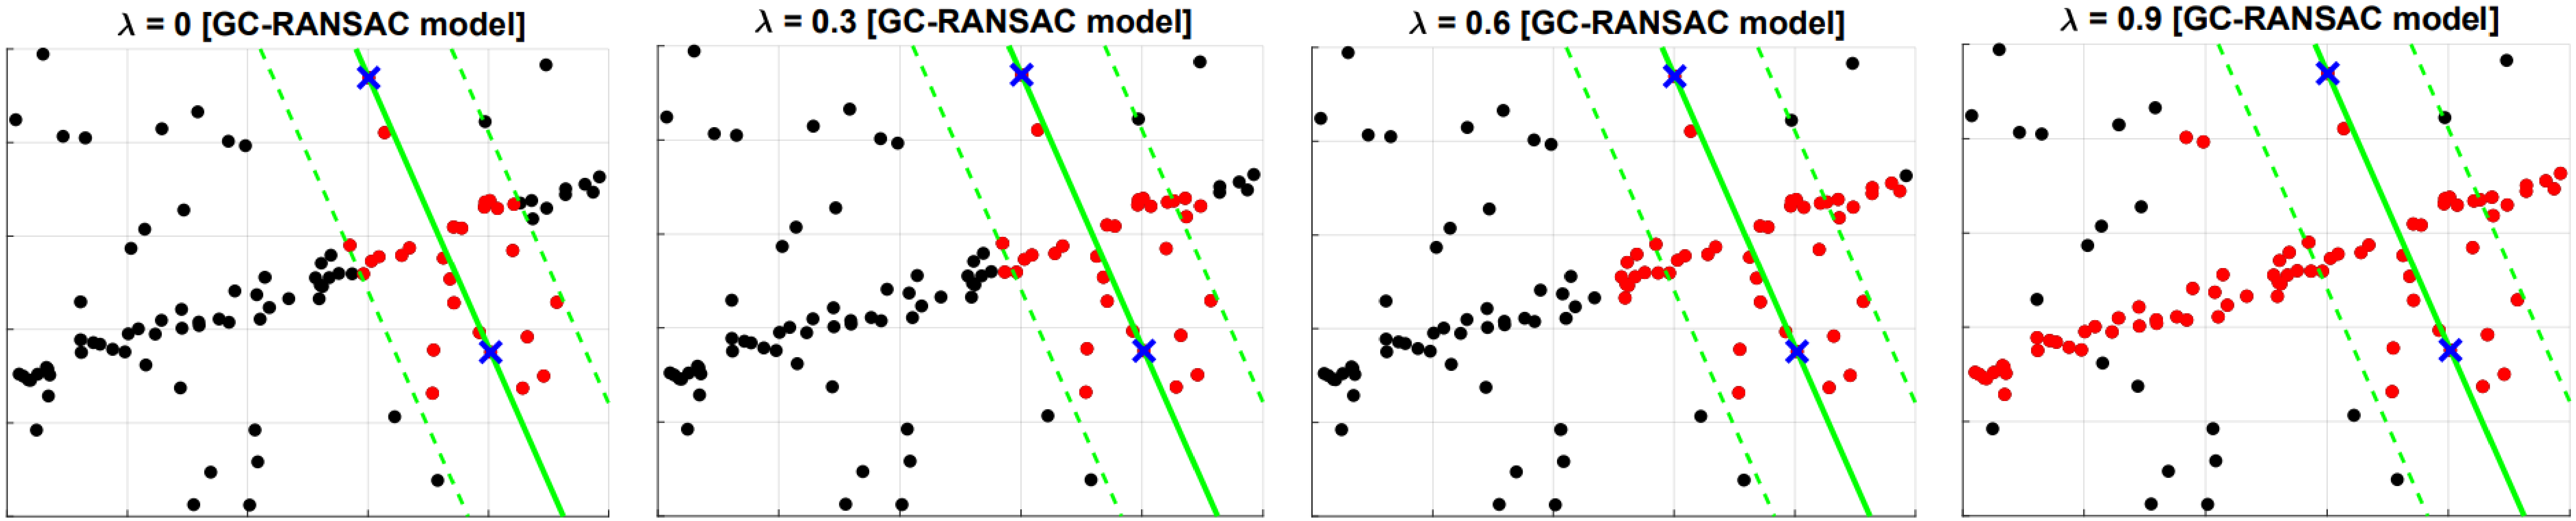

Example labeling using the GC model to interpret spatial coherence.
Parameter $\lambda \in [0, 1]$ is the weight of the term.

Notes:

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

GC-RANSAC Local Optimization

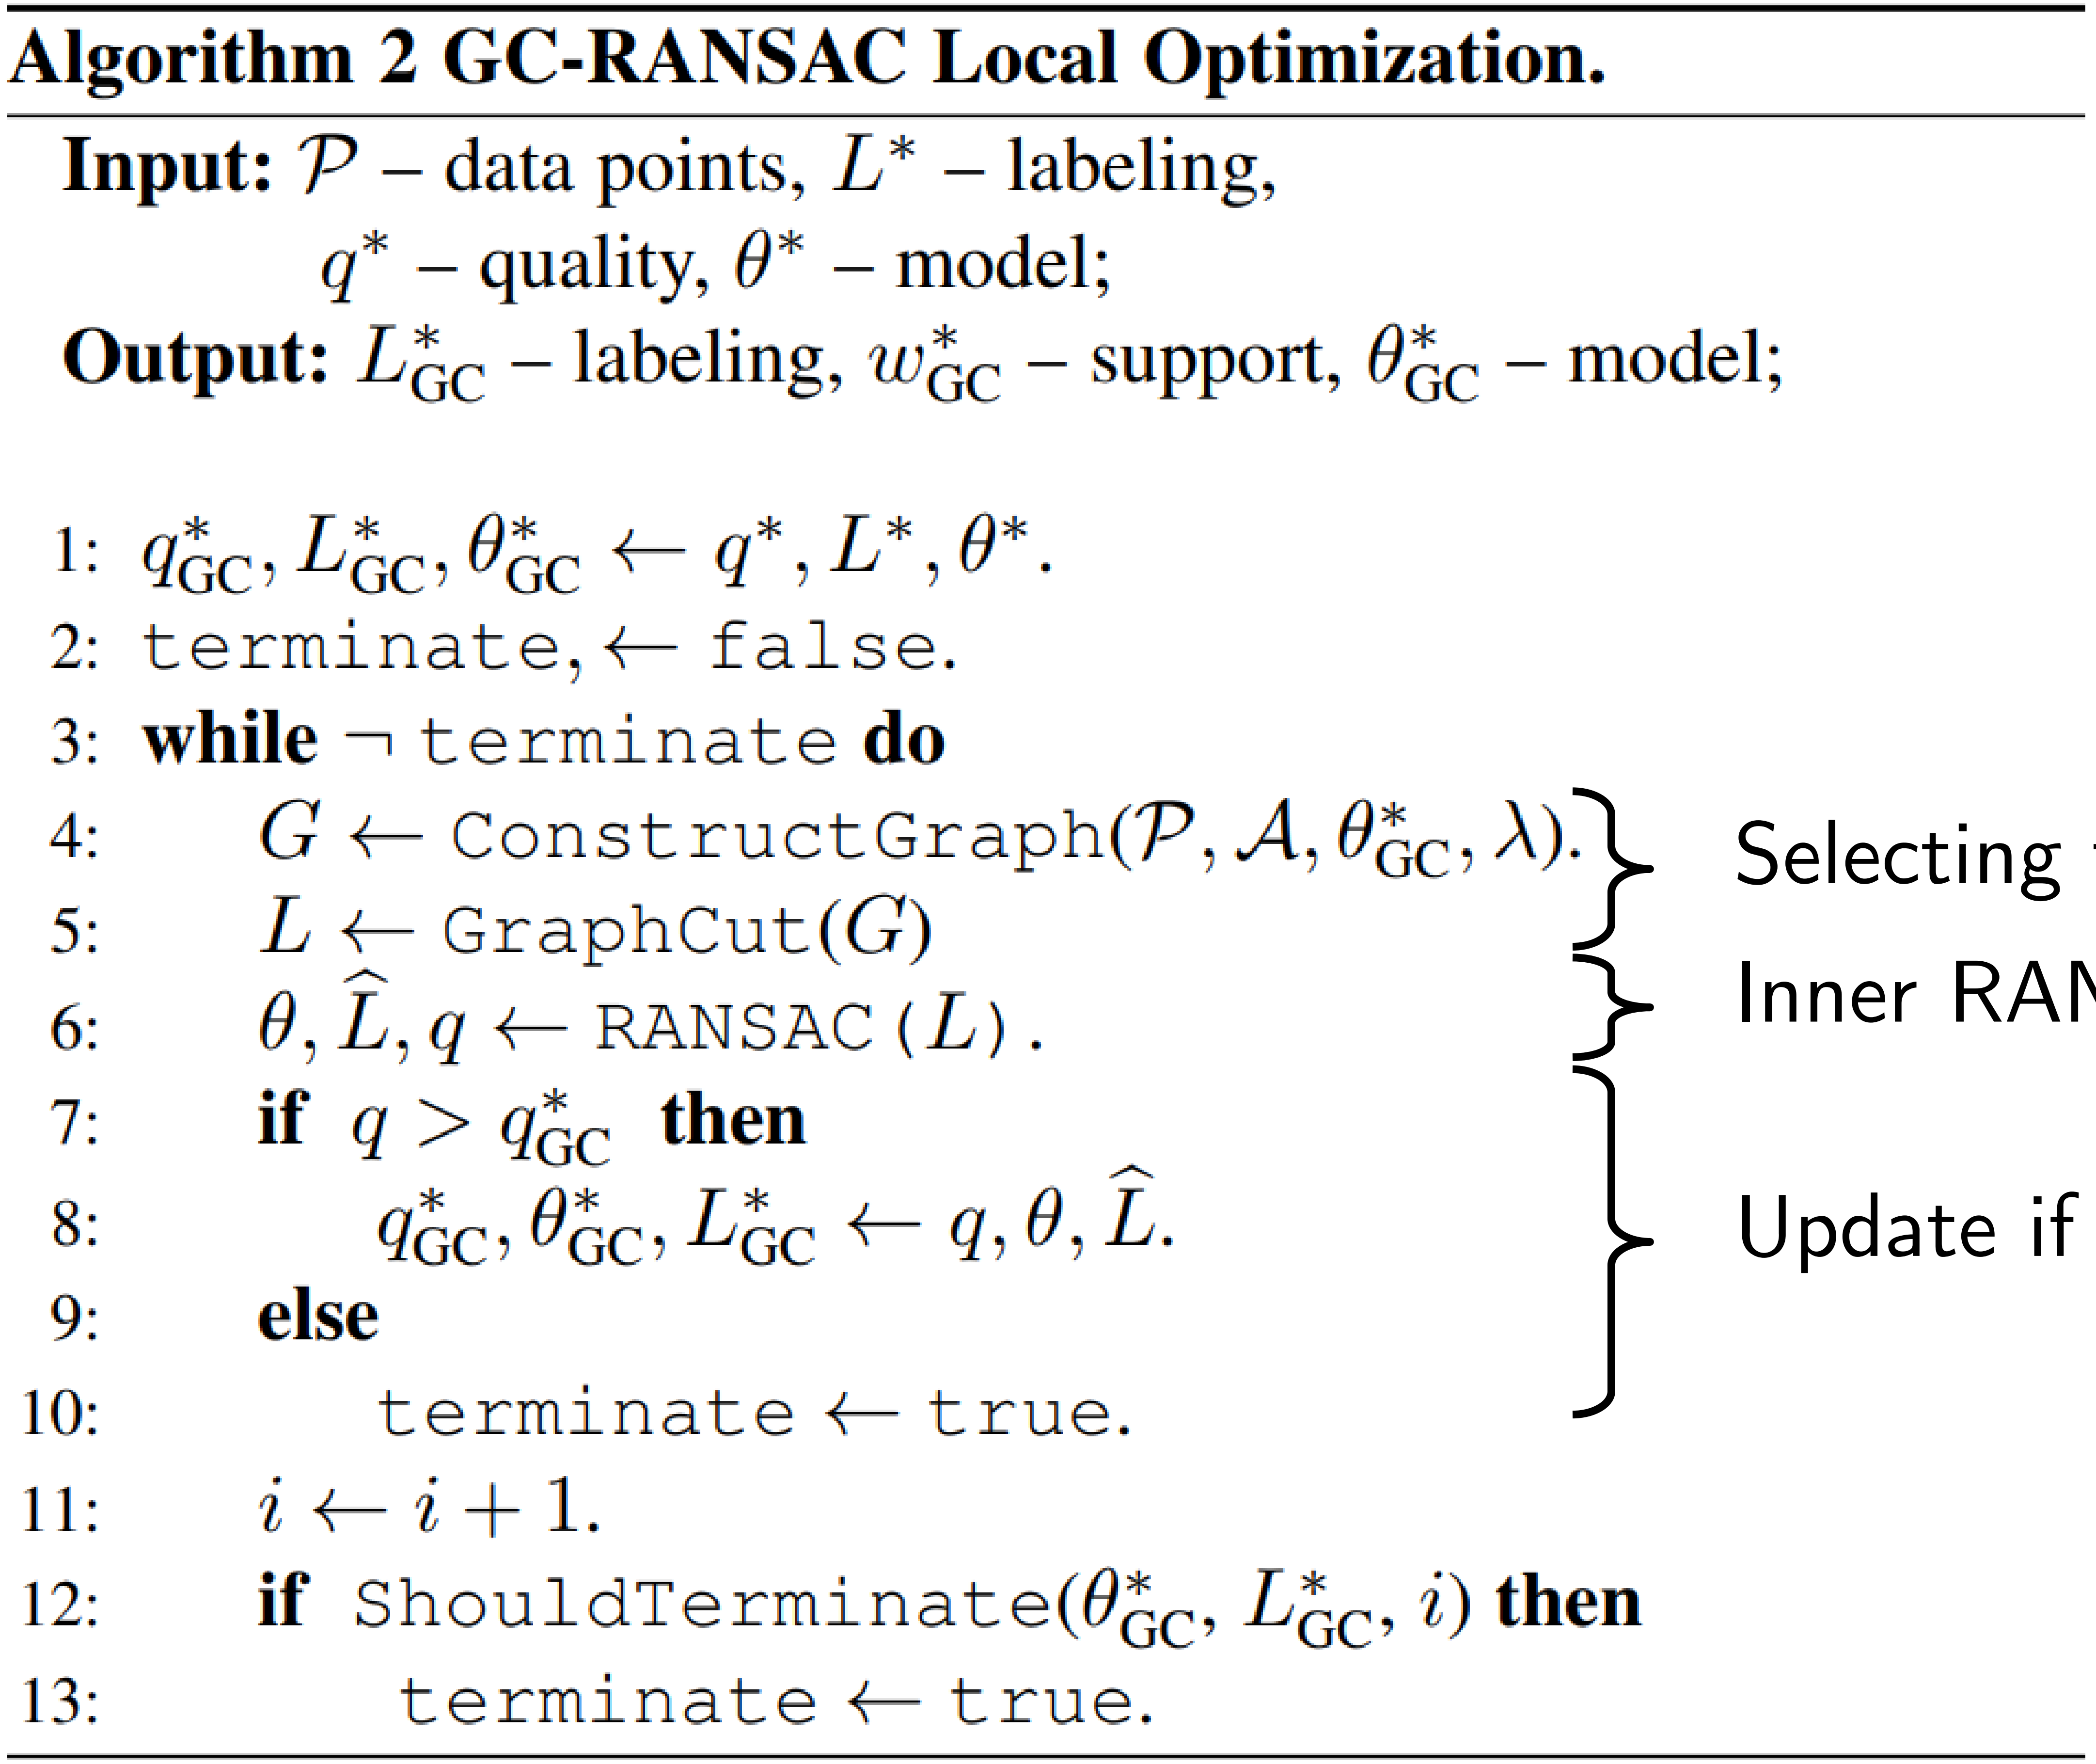

Algorithm 2 GC-RANSAC Local Optimization.

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

```
$q^*_{\text{GC}}, L^*_{\text{GC}}, \theta^*_{\text{GC}} \leftarrow q^*, L^*, \theta^*$.
terminate, ← false.
while ¬ terminate do
    $G \leftarrow$ ConstructGraph($\mathcal{P}, \mathcal{A}, \theta^*_{\text{GC}}, \lambda$).
    $L \leftarrow$ GraphCut($G$)
    $\theta, \widehat{L}, q \leftarrow$ RANSAC($L$).
    if $q > q^*_{\text{GC}}$ then
        $q^*_{\text{GC}}, \theta^*_{\text{GC}}, L^*_{\text{GC}} \leftarrow q, \theta, \widehat{L}$.
    else
        terminate ← true.
    $i \leftarrow i + 1$.
    if ShouldTerminate($\theta^*_{\text{GC}}, L^*_{\text{GC}}, i$) then
        terminate ← true.
```

Selecting the inliers by graph-cut

Inner RANSAC on the spatially coherent inliers

Update if new so-far-the-best

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.

