Discrete Morse Theory and Persistent Homology

Kevin P. Knudson

Department of Mathematics
University of Florida
kknudson@honors.ufl.edu
http://www.math.ufl.edu/∼kknudson/

February 23, 2013

Overview

Overview

Overview

Discrete Morse Theory

Let M be a simplicial complex. A *discrete Morse function* on M is a map from the set of simplices of M to $\mathbb{R}$. We abuse notation and write

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

Discrete Morse Theory

Let M be a simplicial complex. A *discrete Morse function* on M is a map from the set of simplices of M to $\mathbb{R}$. We abuse notation and write

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

It must satisfy the following two conditions, for every p-simplex $\alpha^{(p)}$ in M:

1. $\#\{\beta^{(p+1)} > \alpha^{(p)} | f(\beta) \leq f(\alpha)\} \leq 1$;
2. $\#\{\tau^{(p-1)} < \alpha^{(p)} | f(\tau) \geq f(\alpha)\} \leq 1$.

Think: Function values increase with the dimension of the simplices.

Think: Function values increase with the dimension of the simplices.

Simple example: $f : M \to \mathbb{R}$, $f(\alpha) = \dim \alpha$

Critical Points

A simplex $\alpha^{(p)}$ is *critical* if the following two conditions hold:

1. $\#\{\beta^{(p+1)} > \alpha^{(p)} | f(\beta) \leq f(\alpha)\} = 0$;
2. $\#\{\tau^{(p-1)} < \alpha^{(p)} | f(\tau) \geq f(\alpha)\} = 0$.

That is, α is critical provided f decreases when leaving α via a face, and f increases when leaving α via a coface.

Critical Points

A simplex $\alpha^{(p)}$ is *critical* if the following two conditions hold:

1. $\#\{\beta^{(p+1)} > \alpha^{(p)} | f(\beta) \leq f(\alpha)\} = 0$;
2. $\#\{\tau^{(p-1)} < \alpha^{(p)} | f(\tau) \geq f(\alpha)\} = 0$.

That is, α is critical provided f decreases when leaving α via a face, and f increases when leaving α via a coface.

A simplex that is not critical is called *regular*.

The *index* of the critical simplex α is $\dim \alpha$.

The *index* of the critical simplex α is $\dim \alpha$.

Think:

$$\begin{array}{c}\text{critical simplex } \sigma \\ \text{of dimension } i\end{array} \leftrightarrow \begin{array}{c}\text{critical point of index } i \\ \text{at barycenter of } \sigma\end{array}$$

The *index* of the critical simplex α is $\dim \alpha$.

Think:

$$\begin{array}{c} \text{critical simplex } \sigma \\ \text{of dimension } i \end{array} \leftrightarrow \begin{array}{c} \text{critical point of index } i \\ \text{at barycenter of } \sigma \end{array}$$

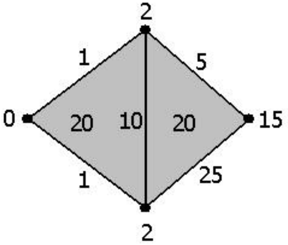

Examples

1. Critical vertex = Local minimum

Examples

1. Critical vertex = Local minimum

2. Critical n-cell = Local maximum

Examples

1. Critical vertex = Local minimum
2. Critical n-cell = Local maximum
3. For $f(\alpha) = \dim \alpha$, every cell is critical.

Here is a discrete Morse function on the circle:

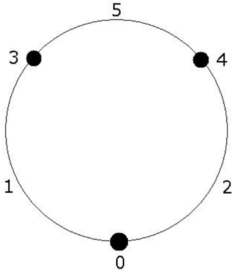

Here is a discrete Morse function on the circle:

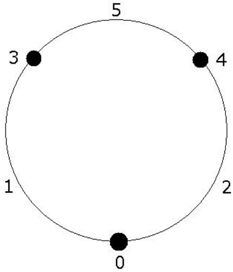

There are two critical cells, $f^{-1}(0)$ and $f^{-1}(5)$.

