Homologically Persistent Skeleton for a 2D cloud of features

Vitaliy Kurlin, kurlin.org, Microsoft Research Cambridge and Durham University, UK

Motivation: fast and robust recognizing visual markers

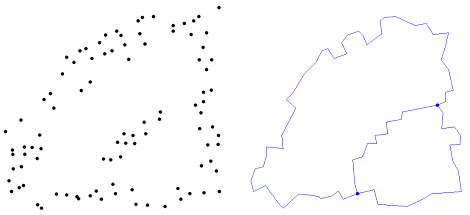

- simple pictures like hieroglyphs are easily readable by humans;
- machines should correctly reconstruct skeletons from noisy scans.

Skeletonization problem:

Given only a cloud C of points, find a graph representing the topology of C across all scales.

All past methods use extra input parameters: a scale, noise bound.

New solution: Homologically Persistent Skeleton HoPeS(C)

- extends a Min Spanning Tree;
- depends only on the cloud C;
- has *min length* among all graphs that span C at any scale and also have most persistent 1D cycles;
- has derived subgraphs with the correct topological type of a graph $G \subset \mathbb{R}^2$ given only by a sample C;
- is *globally stable* under perturbations of C, remains in an offset C^α within a distance α from $C \subset \mathbb{R}^2$.

Def: HoPeS(C) of a cloud C

When the scale α is increasing, the offset C^α is growing, cycles in C^α are born and die (become filled).

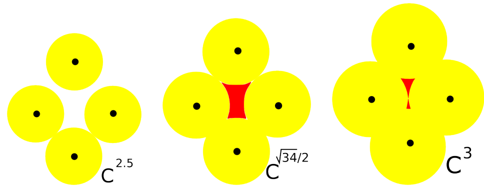

The pairs (birth, death) form the 1D *persistence diagram* PD$\{C^\alpha\}$.

Earlier applications of persistence: counting holes in clouds [2], auto-completing closed contours [3].

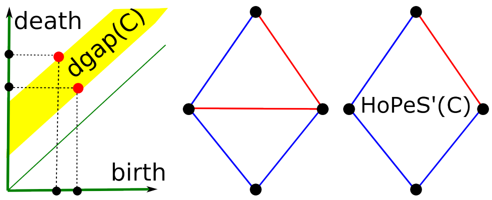

HoPeS(C) is a Min Spanning Tree of C plus all red edges giving birth to new cycles in C^α across all α.

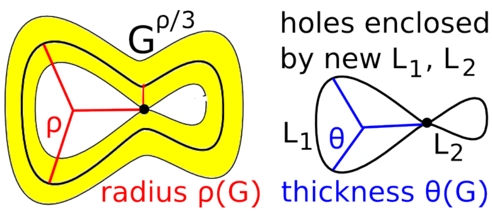

If cycle $L \subset G$ encloses a hole, its *radius* ρ is $\min \alpha$ when L^α is contractible. The *thickness* $\theta(G)$ is max radius of all newborn cycles in G^α.

Reconstruction guarantees

An *ε-sample* C of a graph $G \subset \mathbb{R}^2$ is a cloud $C \subset G^\varepsilon$ with $C^\varepsilon \supset G$.

Th (VK). Let C be an ε-sample of a graph $G \subset \mathbb{R}^2$ with $\theta(G) \geq 0$ and radii $\rho_1 \leq \ldots \leq \rho_m$. If $\rho_1 > 7\varepsilon + \theta(G) + \max\limits_{i=1,\ldots,m-1}\{\rho_{i+1} - \rho_i\}$, then HoPeS$'(C) \sim G$ is 2ε-close to G.

[4]: extension to a metric space.

Summary and References

input: a noisy point cloud $C \subset \mathbb{R}^2$

output: full skeleton HoPeS(C) with all derived subgraphs;

running time: $O(n \log n)$ for any n points with real 2D coordinates;

more details: C++ code, papers, examples at http://kurlin.org.

[1] H. Edelsbrunner and J. Harer. *Computational topology*. AMS'10.

[2] V. Kurlin. A fast and robust algorithm to count persistent holes in noisy 2D clouds. *CVPR 2014*.

[3] V. Kurlin. Auto-completion of contours in sketches, maps and sparse 2D images. *CTIC 2014*.

[4] V. Kurlin. Homologically Persistent Skeleton of a point cloud in any metric space. *Computer Graphics Forum*, v. 34-5 (2015), 253-262.

1) cloud C of Canny edge points. 2) PD$\{C^\alpha\}$. 3) HoPeS$_{1,1}(C)$ for 1st widest gap. 4) HoPeS$_{2,1}(C)$ for 2nd widest gap.