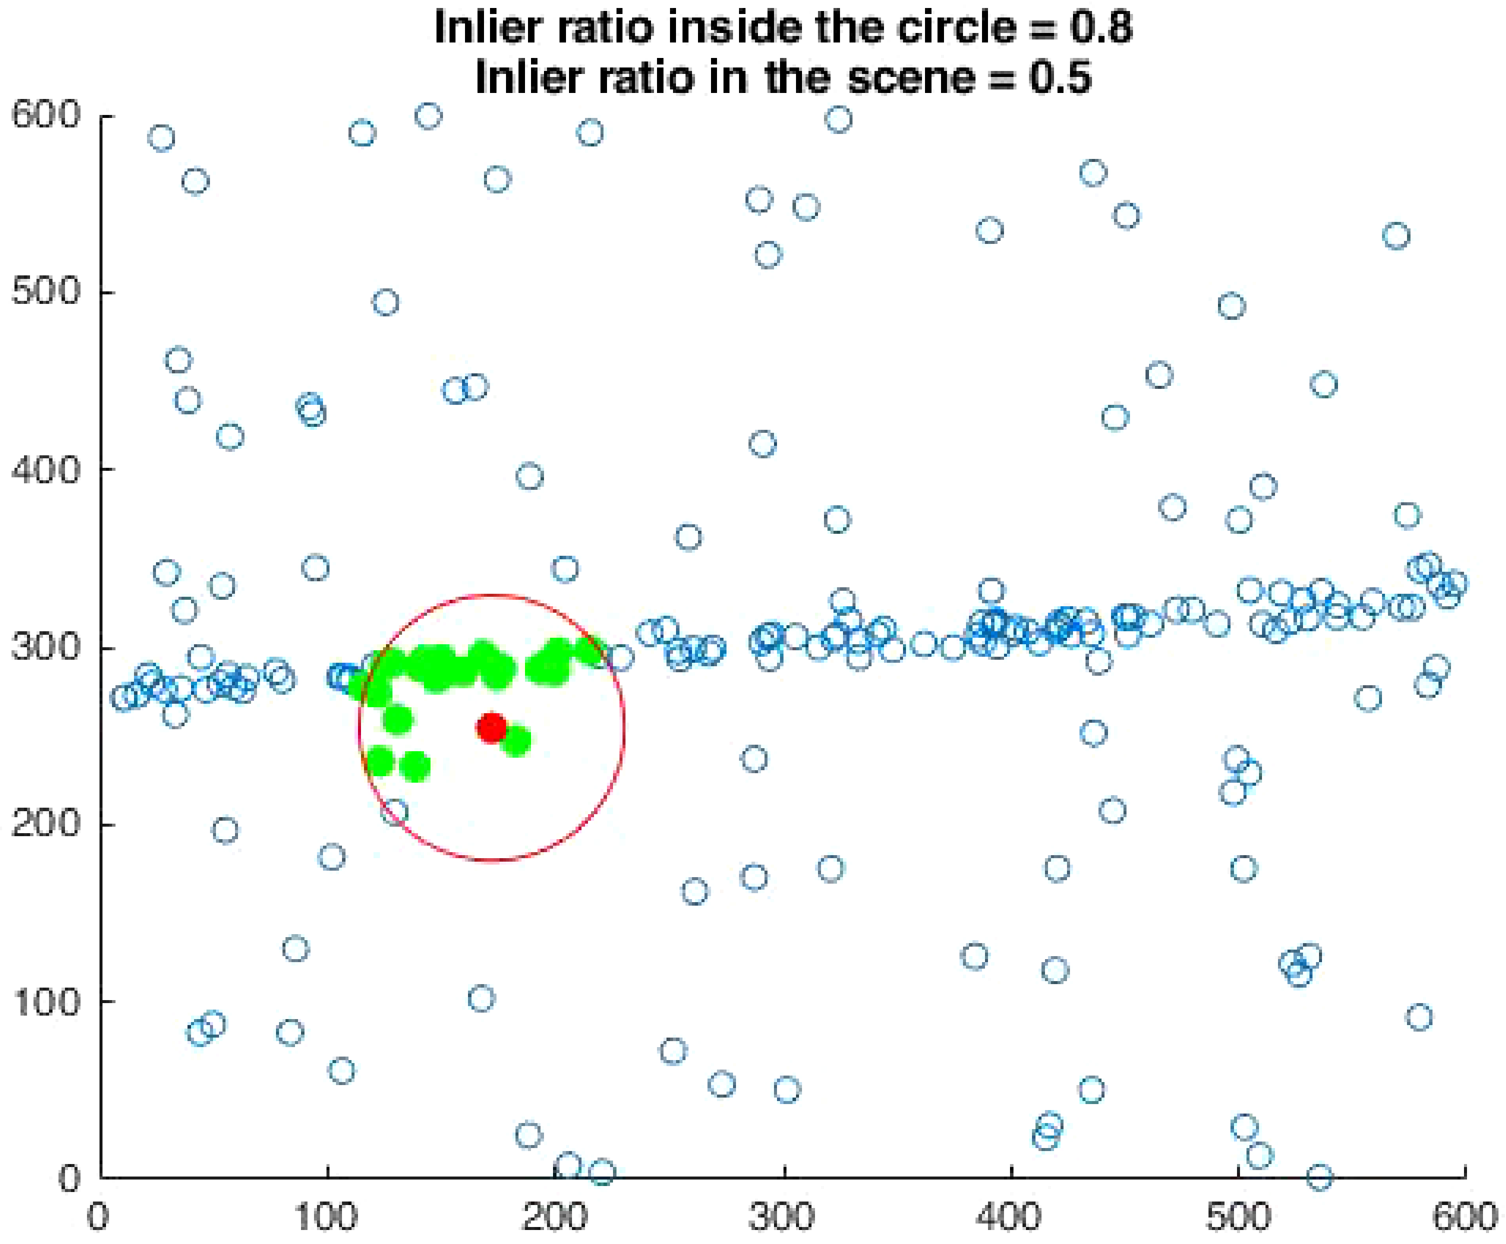

The 1st point (red dot) is selected at random. The points (green) inside the assigned circle (red) are shown.

NAPSAC Sampler

Advantage: high inlier ratio for local samples

Issues:

