Discrete Morse Theory

An overview & recent directions

Born out of classical Morse Theory (Milnor '63)

- Use continuous functions to describe a topological space
- Critical points are meaningful!

Classic Eg.!

Smooth Manifold:

$f: M \to \mathbb{R}$

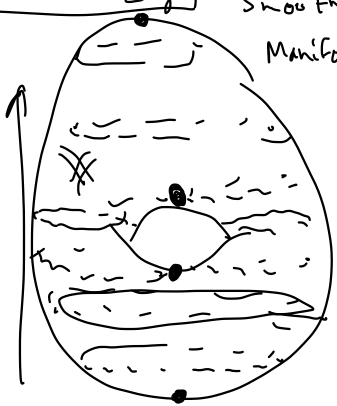

Level sets

Graph Representation

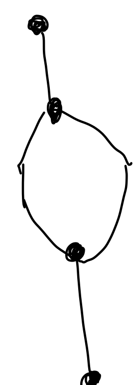

Reeb Graph

Smooth to Discrete:

Turns out, the same ideas work in a discrete setting. (Forman)

- Take same types of functions, but on a simplicial complex.

Eg:

$f: K \to \mathbb{R}$

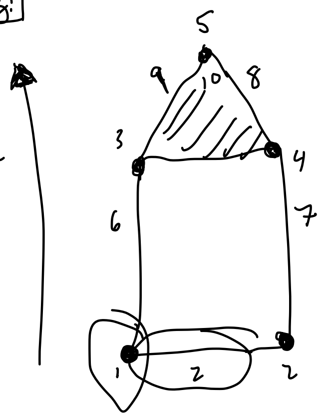

Algebraic Defn:

simplices go up in function value generally as they increase in dimension.
(with at most one exception per simplex)

Link to homology:

of critical i-cells bounds β_i

1 0-cell

1 1-cell

Other Definitions (Part of why DMT is so nifty.)

Topological: Gradient Vector field with arrows on K.
$\sigma \in K$ not part of an arrow is deemed critical.

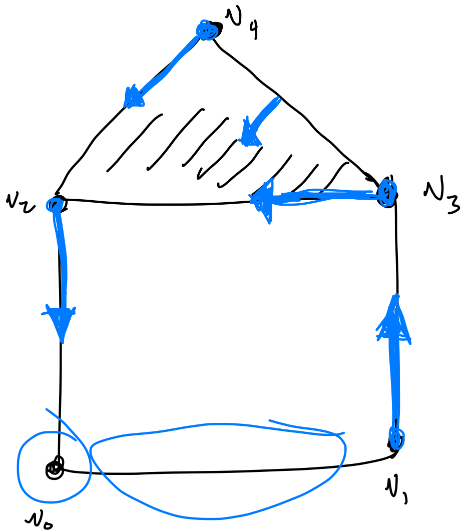

Combinatorial: Matchings on Hasse diagram corresponding to heads & tails of arrows.
unmatched $\rightarrow$ critical.

DAG $\rightarrow$ Morse fcn

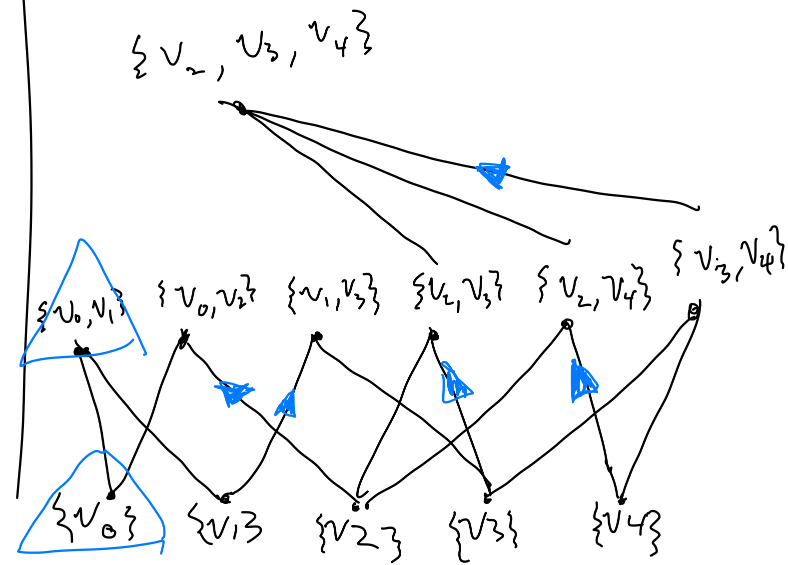

What's the use?