Here is a discrete Morse function on the torus:

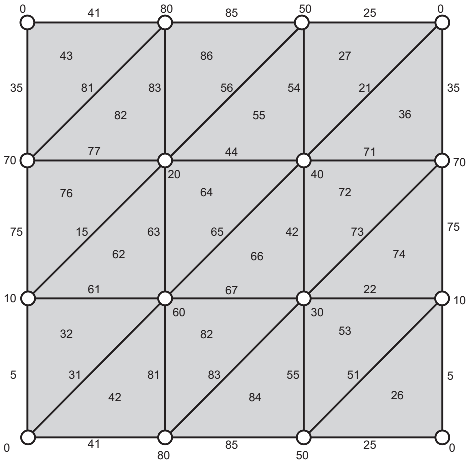

Here is a discrete Morse function on the torus:

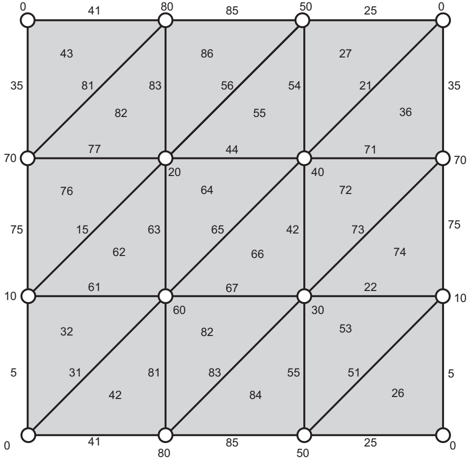

The critical cells are $f^{-1}(0)$, $f^{-1}(42)$, $f^{-1}(44)$, and $f^{-1}(86)$.

Theorem: (Forman) Suppose $f : M \to \mathbb{R}$ is a discrete Morse function. Then M is homotopy equivalent to a CW-complex with exactly one cell of dimension p for each critical simplex of dimension p.

Theorem: (Forman) Suppose $f : M \to \mathbb{R}$ is a discrete Morse function. Then M is homotopy equivalent to a CW-complex with exactly one cell of dimension p for each critical simplex of dimension p.

So, the torus has the homotopy type of a complex with one vertex, two 1-cells, and one 2-cell.

The Associated Gradient Field

Note: Regular simplices occur in pairs.

The Associated Gradient Field

Note: Regular simplices occur in pairs.

A simplex is regular if it has a face (coface) with higher (lower) value.

The Associated Gradient Field

Note: Regular simplices occur in pairs.
A simplex is regular if it has a face (coface) with higher (lower) value.
To visualize this: draw an arrow

$$\alpha^{(p)} \to \beta^{(p+1)}$$

for each such pair.
For any σ in M, exactly one of the following is true:

1. σ is the tail of exactly one arrow;
2. σ is the head of exactly one arrow;
3. σ is neither the head nor tail of an arrow.

The Associated Gradient Field

Note: Regular simplices occur in pairs.
A simplex is regular if it has a face (coface) with higher (lower) value.
To visualize this: draw an arrow

$$\alpha^{(p)} \to \beta^{(p+1)}$$

for each such pair.
For any σ in M, exactly one of the following is true:

1. σ is the tail of exactly one arrow;
2. σ is the head of exactly one arrow;
3. σ is neither the head nor tail of an arrow.

In the last case, σ is critical.

Here is the gradient field on the torus associated to the above discrete Morse function:

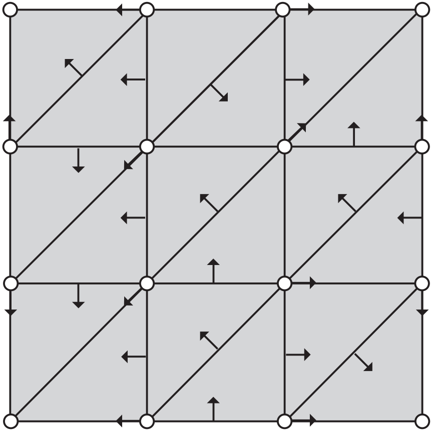

Gradients are easily characterized.

Gradients are easily characterized.