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

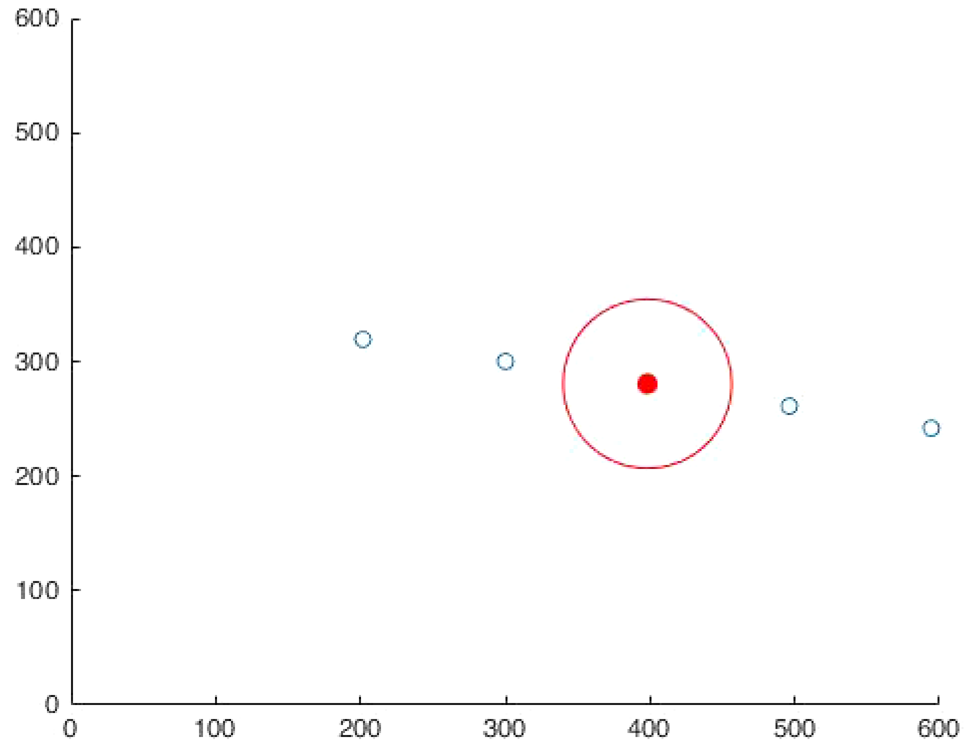

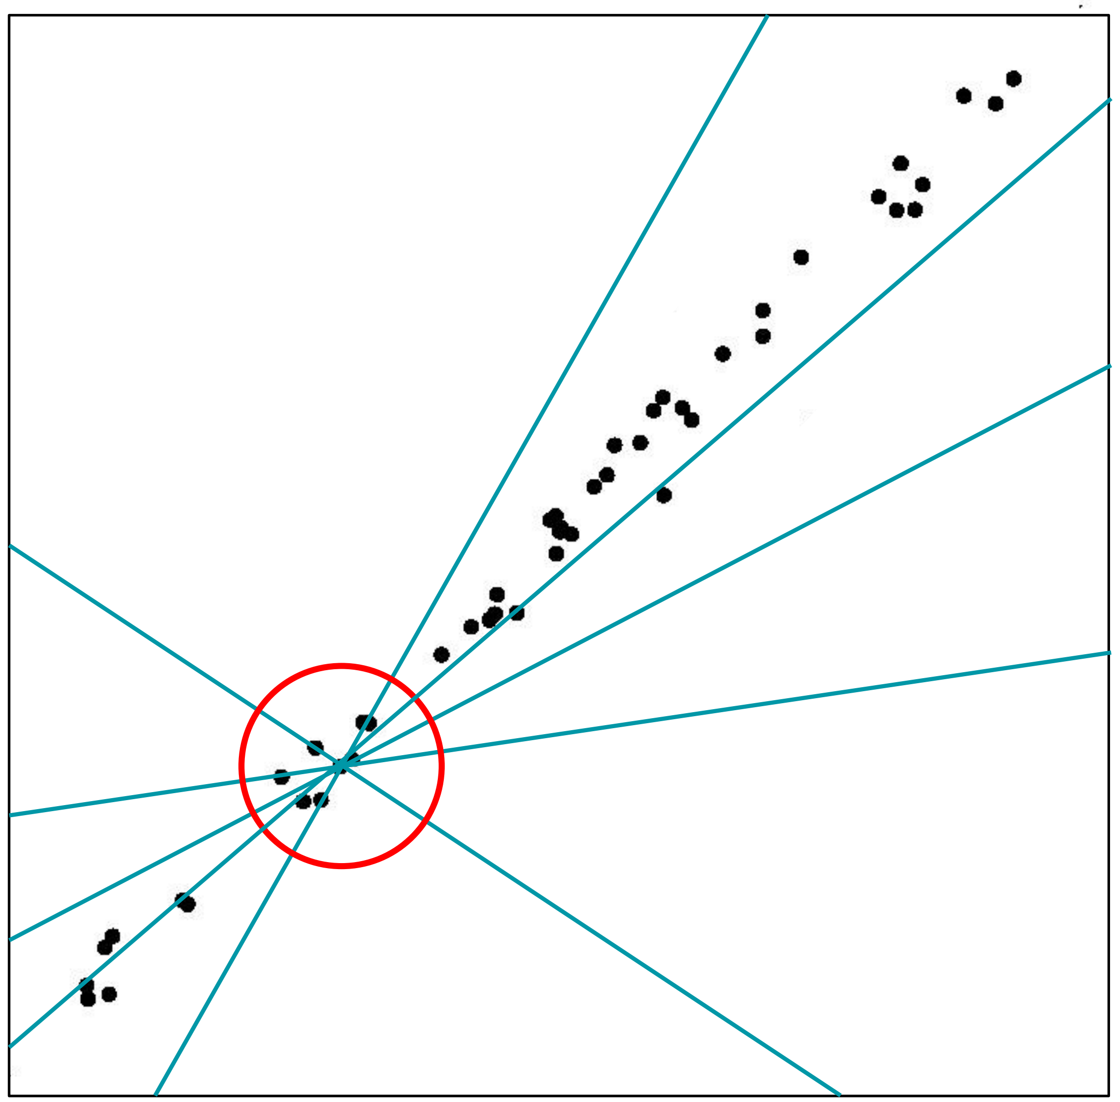

Example line fitting to NAPSAC samples. A red circle is centered on the 1st point. All lines fit using a 2nd point from the circle are inaccurate.

Progressive NAPSAC

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

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 $\boldsymbol{p}$.

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).

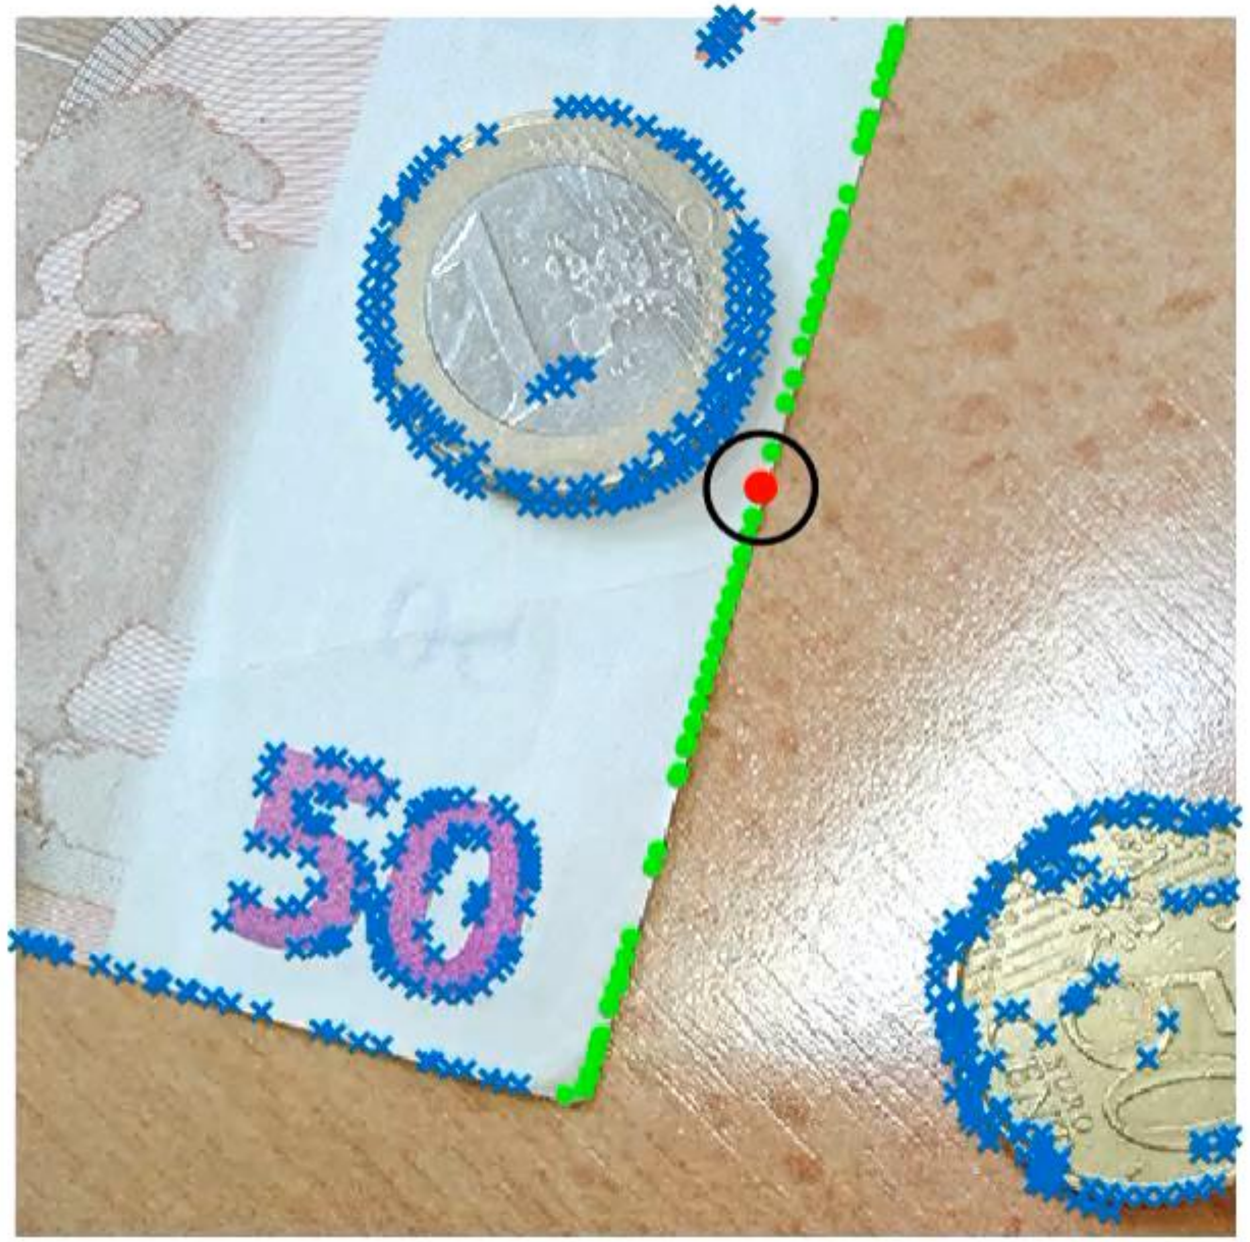

Iteration #1
Inlier ratio inside the circle = 1.00

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.

Conclusion of this Section

Exploiting spatial coherence is beneficial both in terms of

- accuracy (by improving the local optimization)
- processing time (by finding good samples early)

when fitting geometric models since the data often contain spatial structures.

Questions before continuing?

Note: experiments will be shown in the last talk by Dmytro Mishkin.

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)

