Topologically Stable Hough Transform

Stefan Huber[1] Kristóf Huszár[2] Michael Kerber[2] Martin Uray[1]

[1] Josef Ressel Center for Intelligent and Secure Industrial Automation
Department for Information Technologies and Digitalisation
Salzburg University of Applied Sciences

Salzburg University
of Applied Sciences

[2] Institute of Geometry
Graz University of Technology

TU Graz

Thu, Mar 26
EuroCG 2026, Hagen, Germany

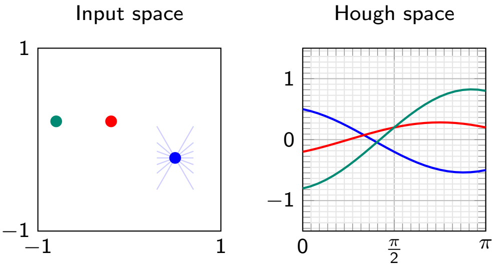

- Hough space = parameterization space of lines
  line in input space $\widehat{=}$ point in Hough space
- Blue curve in Hough space: the set of lines through the blue input point.
  - Intersection point of curves → common line in input space → to be detected
- Line detection: Discretize Hough space in bins, count votes (no. of curves through bin), find bins with many votes

- Classic method in CV, dates back to 1960s [MC15]
- Implemented in standard libraries like OpenCV
- Generalized to other shapes, like circles and “general” shapes
- Can be mathematically understood as Radon transform (1910s)

Classic approach: Threshold on the number of votes to identify peaks.

Inherent limitation:

- Prone to noise: Cleanup heuristics to eliminate false positives
- Bin discretization hurts stability and accuracy. Making discretization finer $\rightarrow$ bad for noise sensitivity and runtime

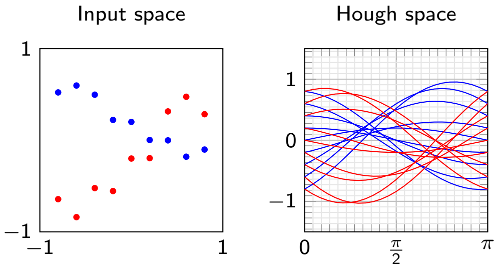

Rationale

Robust identification of peaks is where persistent homology (PH) shines.

Tracking the evolution of subspaces:

- Given a real-valued function $f\colon X \to \mathbb{R}$ over some topological space X
- We can look at the super-levelsets $f^{\geq a} = \{x \in X \mid f(x) \geq a\}$
- As a decreases from ∞, $f^{\geq a}$ grows and "sweeps out" X

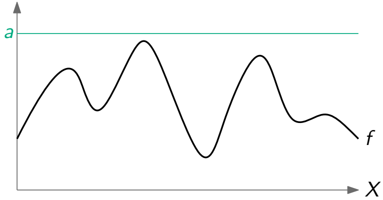

0-dimensional PH:

- Assume local maxima are isolated, with different f-values.
- Point x is dominant in $f^{\geq a}$ if x is f-maximal within its connected component.
  - If x is a local maximum, let $h = f(x)$, then x is dominant in $f^{\geq h}$. We call h the birth level of x.
- The persistence of x is the sup of $t \geq 0$ such that x is dominant in $f^{\geq h-t}$. We call $h - t$ the death level of x.
  - The global maximum has infinite persistence.

Tracking the evolution of subspaces:

- Given a real-valued function $f\colon X \to \mathbb{R}$ over some topological space X
- We can look at the super-levelsets $f^{\geq a} = \{x \in X \mid f(x) \geq a\}$
- As a decreases from ∞, $f^{\geq a}$ grows and “sweeps out” X

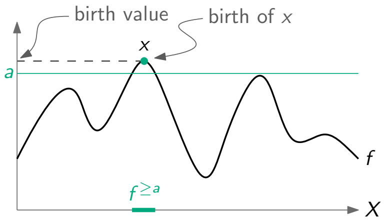

0-dimensional PH:

- Assume local maxima are isolated, with different f-values.
- Point x is dominant in $f^{\geq a}$ if x is f-maximal within its connected component.
  - If x is a local maximum, let $h = f(x)$, then x is dominant in $f^{\geq h}$. We call h the birth level of x.
- The persistence of x is the sup of $t \geq 0$ such that x is dominant in $f^{\geq h-t}$. We call $h - t$ the death level of x.
  - The global maximum has infinite persistence.