Definition. Let V be a discrete vector field on M. A *V-path* is a sequence of simplices

$$\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \ldots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)}$$

with $\{\alpha_i < \beta_i\} \in V$, $i = 0, \ldots, r$, and $\beta_i > \alpha_{i+1} \neq \alpha_i$. Such a path is called *closed* if $\alpha_{r+1} = \alpha_0$.

Gradients are easily characterized.

Definition. Let V be a discrete vector field on M. A *V-path* is a sequence of simplices

$$\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \ldots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)}$$

with $\{\alpha_i < \beta_i\} \in V$, $i = 0, \ldots, r$, and $\beta_i > \alpha_{i+1} \neq \alpha_i$. Such a path is called *closed* if $\alpha_{r+1} = \alpha_0$.

Theorem. (Forman) V is a gradient if and only if there are no closed V-paths.

Gradients are easily characterized.

Definition. Let V be a discrete vector field on M. A *V-path* is a sequence of simplices

$$\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \dots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)}$$

with $\{\alpha_i < \beta_i\} \in V$, $i = 0, \dots, r$, and $\beta_i > \alpha_{i+1} \neq \alpha_i$. Such a path is called *closed* if $\alpha_{r+1} = \alpha_0$.

Theorem. (Forman) V is a gradient if and only if there are no closed V-paths.

Another way to look at this, due to Chari, is to consider the Hasse diagram of M, the directed graph whose vertices are the simplices of M with arrows $\beta^{(p+1)} \to \alpha^{(p)}$ for $\beta > \alpha$. Modify it by reversing an arrow whenever $\{\alpha, \beta\} \in V$. Then V is a gradient if and only if this modified diagram has no directed loops.

Persistent Homology

Persistent Homology

Suppose we have a simplicial complex M and a filtration by subcomplexes

$$\emptyset = M_{-1} \subset M_0 \subset M_1 \subset \cdots \subset M_r = M.$$

Persistent Homology

Suppose we have a simplicial complex M and a filtration by subcomplexes

$$\emptyset = M_{-1} \subset M_0 \subset M_1 \subset \cdots \subset M_r = M.$$

Fix $i \geq 0$ and consider the sequence of homology groups

$$0 \to H_i(M_0) \to H_i(M_1) \to \cdots \to H_i(M_r).$$

Persistent Homology

Suppose we have a simplicial complex M and a filtration by subcomplexes

$$\emptyset = M_{-1} \subset M_0 \subset M_1 \subset \cdots \subset M_r = M.$$

Fix $i \geq 0$ and consider the sequence of homology groups

$$0 \to H_i(M_0) \to H_i(M_1) \to \cdots \to H_i(M_r).$$

If a class α is born in $H_i(M_j)$ and dies in $H_i(M_\ell)$, we say the *persistence* of α is $\ell - j - 1$.

How are classes born?

How are classes born?

A simplex σ enters the complex at filtration level j and creates a cycle. We call such a simplex *positive*.

How are classes born?

A simplex σ enters the complex at filtration level j and creates a cycle. We call such a simplex *positive*.

How do classes die?

How are classes born?

A simplex σ enters the complex at filtration level j and creates a cycle. We call such a simplex *positive*.

How do classes die?

A simplex τ enters the complex at filtration level ℓ and fills in a cycle. We call such a simplex *negative*.

How are classes born?

A simplex σ enters the complex at filtration level j and creates a cycle. We call such a simplex *positive*.

How do classes die?

A simplex τ enters the complex at filtration level ℓ and fills in a cycle. We call such a simplex *negative*.

This creates a pairing of the simplices of M.

The Persistence Pairing

The Persistence Pairing

The pairing is defined as follows:

The Persistence Pairing

The pairing is defined as follows:

1. When a negative simplex τ enters the filtration, it has killed a homology class. Find the youngest positive simplex σ representing that class.

The Persistence Pairing

The pairing is defined as follows:

1. When a negative simplex τ enters the filtration, it has killed a homology class. Find the youngest positive simplex σ representing that class.
2. If σ is not yet paired with anything, then we pair σ and τ.

The Persistence Pairing

The pairing is defined as follows:

1. When a negative simplex τ enters the filtration, it has killed a homology class. Find the youngest positive simplex σ representing that class.
2. If σ is not yet paired with anything, then we pair σ and τ.
3. If σ is already paired, we then search for a homologous cycle represented by a positive simplex which is unpaired (and we look for the youngest such simplex); call this σ'. We then pair σ' with τ.