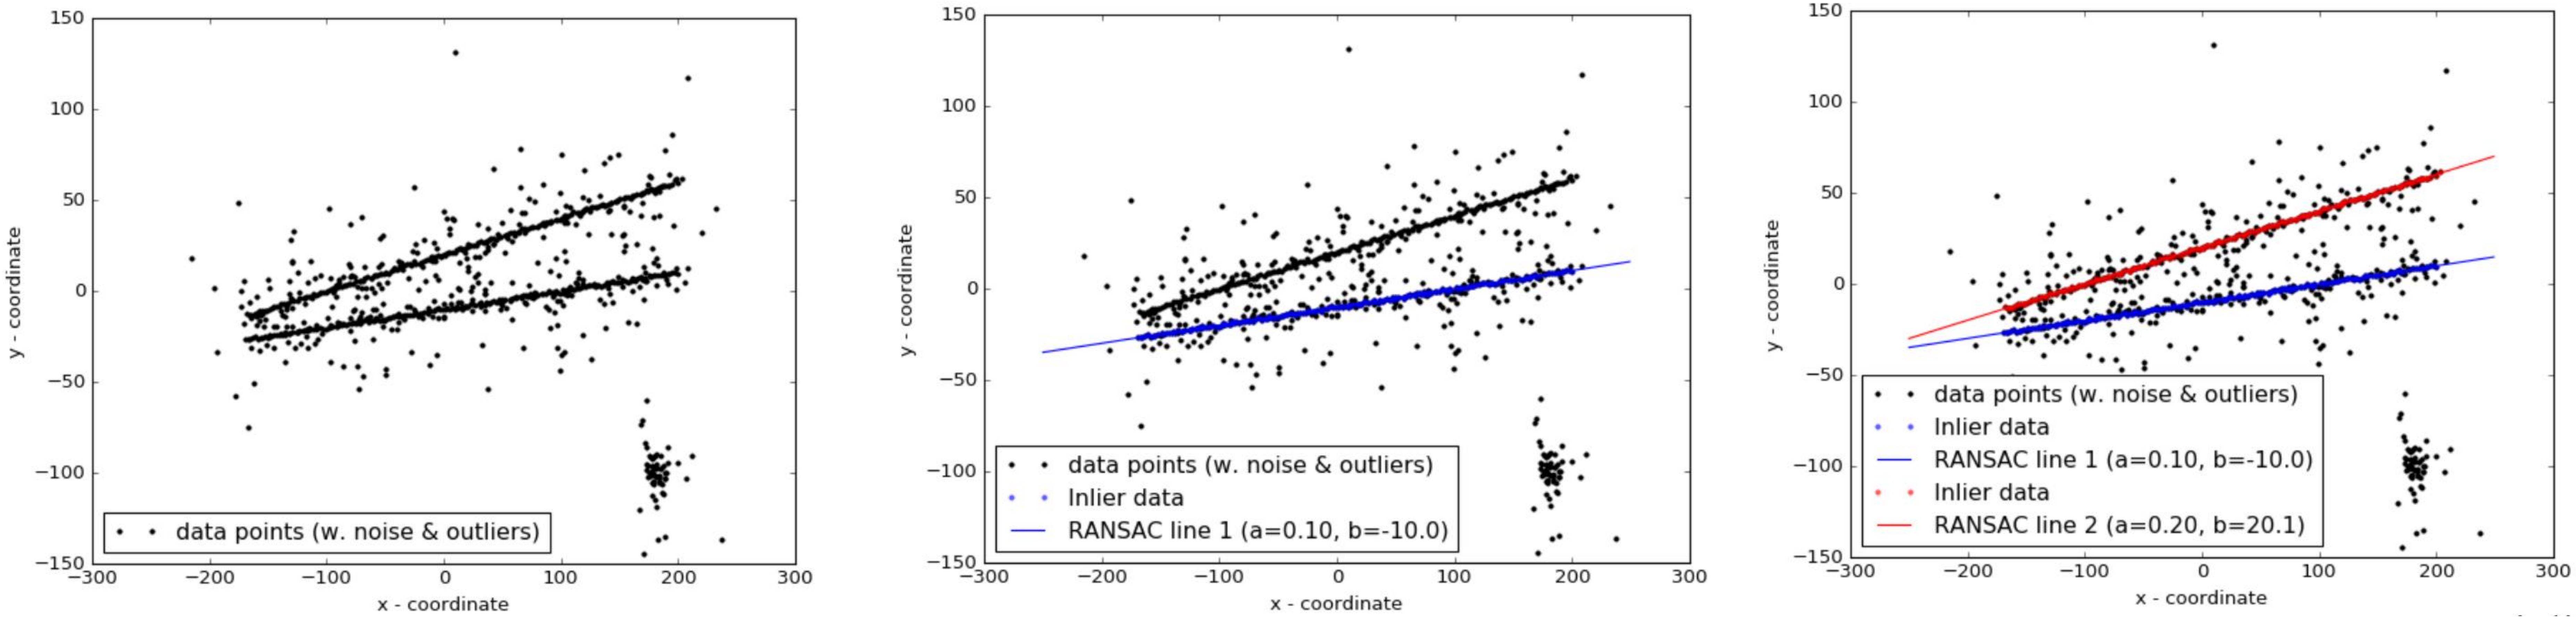

Example fitting by Sequential RANSAC.

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.