Tracking the evolution of subspaces:

- Given a real-valued function $f : X \to \mathbb{R}$ over some topological space X
- We can look at the super-levelsets $f^{\geq a} = \{x \in X \mid f(x) \geq a\}$
- As a decreases from ∞, $f^{\geq a}$ grows and "sweeps out" X

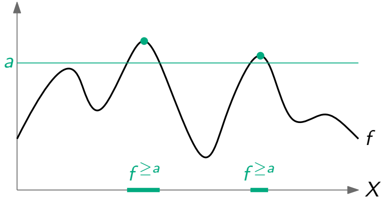

0-dimensional PH:

- Assume local maxima are isolated, with different f-values.
- Point x is dominant in $f^{\geq a}$ if x is f-maximal within its connected component.
  - If x is a local maximum, let $h = f(x)$, then x is dominant in $f^{\geq h}$. We call h the birth level of x.
- The persistence of x is the sup of $t \geq 0$ such that x is dominant in $f^{\geq h-t}$. We call $h - t$ the death level of x.
  - The global maximum has infinite persistence.

Tracking the evolution of subspaces:

- Given a real-valued function $f\colon X \to \mathbb{R}$ over some topological space X
- We can look at the super-levelsets $f^{\geq a} = \{x \in X \mid f(x) \geq a\}$
- As a decreases from ∞, $f^{\geq a}$ grows and "sweeps out" X

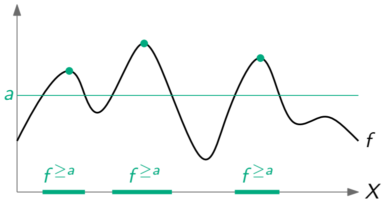

0-dimensional PH:

- Assume local maxima are isolated, with different f-values.
- Point x is dominant in $f^{\geq a}$ if x is f-maximal within its connected component.
  - If x is a local maximum, let $h = f(x)$, then x is dominant in $f^{\geq h}$. We call h the birth level of x.
- The persistence of x is the sup of $t \geq 0$ such that x is dominant in $f^{\geq h-t}$. We call $h - t$ the death level of x.
  - The global maximum has infinite persistence.

Tracking the evolution of subspaces:

- Given a real-valued function $f\colon X \to \mathbb{R}$ over some topological space X
- We can look at the super-levelsets $f^{\geq a} = \{x \in X \mid f(x) \geq a\}$
- As a decreases from ∞, $f^{\geq a}$ grows and "sweeps out" X

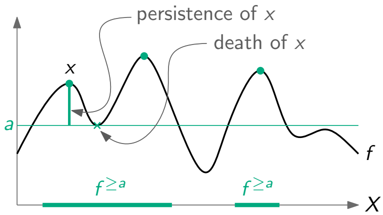

0-dimensional PH:

- Assume local maxima are isolated, with different f-values.
- Point x is dominant in $f^{\geq a}$ if x is f-maximal within its connected component.
  - If x is a local maximum, let $h = f(x)$, then x is dominant in $f^{\geq h}$. We call h the birth level of x.
- The persistence of x is the sup of $t \geq 0$ such that x is dominant in $f^{\geq h-t}$. We call $h-t$ the death level of x.
  - The global maximum has infinite persistence.