The Persistence Pairing

The pairing is defined as follows:

1. When a negative simplex τ enters the filtration, it has killed a homology class. Find the youngest positive simplex σ representing that class.
2. If σ is not yet paired with anything, then we pair σ and τ.
3. If σ is already paired, we then search for a homologous cycle represented by a positive simplex which is unpaired (and we look for the youngest such simplex); call this σ'. We then pair σ' with τ.

Note that a simplex may get paired with something that is not adjacent to it. Also, each simplex in M is in at most one pair (some simplices do not get paired).

Here's one way to do this. Assume the filtration on M is such that each level M_i is obtained from M_{i-1} by attaching a single simplex σ_i.

Let D be the incidence matrix of the complex M:

$$D_{ij} = \begin{cases} 1 & \text{if } \sigma_i \in \partial \sigma_j \\ 0 & \text{otherwise} \end{cases}$$

Here's one way to do this. Assume the filtration on M is such that each level M_i is obtained from M_{i-1} by attaching a single simplex σ_i.

Let D be the incidence matrix of the complex M:

$$D_{ij} = \begin{cases} 1 & \text{if } \sigma_i \in \partial \sigma_j \\ 0 & \text{otherwise} \end{cases}$$

If R is a $0-1$ matrix, let $\mathrm{low}_R(j)$ be the row index of the last 1 in column j of R, and leave $\mathrm{low}_R(j)$ undefined if column j is 0.

Here's one way to do this. Assume the filtration on M is such that each level M_i is obtained from M_{i-1} by attaching a single simplex σ_i.

Let D be the incidence matrix of the complex M:

$$D_{ij} = \begin{cases} 1 & \text{if } \sigma_i \in \partial \sigma_j \\ 0 & \text{otherwise} \end{cases}$$

If R is a $0-1$ matrix, let $\mathrm{low}_R(j)$ be the row index of the last 1 in column j of R, and leave $\mathrm{low}_R(j)$ undefined if column j is 0.

Definition. R is called *reduced* and low_R is a *pairing function* if $\mathrm{low}_R(j) \neq \mathrm{low}_R(j')$ for nonzero columns $j \neq j'$.

Algorithm

Algorithm

Reduce D by adding columns mod 2 to other columns located to the right:

- $R = D$
  - for $j = 1$ to n
  - while $\exists j' < j$ with $\mathrm{low}_R(j') = \mathrm{low}_R(j)$
    - add column j' to column j
  - endwhile
  - endfor

Algorithm

Reduce D by adding columns mod 2 to other columns located to the right:

- $R = D$
  - for $j = 1$ to n
  - while $\exists j' < j$ with $\mathrm{low}_R(j') = \mathrm{low}_R(j)$
    - add column j' to column j
  - endwhile
  - endfor

Then if $i = \mathrm{low}_R(j)$, we pair σ_i with σ_j.

