Frontier Probability Days Tucson, Arizona

Approximate counting of connected components with random walks

Florian Sobieczky

Monday, 19th of May, 2014

This talk is dedicated to Evi Nemeth, lost at sea.

Scale Spaces

Segmentation

Number of Connected Components

Counting with Random Walks

Scale Spaces

- $G = \langle V = \mathbb{Z}^2, E = N.N. \rangle$, and $V_N = \{-L/2+1, \ldots, L\}^2, L^2 = N$
- $f \in \mathcal{S} = \{\eta : V \to R\}$, where R represents 'grey values'

Scale Spaces

- $G = \langle V = \mathbb{Z}^2, E = N.N. \rangle$, and
  $V_N = \{-L/2 + 1, \ldots, L\}^2, L^2 = N$
- $f \in \mathcal{S} = \{\eta : V \to R\}$, where R represents 'grey values'
- $\Phi_t : R^V \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'

Scale Spaces

- $G = \langle V = \mathbb{Z}^2, E = N.N. \rangle$, and $V_N = \{-L/2+1, \ldots, L\}^2, L^2 = N$
- $f \in \mathcal{S} = \{\eta : V \to R\}$, where R represents 'grey values'
- $\Phi_t : R^V \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'
- V discrete (Lindeberg 94) or continuous (Witkin 83, Koenderink 84)

Scale Spaces

- $G = \langle V = \mathbb{Z}^2, E = N.N. \rangle$, and $V_N = \{-L/2+1, \ldots, L\}^2, L^2 = N$
- $f \in \mathcal{S} = \{\eta : V \to R\}$, where R represents 'grey values'
- $\Phi_t : R^V \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'
- V discrete (Lindeberg 94) or continuous (Witkin 83, Koenderink 84)
- Idea: Splitting up information of image into different scales which label different 'derived images' according to different degree of detail (Burt 81, Crowley 81, Witkin 83)

Scale Spaces

- $G = \langle V = \mathbb{Z}^2, E = N.N.\rangle$, and $V_N = \{-L/2+1, \ldots, L\}^2, L^2 = N$
- $f \in \mathcal{S} = \{\eta : V \to R\}$, where R represents 'grey values'
- $\Phi_t : R^V \times \mathcal{S} \to \mathcal{S}$, with $t \in (0, \infty) =:$ 'Scale parameter'
- V discrete (Lindeberg 94) or continuous (Witkin 83, Koenderink 84)
- Idea: Splitting up information of image into different scales which label different 'derived images' according to different degree of detail (Burt 81, Crowley 81, Witkin 83)
- Typical Properties are: **Causality**, Linearity, Scale Invariance, Semi-group property, Isotropy, Homogeneity, ...

Comparison of two scale spaces

Figure: Comparison of $\Phi_t[f](x) := \mathbb{E}^x[f(B_\tau)]$ (top row; B_t is Brownian Motion) with GIMP's 'Selective Gaussian Blurr', where $\tau = \inf\{t > 0 | \int_0^t |f(B_s)|^2 ds > \epsilon\}$. ϵ is a scale parameter of Φ_t corresponding to the tolerance of (pathwise!) greyvalue variance. Φ_t outperforms the Selective Blurr, as seen in third column (optimal case) .

Random Walks and Image Segmentation

- Focus: Scale Spaces with Diffusions:

Random Walks and Image Segmentation

- Focus: Scale Spaces with Diffusions:
- Example: 'Gaussian Scale-Space': $\Phi_t[f] := \mathbb{E}[f(B_t)], t > 0$

Random Walks and Image Segmentation

- Focus: Scale Spaces with Diffusions:
- Example: ‘Gaussian Scale-Space’: $\Phi_t[f] := \mathbb{E}[f(B_t)], t > 0$
- Example: ‘Perona-Malik’ Model: Energy functional with Non-linear Diffusion as Euler-Lagrange equation

Random Walks and Image Segmentation

- Focus: Scale Spaces with Diffusions:
- Example: 'Gaussian Scale-Space': $\Phi_t[f] := \mathbb{E}[f(B_t)], t > 0$
- Example: 'Perona-Malik' Model: Energy functional with Non-linear Diffusion as Euler-Lagrange equation
- Example: Linear Model: 'Normalized Cuts'
  J. Shi, J. Malik: 'Normalized Cuts and Image Segmentation', IEEE PAMI, Vol. 22,20000

Random Walks and Image Segmentation

- Focus: Scale Spaces with Diffusions:
- Example: 'Gaussian Scale-Space': $\Phi_t[f] := \mathbb{E}[f(B_t)], t > 0$
- Example: 'Perona-Malik' Model: Energy functional with Non-linear Diffusion as Euler-Lagrange equation
- Example: Linear Model: 'Normalized Cuts'
  J. Shi, J. Malik: 'Normalized Cuts and Image Segmentation', IEEE PAMI, Vol. 22,20000
- Example: Grady-Model: – Labelling technique (linear, solves Bottlenecks' problem)
  L.Grady, E.L. Schwarz: 'Isoperimetric Graph partitioning for Data Clustering and Image Segmentation' PAMI, 2004