Tracking the evolution of subspaces:

- Given a real-valued function $f\colon X \to \mathbb{R}$ over some topological space X
- We can look at the super-levelsets $f^{\geq a} = \{x \in X \mid f(x) \geq a\}$
- As a decreases from ∞, $f^{\geq a}$ grows and "sweeps out" X

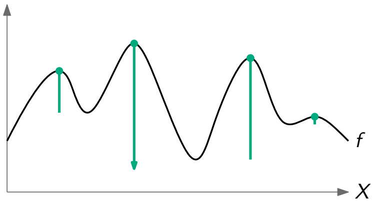

0-dimensional PH:

- Assume local maxima are isolated, with different f-values.
- Point x is dominant in $f^{\geq a}$ if x is f-maximal within its connected component.
  - If x is a local maximum, let $h = f(x)$, then x is dominant in $f^{\geq h}$. We call h the birth level of x.
- The persistence of x is the sup of $t \geq 0$ such that x is dominant in $f^{\geq h-t}$. We call $h-t$ the death level of x.
  - The global maximum has infinite persistence.

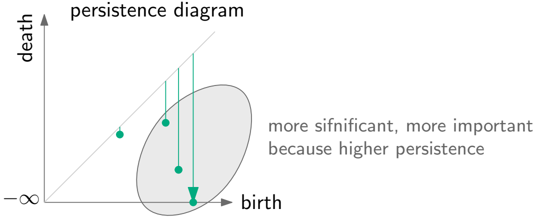

If we change f a little (e.g., noise) then also persistence (diagrams) changes only a little.

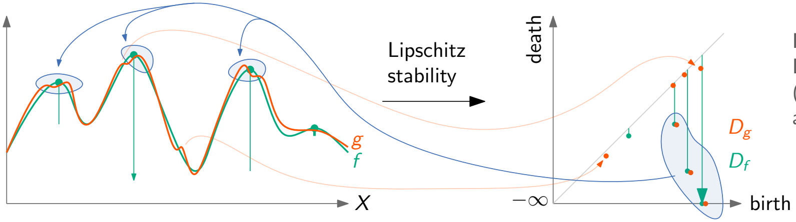

If f, g are close then the Diagram points are close. (Points close to diagonal are insignificant.)

Distance on persistence diagrams:

- D_f is the persistence diagram of f. (Formally a multiset, with ∞-multiplicity at each diagonal point.)
- $d_B(D_f, D_g)$ is the bottleneck distance. (Which is $\inf_{\gamma\colon D_f \leftrightarrow D_g} \sup_{x \in D_f} \|x - \gamma(x)\|_\infty$.)

Theorem (Cohen-Steiner, Edelsbrunner, Harer 2007 [CEH07])

For any two (tame) functions $f, g\colon X \to \mathbb{R}$,

$$d_B(D_f, D_g) \leq \|f - g\|_\infty.$$

Why PH for Hough transform?

1

Understanding:

- As a tool to understand the limitations of the heuristics to identify noisy peaks.

2

Algorithm construction:

- Replace heuristics to robustly identify noisy peaks by Lipschitz-stable mathematical methods.

Programme

Lipschitz-stable sequence

$$\text{input point cloud} \mapsto \text{Hough image} \mapsto \text{persistence diagram (persistent points} \mathrel{\widehat{=}} \text{peaks} \mathrel{\widehat{=}} \text{input lines)}$$

For the first map to be continuous, we look for a continuous Hough image, not votes on a discrete bin grid.

Hough space

As the point-line dual space with Möbius topology

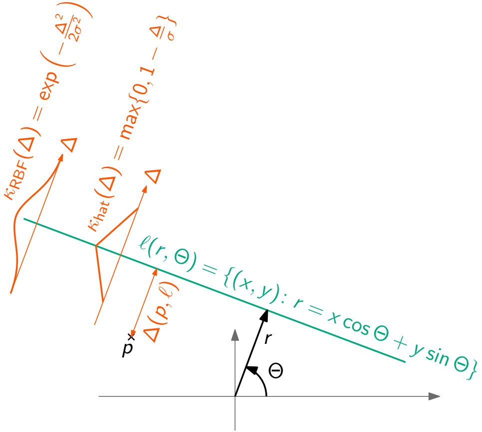

Hough space:

- $M = \mathbb{R} \times [0, \pi]$, where $(r, \Theta) \in M$ parametrizes a line $\ell(r, \Theta)$
- Note that $\ell(r, \Theta) = \ell(-r, \Theta + \pi)$, leading to well-known open Möbius strip topology[1] of M

Each point p spends a vote $S(p, \ell)$ to a line ℓ:

- $S(p, \ell) = \kappa(\Delta(p, \ell))$. Examples for κ are κ_{hat} or κ_{RBF}
- Require $\kappa \colon [0, \infty) \to [0, 1]$ with $\kappa(0) = 1$, monotonously decreasing, $\kappa(x) \to 0$ as $x \to \infty$

The score function $S \colon M \to [0, 1]$ collects all such votes for an input point set P:

- $S_P(\ell) = \frac{1}{|P|} \sum_{p \in P} S(p, \ell)$.
- Goal: Detecting peaks in S_P.

[1] Defining $(r, \Theta) \equiv (-r, \Theta + \pi)$ then $M/\equiv$ is homeomorphic to the open Möbius strip.

Persistence-based selection of peaks

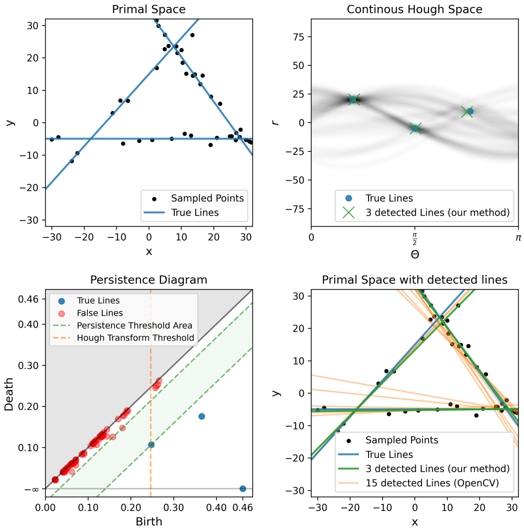

Theory for persistence-based peak computation:

- Stability: Let κ be λ-Lipschitz and P' be an ε-perturbation of P. Then $\|S_P - S_{P'}\|_\infty \leq \lambda\varepsilon$.
- Matching maxima: Let $\|S_P - S_{P'}\| = \varepsilon$ then there is a (partial) matching between local maxima (if persistence at least ε) and persistence differs by at most 2ε.
- Location of matching maxima: Let x be local maximum of S_P with birth level h and persistence α. Let $S_{P'}$ with $\|S_P - S_{P'}\| = \varepsilon < \alpha/2$. Then $S_{P'}$ contains local maximum x' with persistence $\geq \alpha - 2\varepsilon$, $|S_P(x') - h| \leq 2\varepsilon$ and x' in connected component of x in $S_P^{\geq h-\alpha}$.

Algorithm sketch:

- Discretize Hough space, with grid size by above theorems.
- Compute persistence diagram, e.g., cubical complexes, but respect the Möbius topology of M.
- Pick local maxima by thresholding persistence. (Classic methods pick by birth level.)

Approximate S_P based on a quad-tree subdivision that yields $\tilde{S}_P$

- with $\|S_P - \tilde{S}_P\|_\infty \leq \varepsilon$ and
- is constant on each quad-tree cell.

Subdivide cell B if

- $\lambda \cdot \mathrm{diam}(B)/2\varepsilon$.
- Trivial λ is (global) $\lambda_\kappa\sqrt{2}$, but we show how to improve by a localized version.