An Example

An Example

s^+

An Example

s^+, t^+

An Example

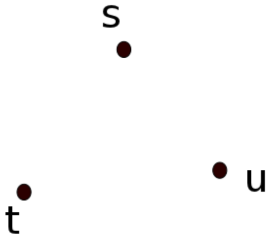

s^+, t^+, u^+

An Example

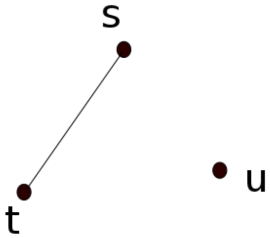

s^+, t^+, u^+, st^-

An Example

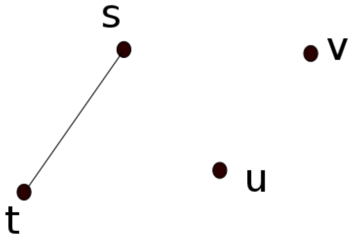

s^+, t^+, u^+, st^-, v^+

An Example

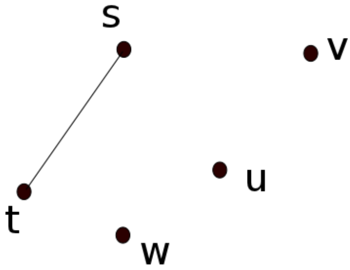

$$s^+, t^+, u^+, st^-, v^+, w^+$$

An Example

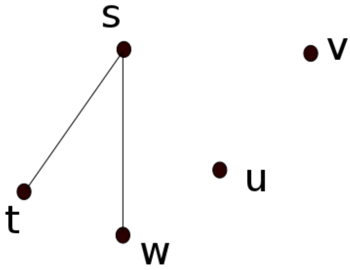

$$s^+, t^+, u^+, st^-, v^+, w^+, sw^-$$

An Example

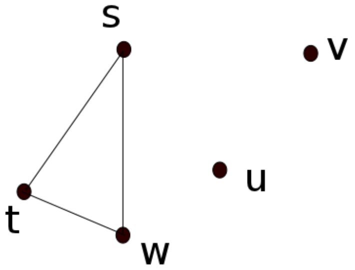

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+$

An Example

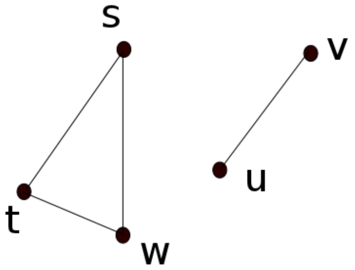

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
uv^-

An Example

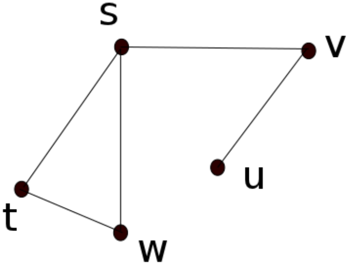

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
uv^-, sv^-

An Example

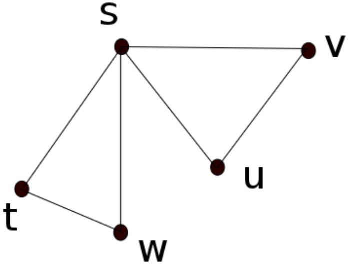

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
uv^-, sv^-, su^+

An Example

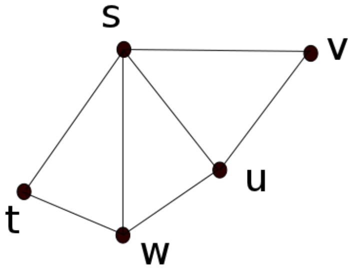

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
uv^-, sv^-, su^+, uw^+

An Example

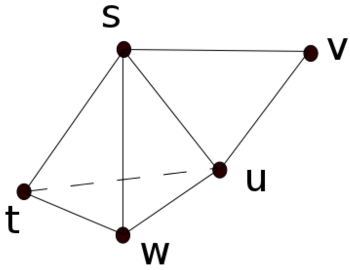

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
$uv^-, sv^-, su^+, uw^+, tu^+$

An Example

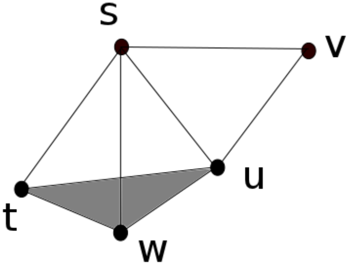

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
$uv^-, sv^-, su^+, uw^+, tu^+, tuw^-$

An Example

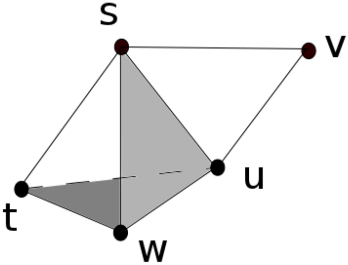

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
$uv^-, sv^-, su^+, uw^+, tu^+, tuw^-, suw^-$

An Example

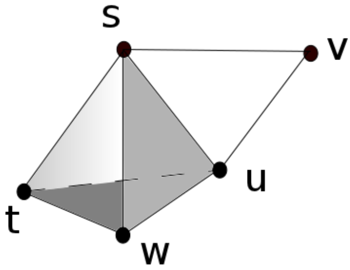

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
$uv^-, sv^-, su^+, uw^+, tu^+, tuw^-, suw^-,$
stu^-