Description of the Hoshen-Kopelman Algorithm

- Hoshen-Kopelman: Phys.Rev B **14**, p 3438, 1976

Description of the Hoshen-Kopelman Algorithm

- Hoshen-Kopelman: Phys.Rev B **14**, p 3438, 1976
- Proceed row by row: (First, only *labeling*)
  - Does former pixel or pixel from last row have some greyvalue?
  - If Yes, assign same label. If No, pick new label.

Description of the Hoshen-Kopelman Algorithm

- Hoshen-Kopelman: Phys.Rev B **14**, p 3438, 1976
- Proceed row by row: (First, only *labeling*)
  - Does former pixel or pixel from last row have some greyvalue?
  - If Yes, assign same label. If No, pick new label.
- After first sweep, do second one to identify labes if same cluster.

Description of the Hoshen-Kopelman Algorithm

- Hoshen-Kopelman: Phys.Rev B **14**, p 3438, 1976
- Proceed row by row: (First, only *labeling*)
  - Does former pixel or pixel from last row have some greyvalue?
  - If Yes, assign same label. If No, pick new label.
- After first sweep, do second one to identify labes if same cluster.
- For Image of Order N there are $O(N)$ steps

Description of the Hoshen-Kopelman Algorithm

- Hoshen-Kopelman: Phys.Rev B **14**, p 3438, 1976
- Proceed row by row: (First, only *labeling*)
  - Does former pixel or pixel from last row have some greyvalue?
  - If Yes, assign same label. If No, pick new label.
- After first sweep, do second one to identify labes if same cluster.
- For Image of Order N there are $O(N)$ steps
- Drawback: Doesn't treat 'almost separated clusters' as two.

L. Grady's Model:

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x)-f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed
- Point x obtains label of exit with highest exit measure.

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed
- Point x obtains label of exit with highest exit measure.
- Advantage: L with boundary is invertible (RW properly substochastic):
  Solving a linear system, instead of computing eigenvectors

L. Grady's Model:

- In Grady's Model, graph Laplacian L is set up for a weighted lattice graph, wheights $\sim \exp(-a|f(x) - f(y)|)$, $x, y \in V$.
- Several exit points are defined (RW is 'killed' there), one for each Segment: 'Boundary of the Graph'.
- Each exit carries label.
- Instead of computing the eigenvectors, for each point $x \in V$ and time $t > 0$ the exit measure (harmonic measure) is computed
- Point x obtains label of exit with highest exit measure.
- Advantage: L with boundary is invertible (RW properly substochastic):
  Solving a linear system, instead of computing eigenvectors
- Solves 'Bottleneck' Problem.

Counting the Number of Connected Components

- Bernoulli-Percolation: *Number of open clusters per vertex* (such as: G.R.Grimmett: 'On the number of clusters in the percolation model', J.London Math.Soc. (2), 13(1076), 346-350)

Counting the Number of Connected Components

- Bernoulli-Percolation: *Number of open clusters per vertex* (such as: G.R.Grimmett: 'On the number of clusters in the percolation model', J.London Math.Soc. (2), 13(1076), 346-350)
- Eigenvector bisection techniques: (e.g. *Normalized Cuts*) Problem: High computational complexity (eigenvectors) or 'interactive'

Counting the Number of Connected Components

- Bernoulli-Percolation: *Number of open clusters per vertex* (such as: G.R.Grimmett: 'On the number of clusters in the percolation model', J.London Math.Soc. (2), 13(1076), 346-350)
- Eigenvector bisection techniques: (e.g. *Normalized Cuts*) Problem: High computational complexity (eigenvectors) or 'interactive'
- Grady's Harmonic Measure Technique:
  Problem: Pre-assigned exits - 'Interactive' approach, 'Exits' not initially defined

Counting the Number of Connected Components

- Bernoulli-Percolation: *Number of open clusters per vertex* (such as: G.R.Grimmett: 'On the number of clusters in the percolation model', J.London Math.Soc. (2), 13(1076), 346-350)
- Eigenvector bisection techniques: (e.g. *Normalized Cuts*) Problem: High computational complexity (eigenvectors) or 'interactive'
- Grady's Harmonic Measure Technique:
  Problem: Pre-assigned exits - 'Interactive' approach, 'Exits' not initially defined
- Union-Find Algorithms (HK, etc.)
  Problem: Does not detect 'vague boundaries between segments'