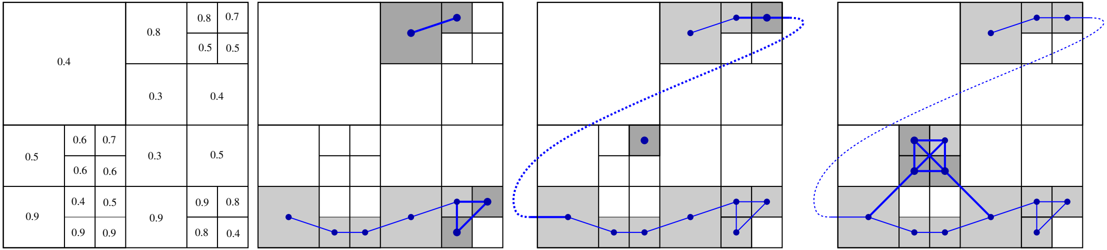

Figure: Sequence of super-levelsets and the nerve (blue) of the quad-tree subdivision, with Möbius topology, yielding persistence diagram.

Conclusion

Why the problem is nice:

- Tools from computational topology to methodologically revisit and improve a classic algorithm in CV.
- Tools from computational geometry with a grain of topology to improve runtime performance.
- A nice application setting for PH that is
  - not about point clouds,
  - not in "standard" Euclidean $\mathbb{R}^n$,
  - in a dual space of the original problem.

What's next:

- Method: General shapes
- Theory: Improving on stability of locus of lines
- Experiments: Extensive evaluation (toy examples for now)

arXiv paper

GitHub code repository

Thank you

Strudel pictures

For $P = \{p\}$, the function graph of S_P has a strudel[2] shape: the kernel κ slides along the sinusoidal curve for p.

$P = \{(0.5, 0.2)\}, \kappa = \kappa_{\mathrm{RBF}}, \sigma = 0.1$

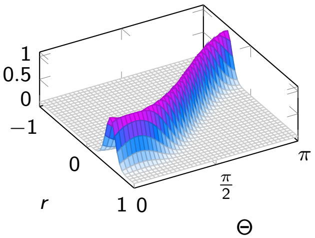

$\{(0.5, 0.2), (-0.4, 0.3)\}$

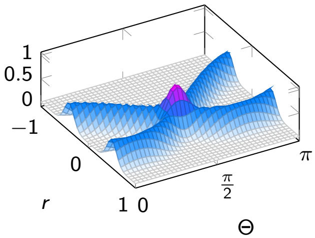

$\{(0.5, 0.2), (-0.4, 0.3), (0.2, 0.25)\}$

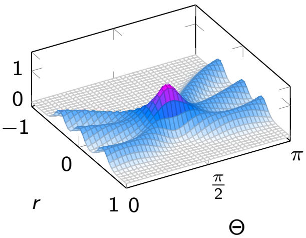

[2] https://en.wikipedia.org/wiki/Strudel

Input in square of side-length 1.

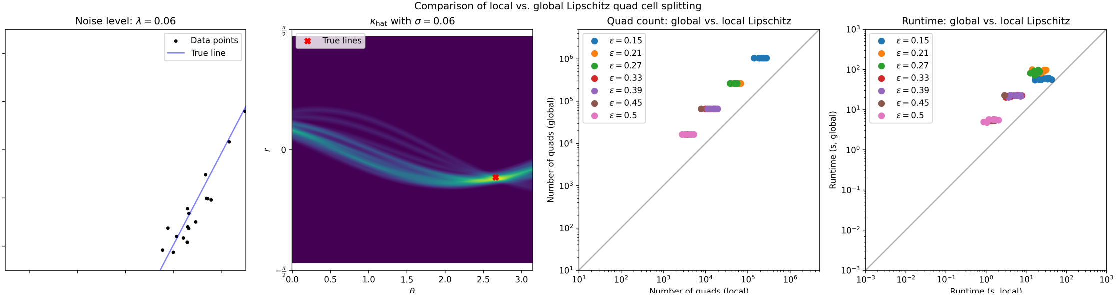

Input in square of side-length 1.