An Example

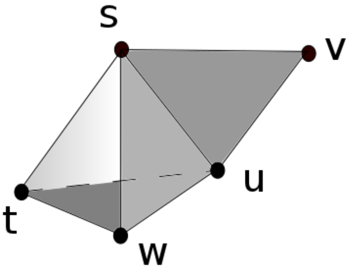

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
$uv^-, sv^-, su^+, uw^+, tu^+, tuw^-, suw^-,$
stu^-, suv^-

An Example

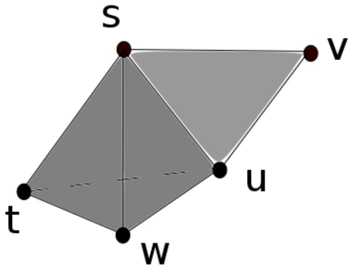

$s^+, t^+, u^+, st^-, v^+, w^+, sw^-, tw^+,$
$uv^-, sv^-, su^+, uw^+, tu^+, tuw^-, suw^-,$
stu^-, suv^-, stw^+

The resulting pairing arising from this filtration is then

$$\begin{array}{cc} (t, st) & (tu, tuw) \\ (w, sw) & (uw, suw) \\ (v, uv) & (tw, stu) \\ (u, sv) & (su, suv) \end{array}$$

The resulting pairing arising from this filtration is then

$$\begin{array}{cc} (t, st) & (tu, tuw) \\ (w, sw) & (uw, suw) \\ (v, uv) & (tw, stu) \\ (u, sv) & (su, suv) \end{array}$$

The following simplices remain unpaired: s and stw. Note that this makes sense: the Betti numbers of this complex are $1, 0, 1$.

The resulting pairing arising from this filtration is then

$$\begin{array}{cc} (t, st) & (tu, tuw) \\ (w, sw) & (uw, suw) \\ (v, uv) & (tw, stu) \\ (u, sv) & (su, suv) \end{array}$$

The following simplices remain unpaired: s and stw. Note that this makes sense: the Betti numbers of this complex are $1, 0, 1$.

Note that we wound up with (u, sv) and (tw, stu). In both cases, the lower dimensional simplex is not a face of the larger one. What that means is that it took a while for the homology classes represented by u and tw to die. This is information you might want to know.

BUT

BUT This looks like a vector field on M if we ignore these anomalies.

Persistence to Discrete Vector Fields

Persistence to Discrete Vector Fields

Given a filtered simplicial complex M and its associated persistence pairing P, we define a discrete vector field V_P on M as follows:

Persistence to Discrete Vector Fields

Given a filtered simplicial complex M and its associated persistence pairing P, we define a discrete vector field V_P on M as follows:

1. If $\{\alpha, \beta\} \in P$, and if α is a codimension-one face of β, then $\{\alpha, \beta\} \in V_P$.

Persistence to Discrete Vector Fields

Given a filtered simplicial complex M and its associated persistence pairing P, we define a discrete vector field V_P on M as follows:

1. If $\{\alpha, \beta\} \in P$, and if α is a codimension-one face of β, then $\{\alpha, \beta\} \in V_P$.
2. If $\{\alpha, \beta\} \in P$, and if α is not a face of β, then leave α and β unpaired.

Persistence to Discrete Vector Fields

Given a filtered simplicial complex M and its associated persistence pairing P, we define a discrete vector field V_P on M as follows:

1. If $\{\alpha, \beta\} \in P$, and if α is a codimension-one face of β, then $\{\alpha, \beta\} \in V_P$.
2. If $\{\alpha, \beta\} \in P$, and if α is not a face of β, then leave α and β unpaired.

Note that each pair $\{\alpha, \beta\} \in V_P$ has α positive and β negative.

Theorem. V_P is a gradient.

Theorem. V_P is a gradient.

Here's the idea. Suppose we had a loop in V_P:

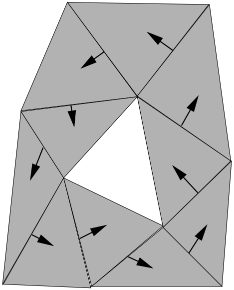

Theorem. V_P is a gradient.