Feature detection on Satellite Images

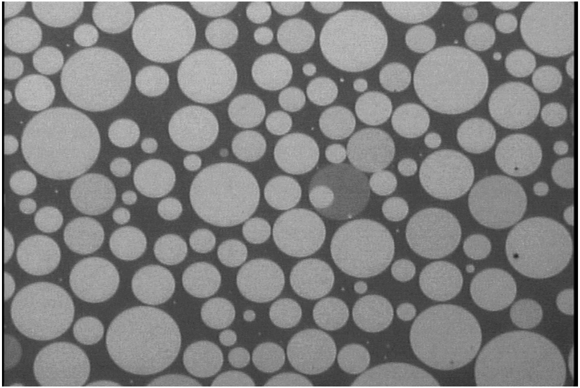

Figure: From courtesy of J.E.Maclennan: ‘Liquid Crystallography’ (Noisy Picture)

Feature detection on Satellite Images

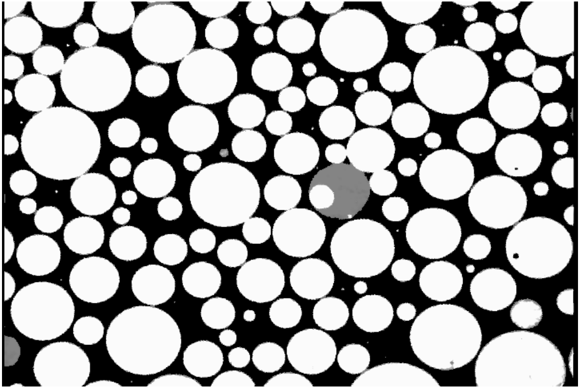

Figure: From courtesy of J.E.Maclennan: 'Liquid Crystallography' (Noise removal induces bottlenecks)

Feature detection on Satellite Images

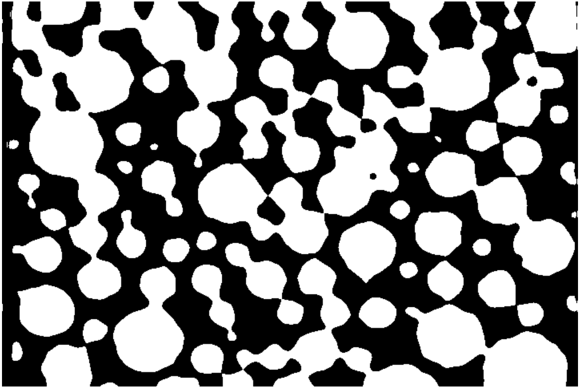

Figure: Strong Noise removal induces many new bottlenecks (Still want to find approximately same number of components)

Counting 'nearly' separate components with RW's

- General Setting: $G = \langle V, E \rangle$ and $\{V_N\}$: $V_N \subset V_{N+1} \to V$
  Let G be infinite, transitive, amenable.
  $\Omega = 2^E, \mu \in \mathcal{M}_{1,+}(\Omega)$, Aut($G$)-invariant
  $\omega \in \Omega$ and $H_N(\omega) = \langle V, \omega^{-1}(\{1\}) \rangle | V_N$.

Counting 'nearly' separate components with RW's

- General Setting: $G = \langle V, E\rangle$ and $\{V_N\}$: $V_N \subset V_{N+1} \to V$
  Let G be infinite, transitive, amenable.
  $\Omega = 2^E, \mu \in \mathcal{M}_{1,+}(\Omega)$, Aut($G$)-invariant
  $\omega \in \Omega$ and $H_N(\omega) = \langle V, \omega^{-1}(\{1\})\rangle | V_N$.
- Let # of Conn. Comp. of finite graph H_N be M_N.

Counting 'nearly' separate components with RW's

- General Setting: $G = \langle V, E \rangle$ and $\{V_N\}: \ V_N \subset V_{N+1} \to V$
  Let G be infinite, transitive, amenable.
  $\Omega = 2^E, \mu \in \mathcal{M}_{1,+}(\Omega)$, Aut($G$)-invariant
  $\omega \in \Omega$ and $H_N(\omega) = \langle V, \omega^{-1}(\{1\}) \rangle | V_N$.
- Let # of Conn. Comp. of finite graph H_N be M_N.
- $M_N = \mathrm{Tr}[\prod_{N,0}] = \lim_{t \to \infty} \mathrm{Tr}[\exp(-t\mathcal{L}_N)]$
  where $\mathcal{L}_N = \mathbb{I}_N - P_N$ with P_N delayed random walk X_t on H_N and $\prod_{N,0}$ the orthogonal projector into the invariant subspace of P_N.