- simplify data w/o losing homology.
(set of instructions to collapse K)

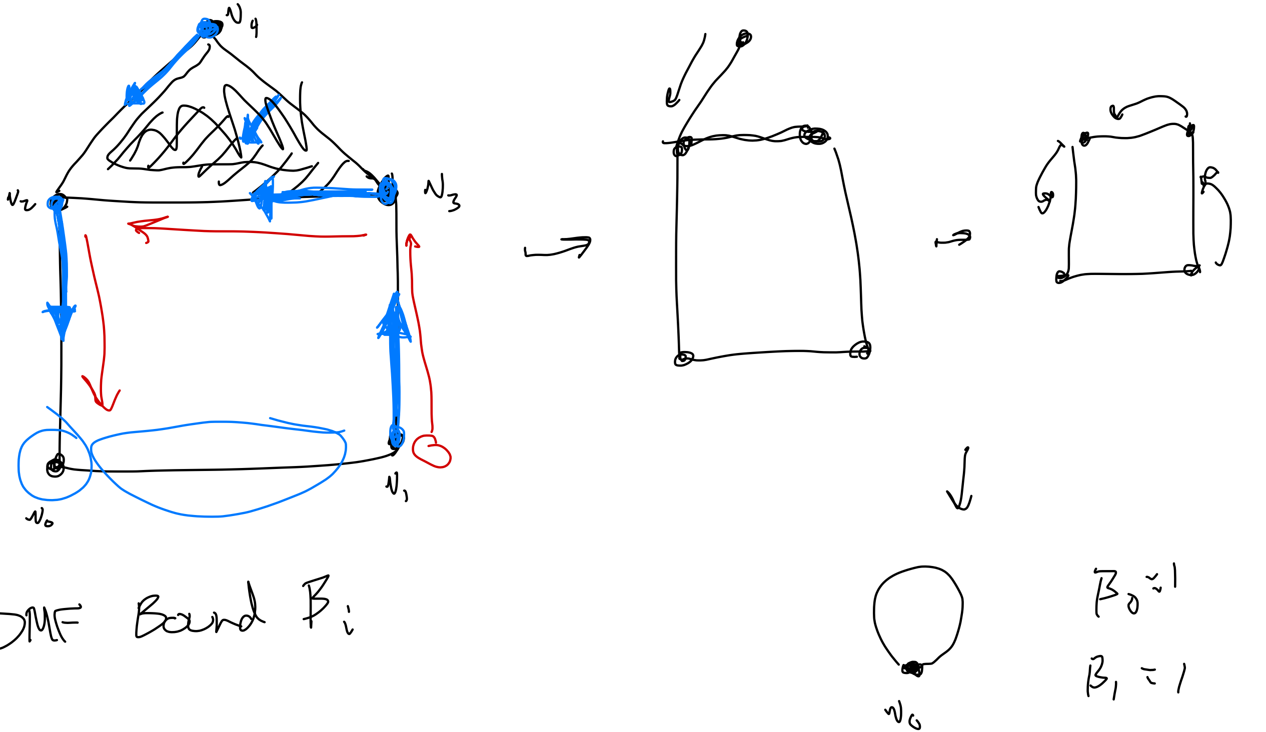

$\beta_0 = 1$

$\beta_1 = 1$

DMF Bound β_i

Problem: Creating a discrete Morse function w/ minimal critical cells is NP-Hard. (Joswig et al.)

→ Get as close to β_i as possible!

Idea: (King et al.)

Add vertex data to K and use that. (Poly time) - Extract

Given $F: K_0 \to \mathbb{R}$

2 parts:

(1) - Raw extraction → pair smallest lex σ_j largest lex σ_{j-1} coface

- Cancel extra critical cells → examine gradient paths from σ_j + C_j to σ_{j-1} + C_{j-1}

dim ≤ 2

Extract Right Child: $\Theta(dn)$ w/ $O(n)$ space

- Do Extract Raw directly, without recursion

n = # of simp. of K

$d = \dim(K)$

Eg: Extract Right Child:

$f: K \to \mathbb{R}$ as follows:

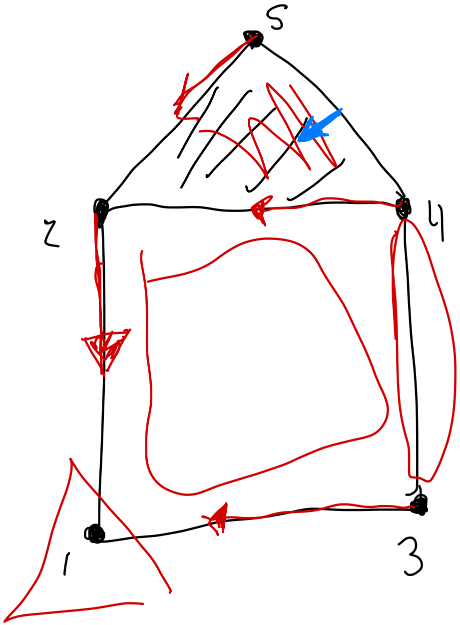

(Collapse)

sorted: Left → Right
smallest lex. → largest

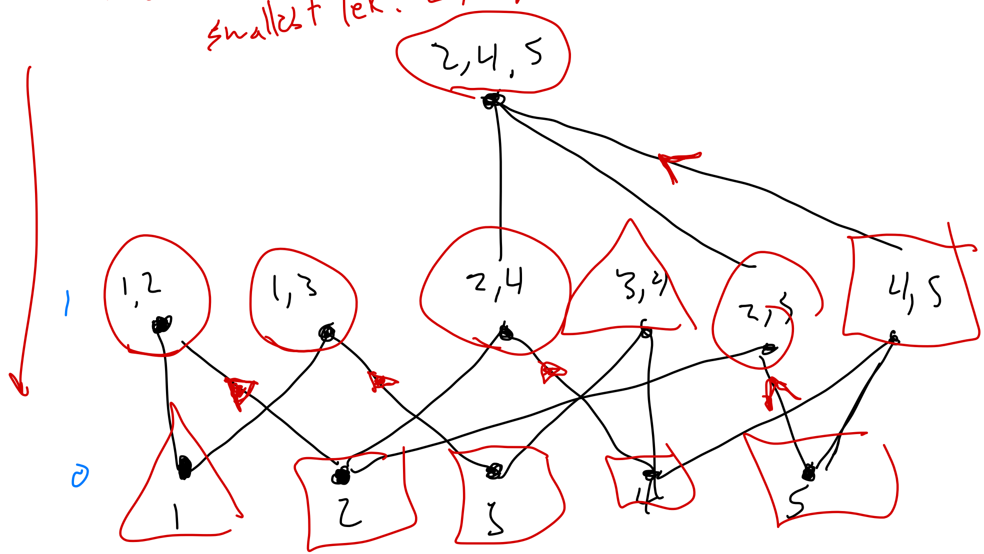

Eg 2: Problem Vertex Data:

$f: K_0 \to \mathbb{R}$

$\beta_0 \leq 2$

$\beta_1 \leq 2$

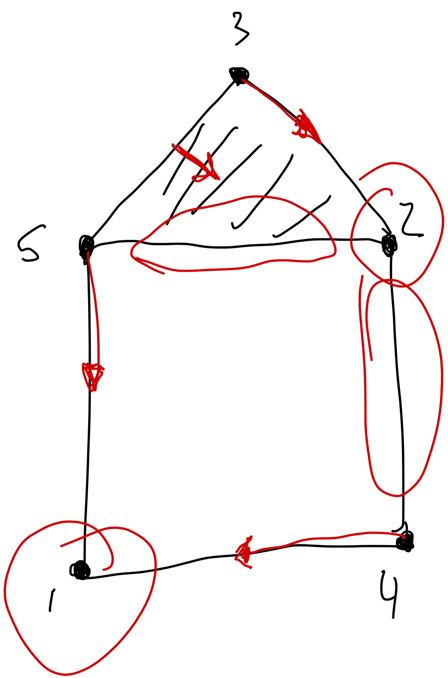

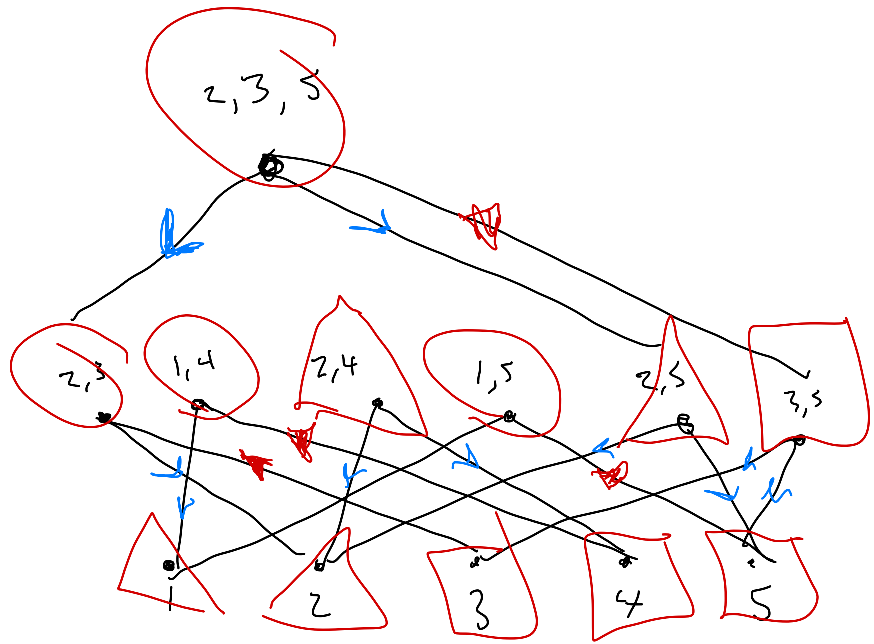

What to do now? Cancel!

Thm: If $\exists$ a gradient path from $\sigma \in C_i$ to unique $\sigma' \in C_{i-1}$, one may cancel both σ and σ' by reversing the path between them.
(Bauer) (King)

Eg:

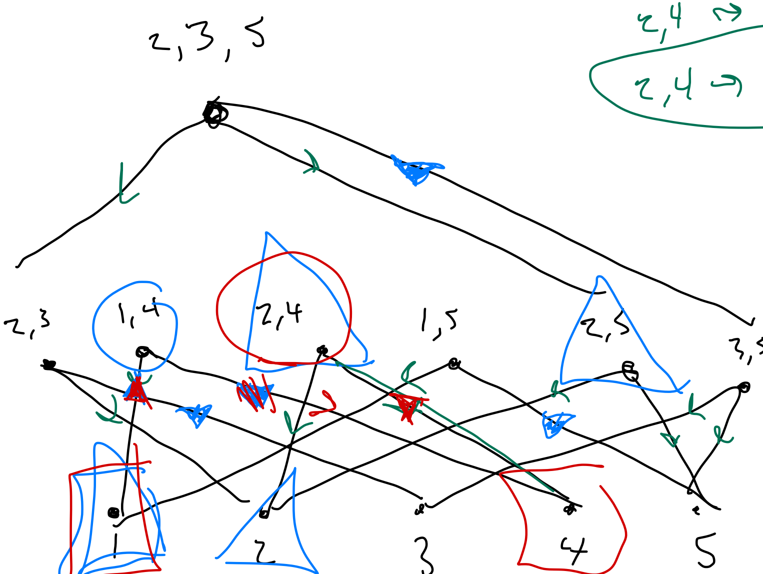

$2,4 \to 2$

$2,4 \to 4 \to 1,4 \to 1$

$2,4 \notin C$

$1 \notin C$

Cancelling, though "just" the final step, has proved to be tricky quickly.

- King does a search $\forall\, \sigma \in C_j$, which has quadratic cost.
- Simple memoization would work, except that storage would blow up.

Bad Eg:

Evil spider!