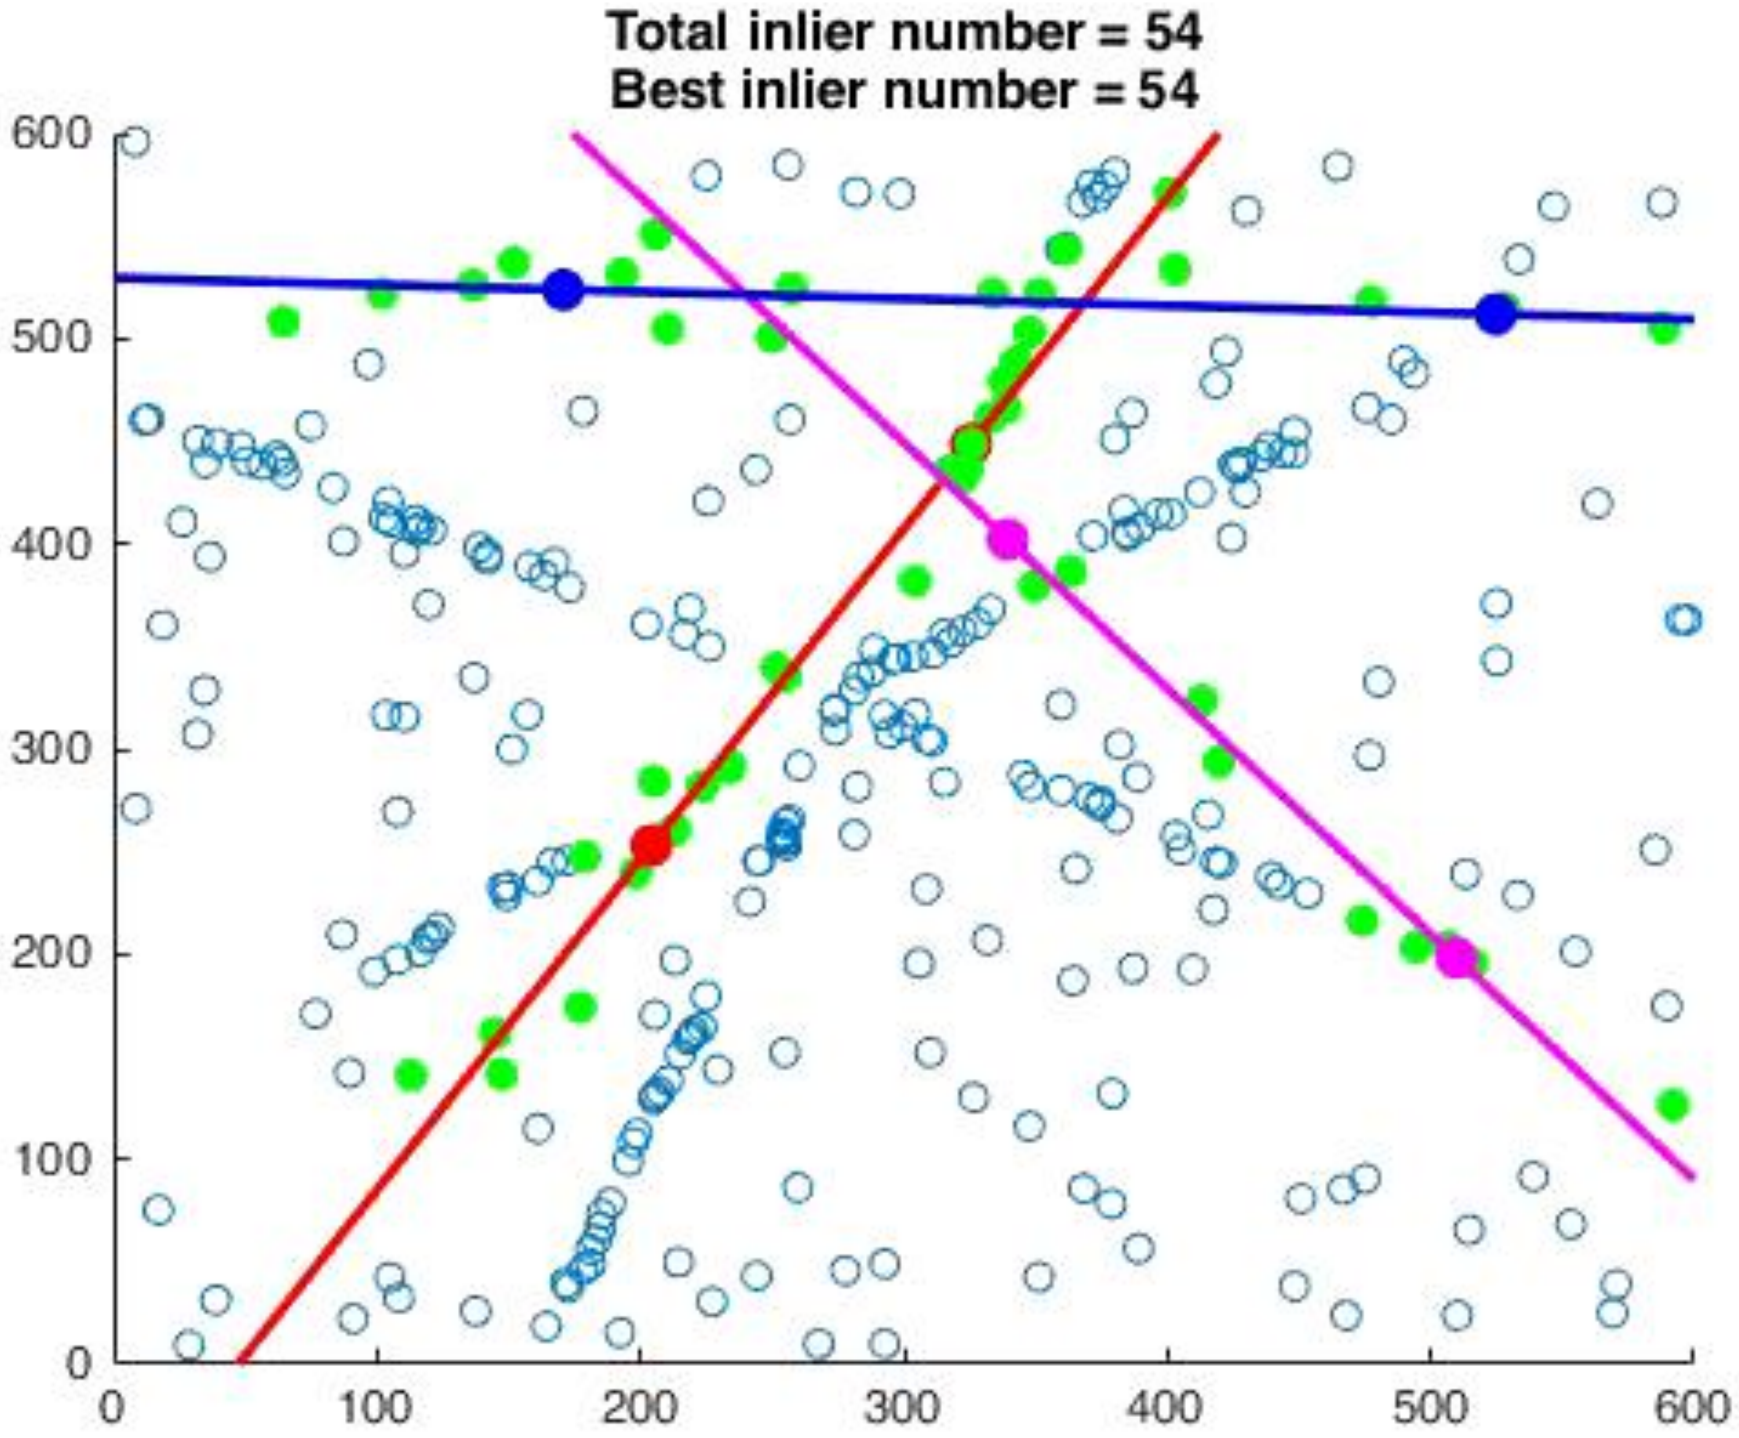

Example fitting by MultiRANSAC.

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

Recent Methods (RANSAC-like initialization)

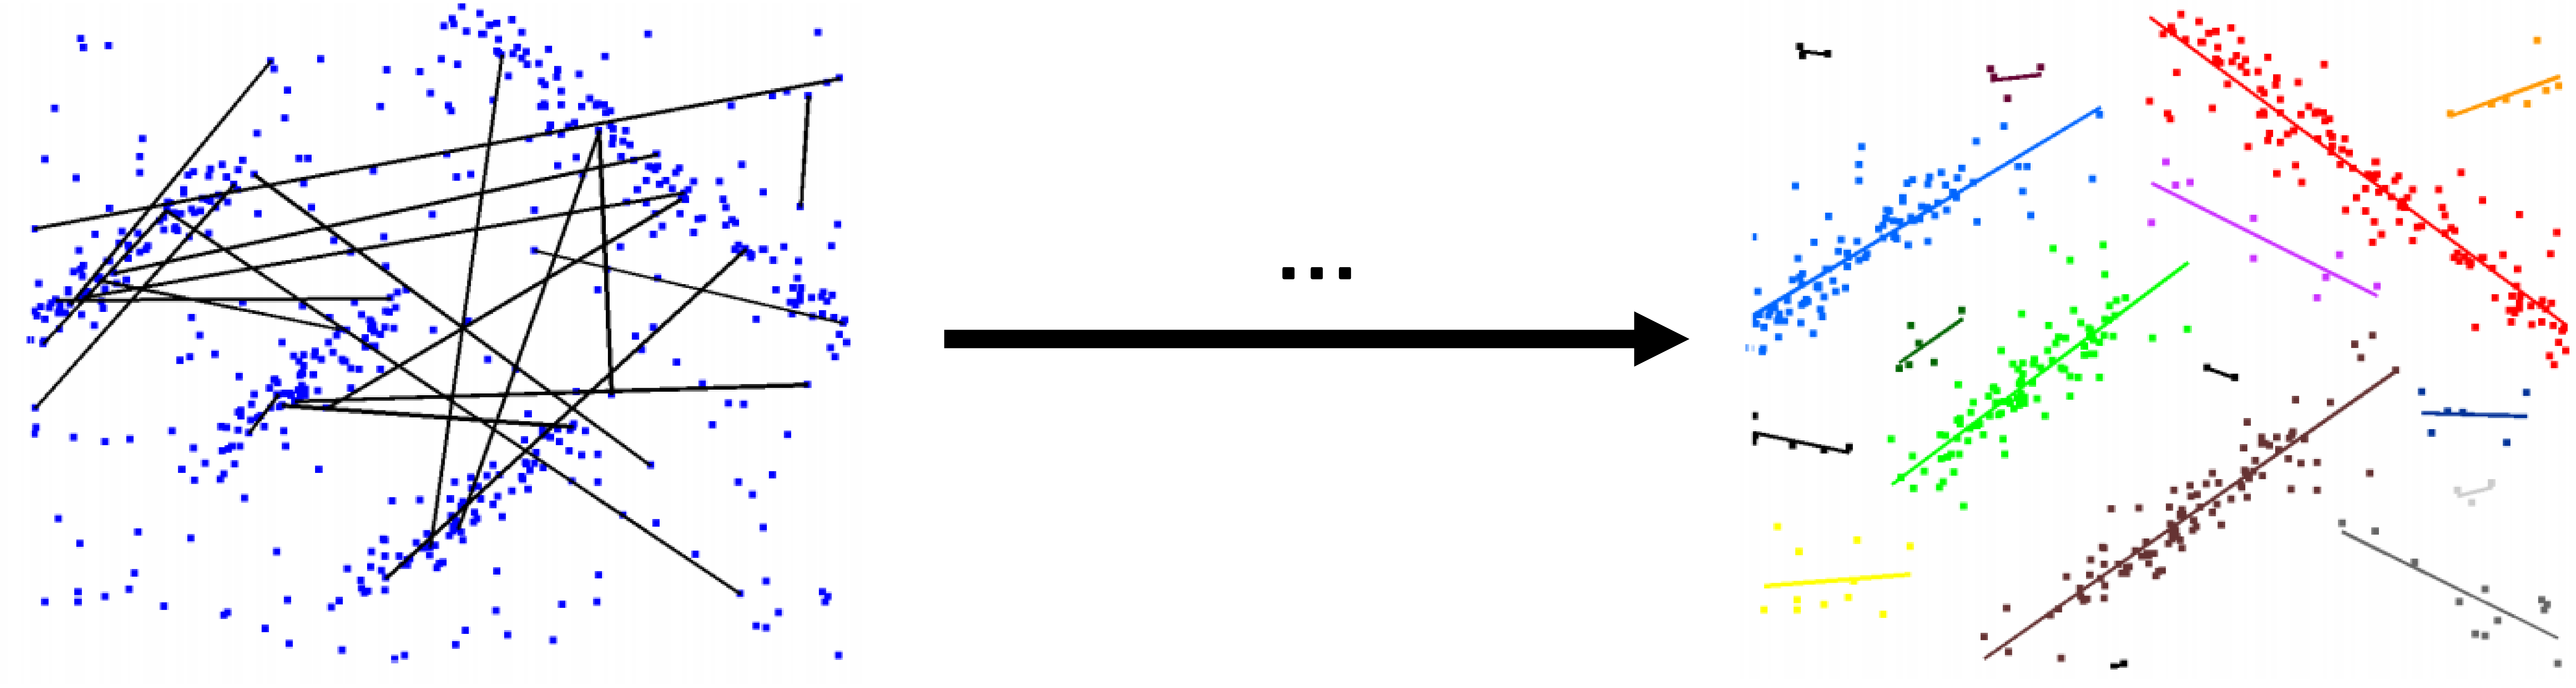

Pipeline:

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

...

Recent methods (RANSAC Inside the Optimization)