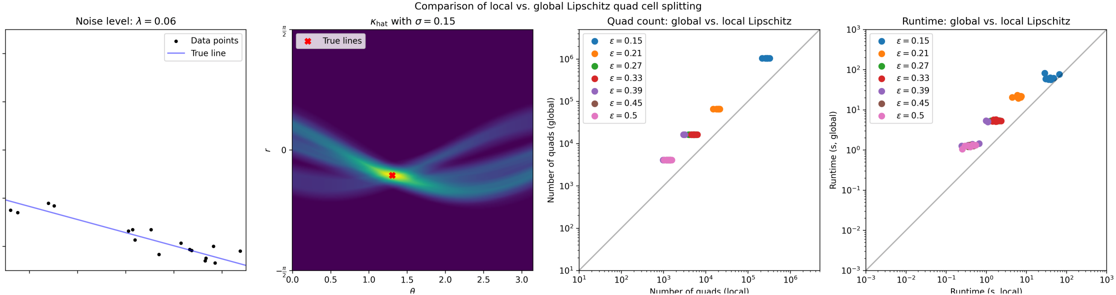

Quadtree cell splitting: Local versus global Lipschitz constant

Input in square of side-length 1.

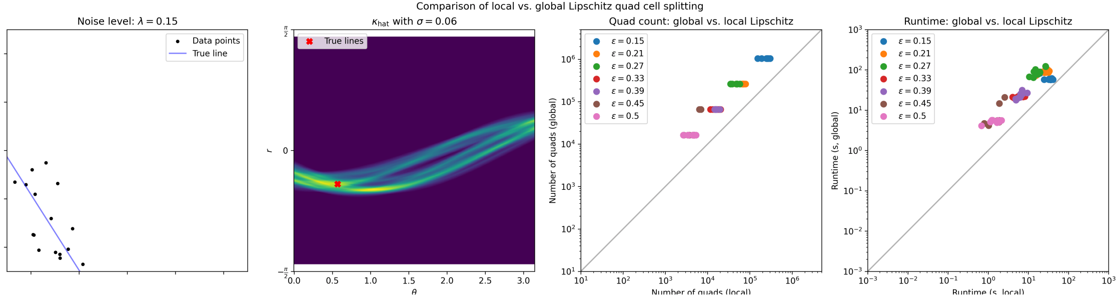

Input in square of side-length 1.

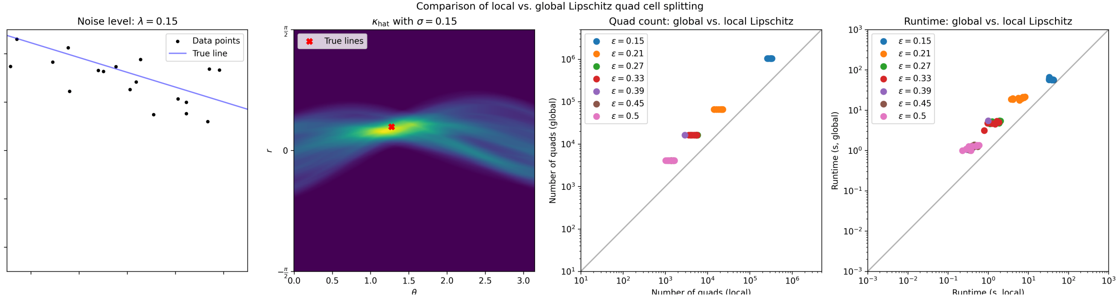

Input in square of side-length 1.

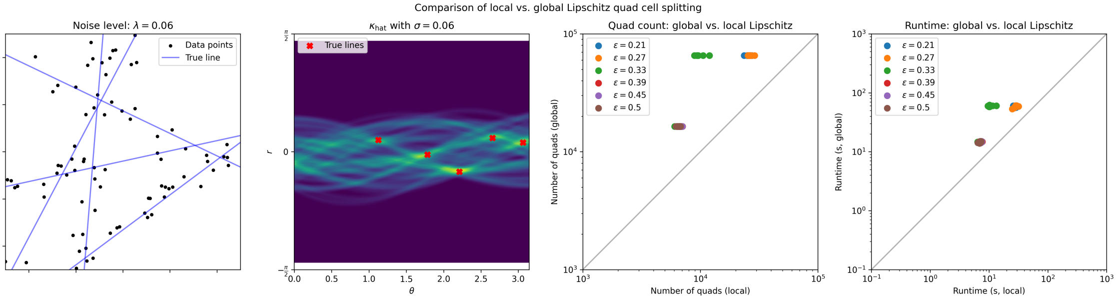

Input in square of side-length 1.