Here's the idea. Suppose we had a loop in V_P:

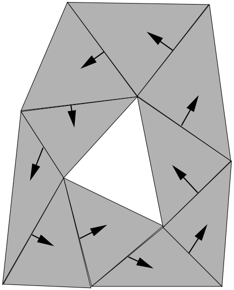

Each of the edges on the interior is positive by assumption. But that can't be—one of those edges had to join the outer loop with the inner triangle and therefore be negative. (N.B. This is a gross oversimplification, but it illustrates the point.)

An Example

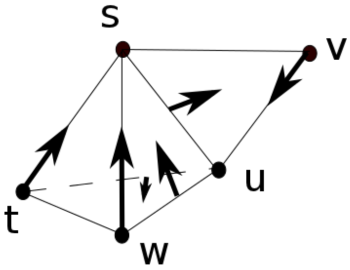

An Example

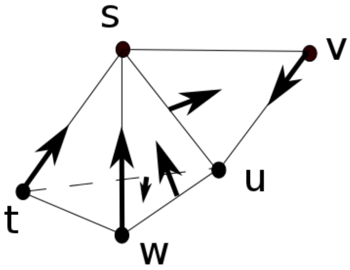

Note: *stu*, *sv*, *tw*, *u* are all critical and "shouldn't" be.

An Example

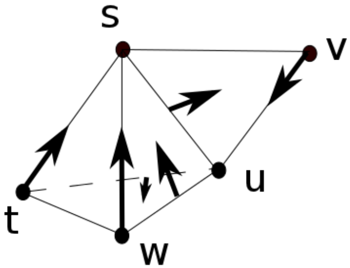

Note: *stu*, *sv*, *tw*, *u* are all critical and "shouldn't" be.
But we have unique gradient paths

$$sv > v < uv > u$$

and

$$stu > tu < tuw > tw$$

We can cancel these by reversing the arrows to yield a gradient with only s and stw critical.

We can cancel these by reversing the arrows to yield a gradient with only s and stw critical.

Conjecture. Those simplices that are paired by P but not in V_P have a unique gradient path joining them.

Discrete Vector Fields to Persistence

Discrete Vector Fields to Persistence

Conversely, given a discrete gradient V, we'd like to associate to it a filtration $M_\bullet$ of M and then compute the persistence pairing P_V.

Discrete Vector Fields to Persistence

Conversely, given a discrete gradient V, we'd like to associate to it a filtration $M_\bullet$ of M and then compute the persistence pairing P_V.

The natural thing to try is to filter M by sublevel sets of a discrete Morse function associated to V (there's a fairly canonical choice of such a map). This will lead to a filtration where more than one simplex enters at a time, but that can be dealt with. I'm still working on this.

What You Want:

What You Want: $P_{V_P} = P$

What You Want: $P_{V_P} = P$

What You Can't Have:

What You Want: $P_{V_P} = P$

What You Can't Have: $V_{P_V} = V$

What You Want: $P_{V_P} = P$

What You Can't Have: $V_{P_V} = V$
Example: Take $V = \emptyset$. Then $P_V \neq \emptyset$; in fact some vertex will get paired with one of its edges, and this pair will belong to V_{P_V}.

Future Work

Future Work

1. Nail down the proper definition of P_V.

Future Work

1. Nail down the proper definition of P_V.
2. Given any discrete vector field on M, we get a "persistence diagram" as in Edelsbrunner, Letscher, Zomorodian. Is there a characterization of those diagrams that correspond to gradients?

Future Work

1. Nail down the proper definition of P_V.
2. Given any discrete vector field on M, we get a "persistence diagram" as in Edelsbrunner, Letscher, Zomorodian. Is there a characterization of those diagrams that correspond to gradients?
3. Kozlov has characterized discrete Morse functions as poset maps with "small fibers" (fibers of cardinality ≤ 2). Since persistent homology roughly corresponds to discrete Morse theory, what is the theory associated with other poset maps? Is there a connection to multidimensional persistence?

Fin