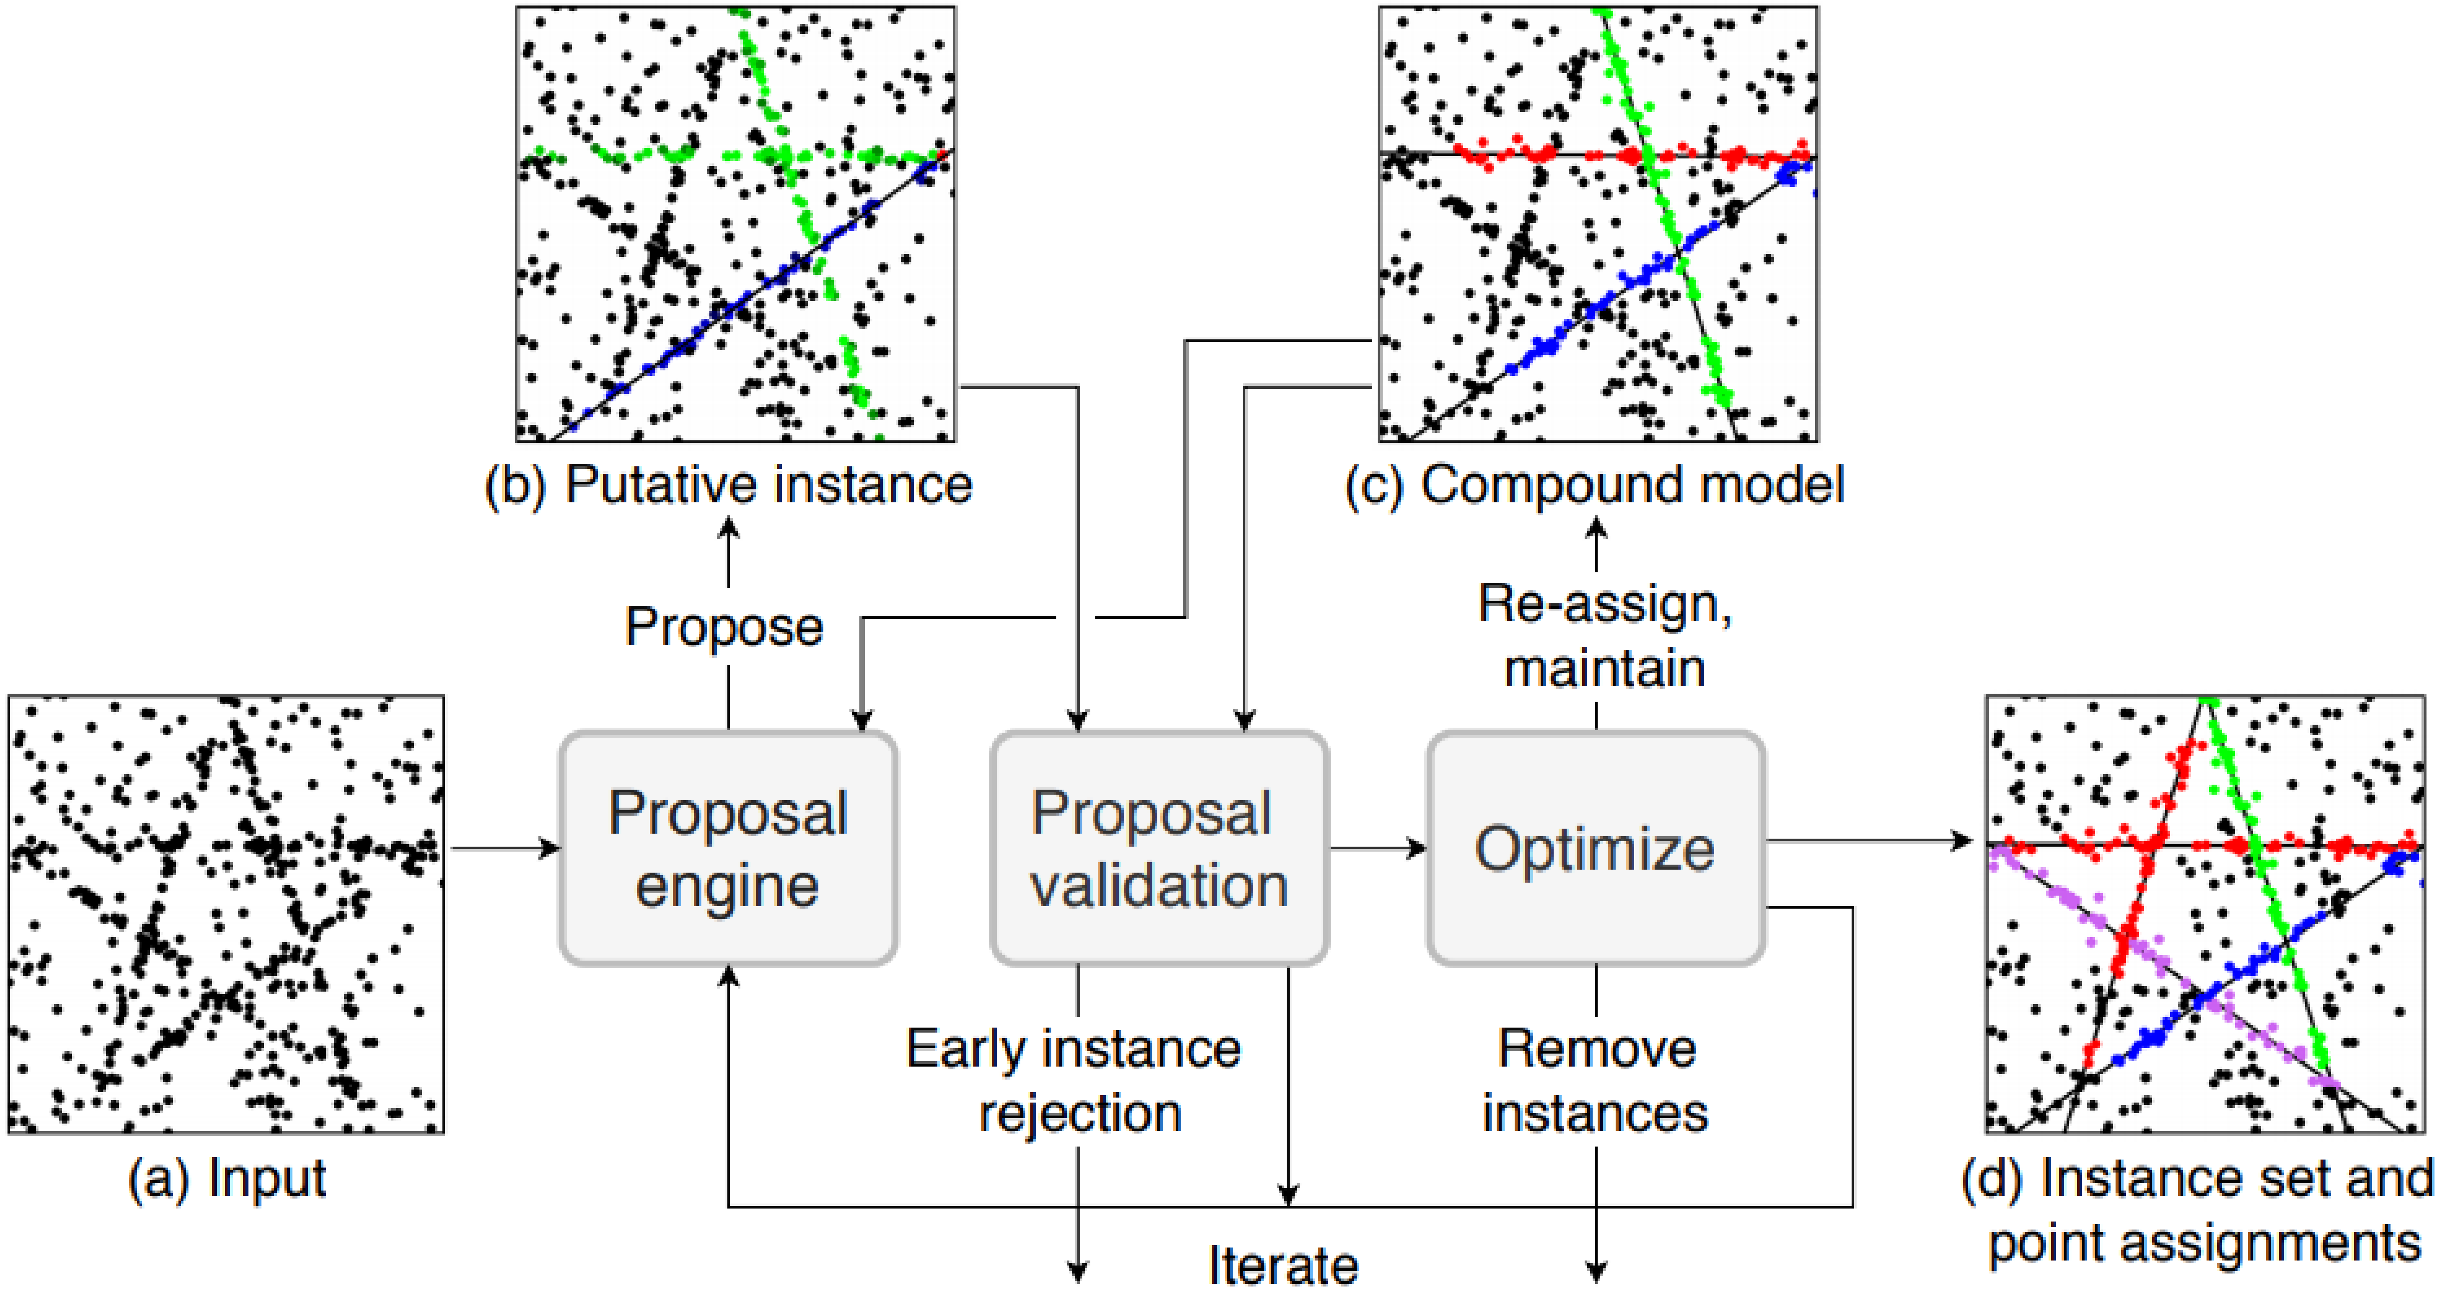

Pipeline of Progressive-X.