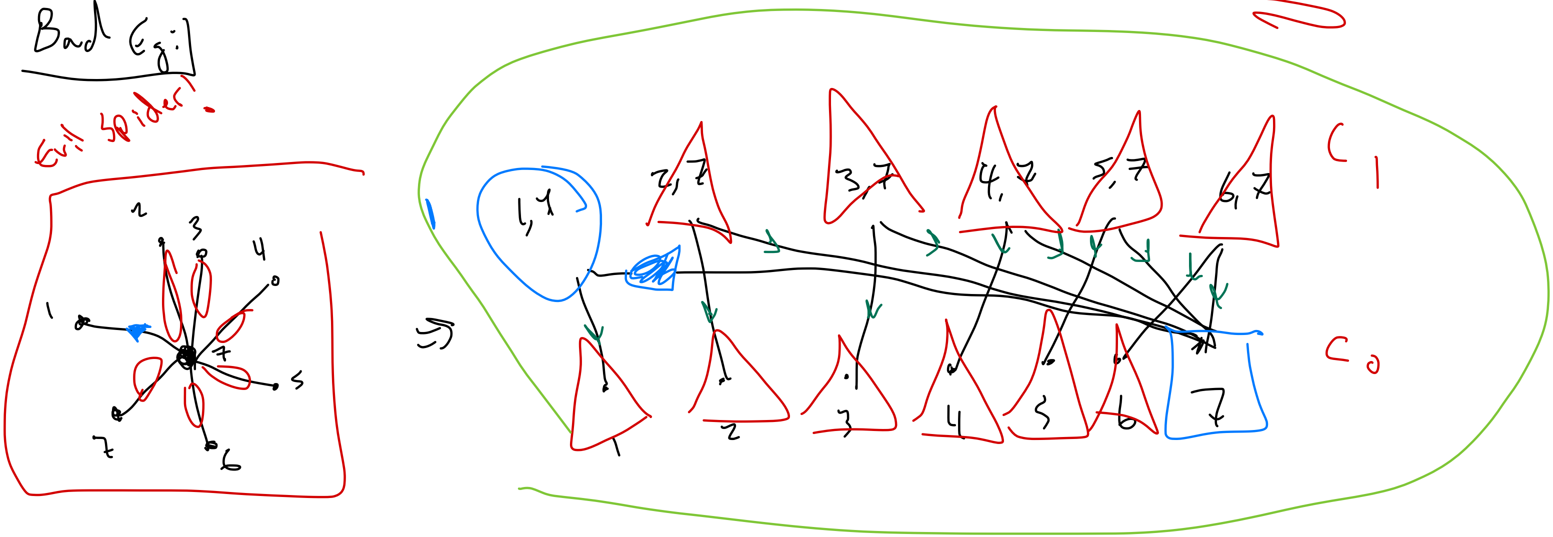

C is roughly linear!

Nasty!

What would King et al. do??

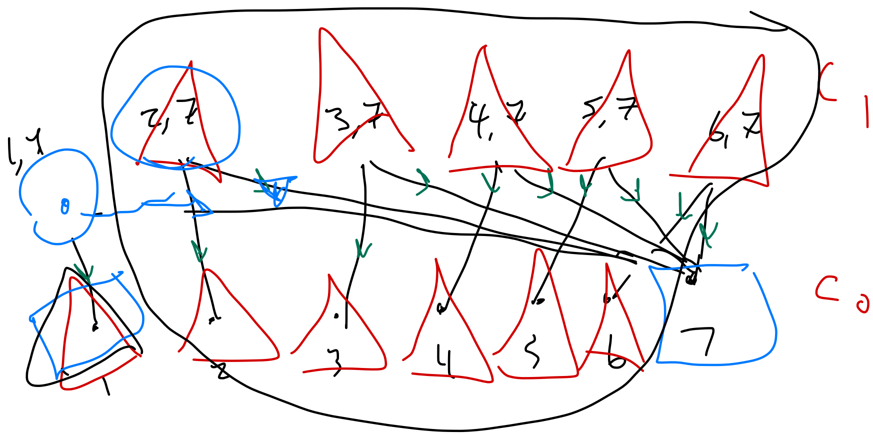

$2,7 \to 7 \to 1,7 \to 1$

$2,7 \to 2$

$3,7 \to \dots \to 2$ reverse 1

cancel 3,7 and 2

$3,7 \to 3$

$2,7 \notin C$

$1 \notin C$

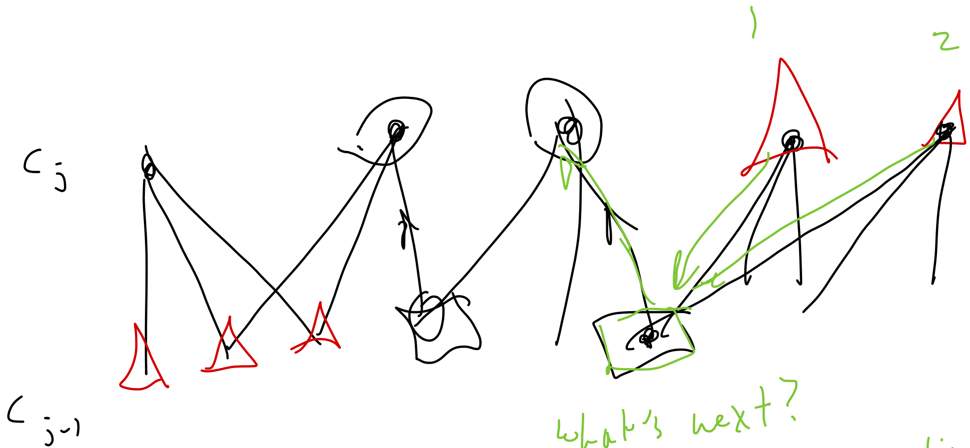

what's next?

Gradient path ⇝ proportional to n simp. long!

Task: really want subquadratic time & space

How can we cancel for these worst cases?