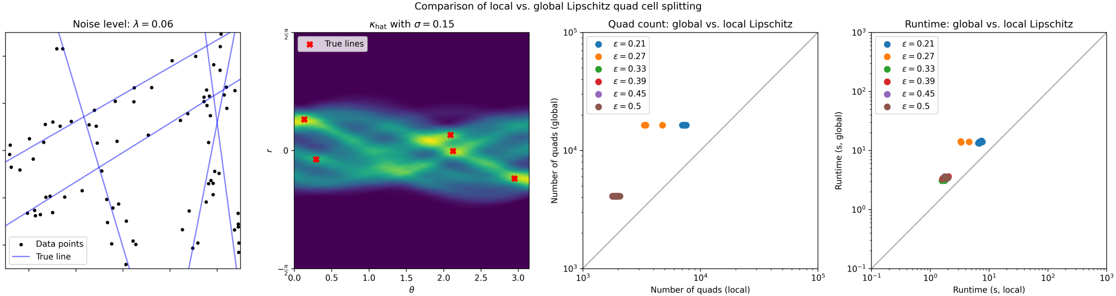

Input in square of side-length 1.

Comparison of local vs. global Lipschitz quad cell splitting

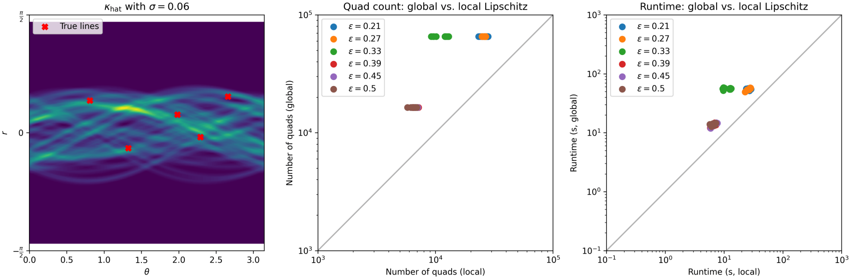

Input in square of side-length 1.

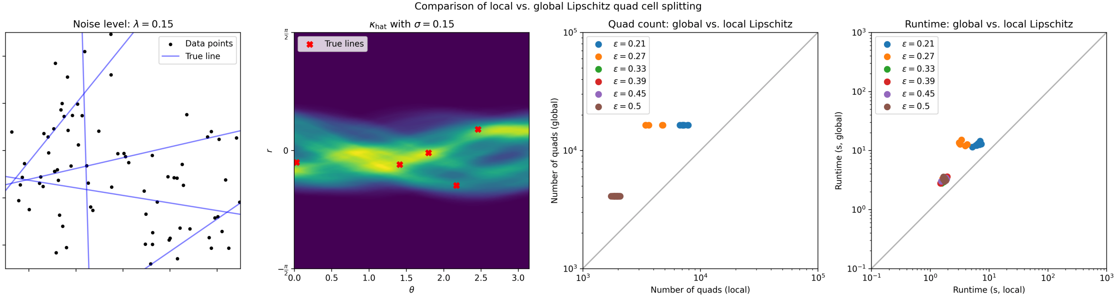

[CEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. "Stability of persistence diagrams." In: *Discrete & Computational Geometry* 37 (2007), pp. 103–120. ISSN: 0179-5376. DOI: 10.1007/s00454-006-1276-5.

[MC15] Priyanka Mukhopadhyay and Bidyut B. Chaudhuri. "A Survey of Hough Transform." In: *Pattern Recognition* 48.3 (2015-03), pp. 993–1010. ISSN: 0031-3203. DOI: 10.1016/j.patcog.2014.08.027.