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)

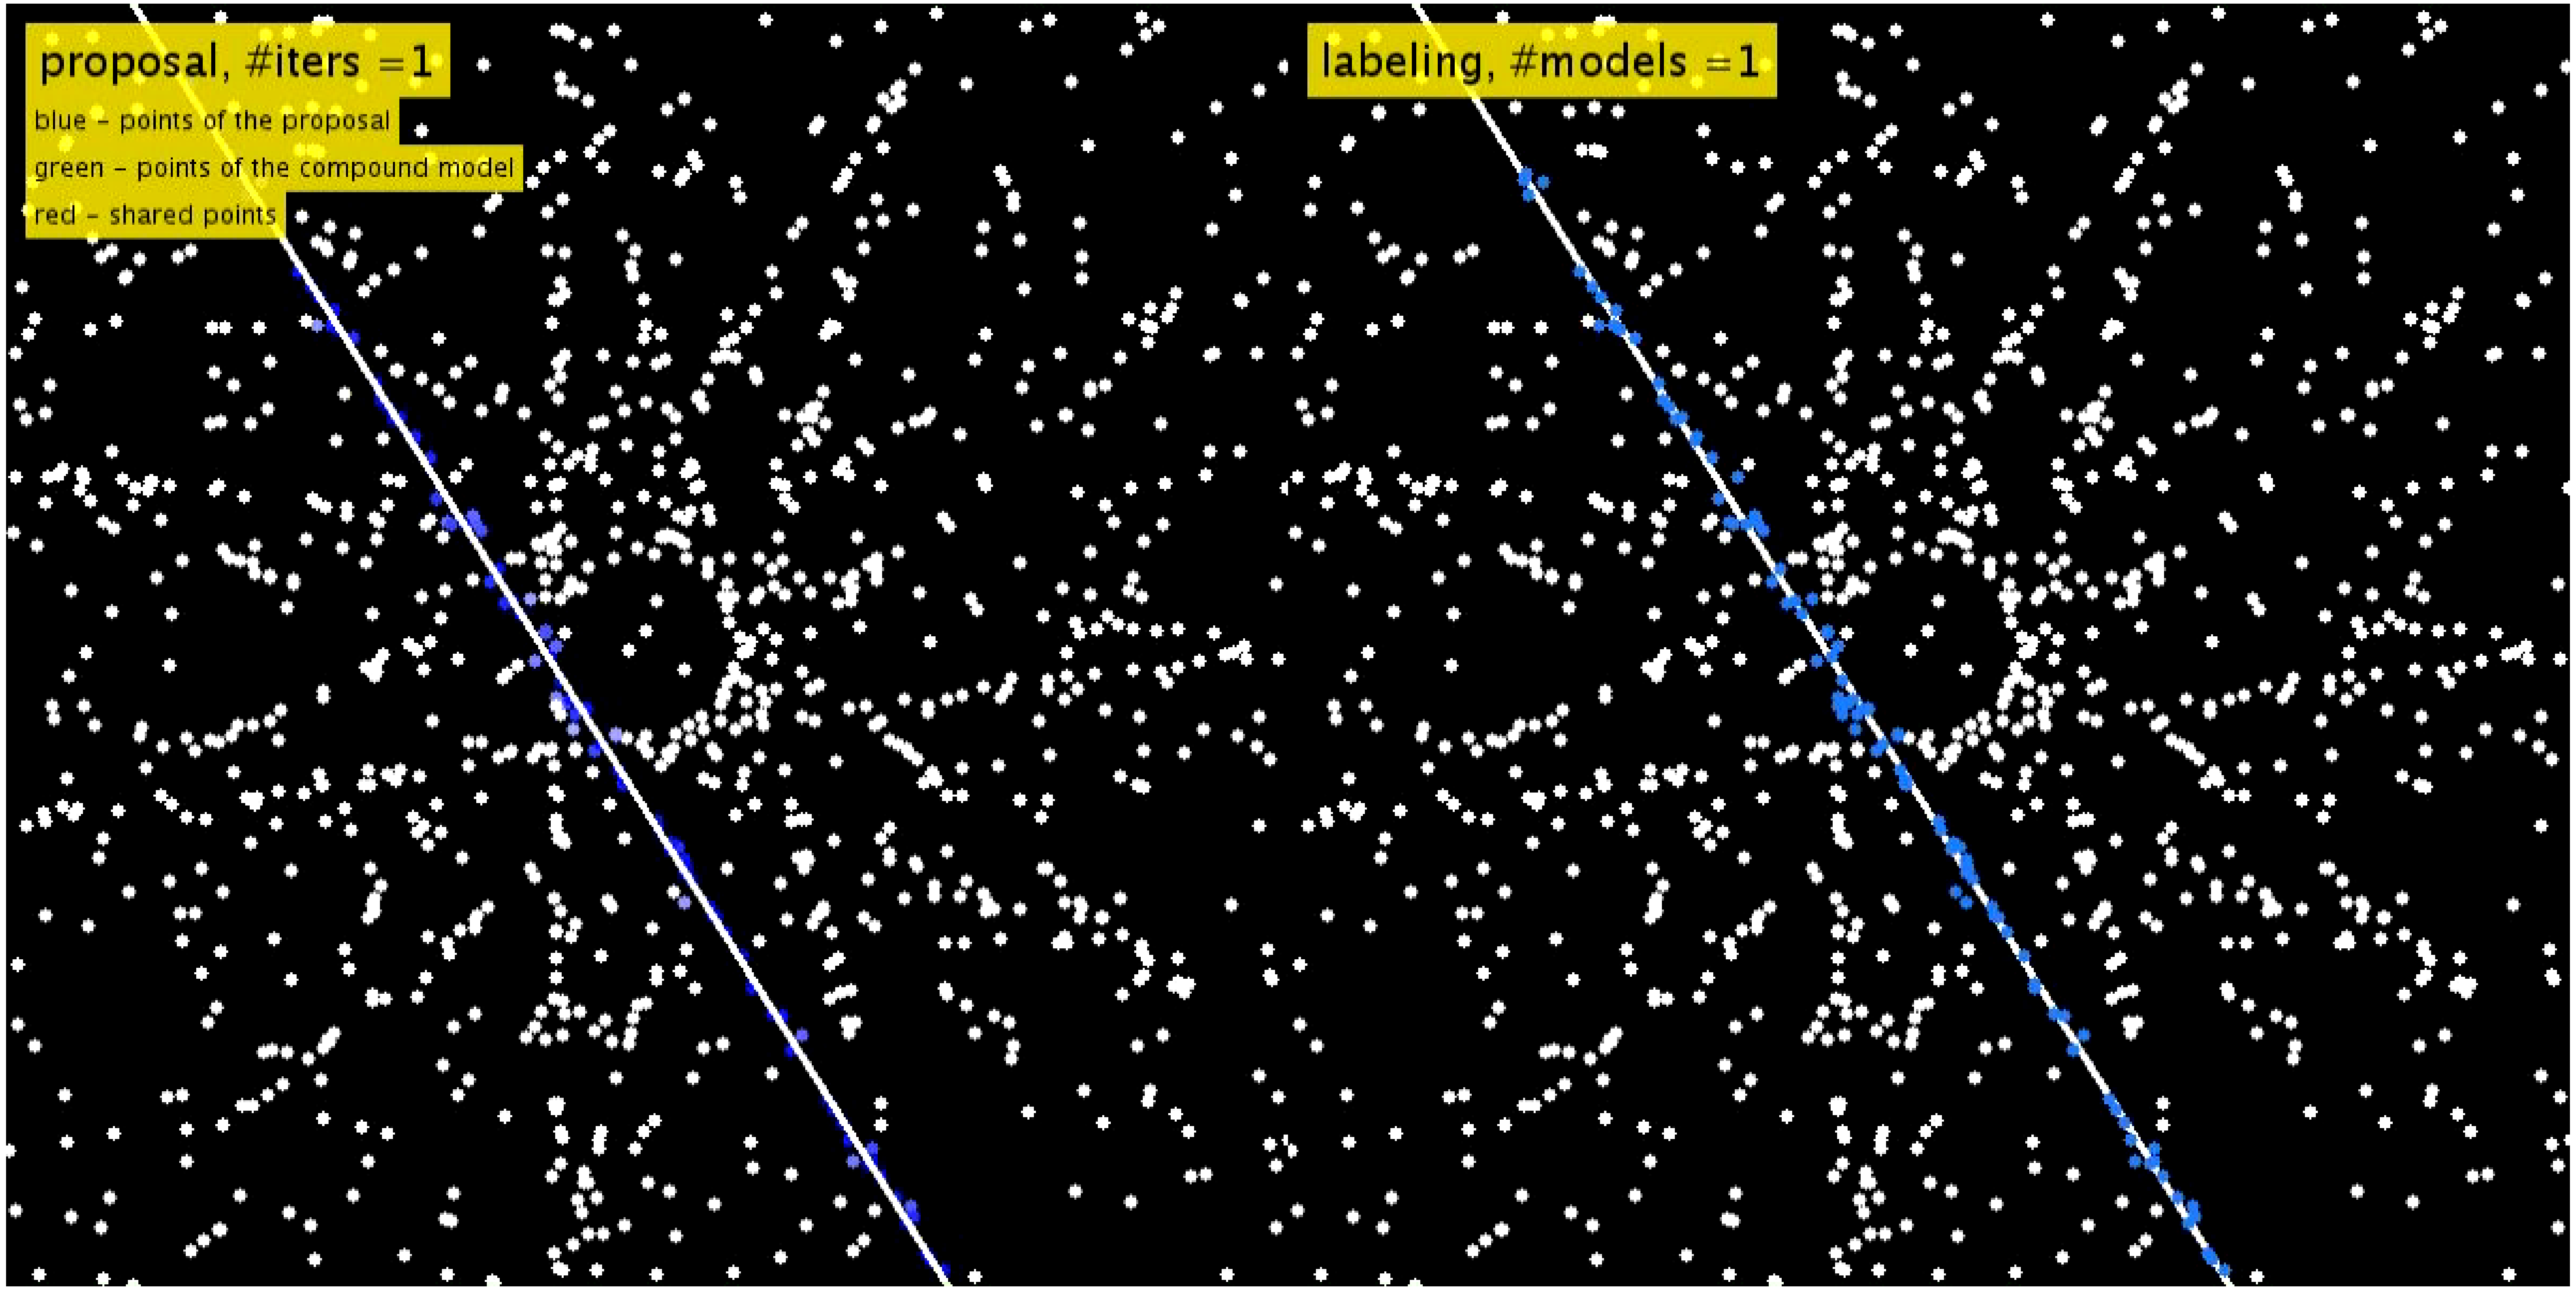

Example line fitting by Progressive-X.

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

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.*

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 *k*th closest point