Counting 'nearly' separate components with RW's

- General Setting: $G = \langle V, E \rangle$ and $\{V_N\}$: $V_N \subset V_{N+1} \to V$
Let G be infinite, transitive, amenable.
$\Omega = 2^E, \mu \in \mathcal{M}_{1,+}(\Omega)$, Aut($G$)-invariant
$\omega \in \Omega$ and $H_N(\omega) = \langle V, \omega^{-1}(\{1\}) \rangle | V_N$.
- Let # of Conn. Comp. of finite graph H_N be M_N.
- $M_N = \mathrm{Tr}[\prod_{N,0}] = \lim_{t \to \infty} \mathrm{Tr}[\exp(-t\mathcal{L}_N)]$
where $\mathcal{L}_N = \mathbb{I}_N - P_N$ with P_N delayed random walk X_t on H_N and $\prod_{N,0}$ the orthogonal projector into the invariant subspace of P_N.
- $\frac{1}{N}\mathrm{Tr}[e^{-t\mathcal{L}_N}] = \mathbb{P}[X_t = X_0 \mid X_0 \sim \mathrm{UNIF}(V)] =: \overline{P}_N(t)$
This is the return probability if the initial distribution is uniform on V_N.

Counting 'nearly' separate components with RW's

- Observe: For $G = \langle \mathbb{Z}^2, N.N. \rangle$, with $0 = \tilde{\lambda}_0 < \tilde{\lambda}_1 < \cdots < \tilde{N}_j$ the spectrum of j-th component:

Shape of jth component	Order of mag of $\tilde{\lambda}_1$ in $\tilde{N} := \tilde{N}_j$:
Circle	$\tilde{N}^{-1}$
Dumbell	$(\tilde{N} \log \tilde{N})^{-1}$

- Ask: How many 'nearly separated components' are there?

Counting 'nearly' separate components with RW's

- Observe: For $G = \langle \mathbb{Z}^2, N.N. \rangle$, with $0 = \tilde{\lambda}_0 < \tilde{\lambda}_1 < \cdots < \tilde{N}_j$ the spectrum of j-th component:

Shape of jth component	Order of mag of $\tilde{\lambda}_1$ in $\tilde{N} := \tilde{N}_j$:
Circle	$\tilde{N}^{-1}$
Dumbell	$(\tilde{N} \log \tilde{N})^{-1}$

- Choose $N_1(N)$ to be the minimum cardinality of a dumbbell's vertex-set, which should still be counted as two components. Let $\widehat{N}(N)$ be the maximal cardinality of any *connected* component. Then, choose $\epsilon_N, c > 0$ such that

$$\frac{c}{N_1 \log N_1} < \epsilon_N < \frac{1}{\widehat{N}}.$$

- Ask: How many 'nearly separated components' are there?

Counting 'nearly' separate components: Results

- **Prop. 1** If N suff. large, $M_N^{(\epsilon_N)}$, defined by

$$M_N^{(\epsilon)} \;=\; \mathrm{dim}(\mathrm{span}\{v \in R^{V_N} \mid (1-P_N)v = \lambda v \;\text{ and }\; \lambda < \epsilon\})$$

is the number of connected or nearly separated components if they can only be disks, dumbbells, or touching disks with the smaller of the two touching disks with cardinality N_1 fulfilling

$$N_1 \log N_1 \;>\; c \cdot \widehat{N}.$$

Counting 'nearly' separate components: Results

- **Prop. 1** If N suff. large, $M_N^{(\epsilon_N)}$, defined by

$$M_N^{(\epsilon)} \;=\; \mathrm{dim}(\mathrm{span}\{v \in R^{V_N} \mid (1-P_N)v = \lambda v \;\text{ and }\; \lambda < \epsilon\})$$

  is the number of connected or nearly separated components if they can only be disks, dumbbells, or touching disks with the smaller of the two touching disks with cardinality N_1 fulfilling

$$N_1 \log N_1 \;>\; c \cdot \widehat{N}.$$

- **Prop. 2** $M_N^{(\epsilon_N)} \geq N \cdot \overline{P}_N(t)$, where $t \;=\; \frac{\ln 2}{c} N_1 \ln N_1 \;\geq\; \frac{\ln 2}{\epsilon_N}$.

Counting 'nearly' separate components: Results

- **Prop. 1** If N suff. large, $M_N^{(\epsilon_N)}$, defined by

$$M_N^{(\epsilon)} \;=\; \dim(\mathrm{span}\{v \in R^{V_N} \mid (1-P_N)v = \lambda v \;\text{ and }\; \lambda < \epsilon\})$$

is the number of connected or nearly separated components if they can only be disks, dumbbells, or touching disks with the smaller of the two touching disks with cardinality N_1 fulfilling

$$N_1 \log N_1 \;>\; c \cdot \widehat{N}.$$

- **Prop. 2** $M_N^{(\epsilon_N)} \geq N \cdot \overline{P}_N(t)$, where $t \;=\; \frac{\ln 2}{c} N_1 \ln N_1 \;\geq\; \frac{\ln 2}{\epsilon_N}$.
- **Prop. 3** $M_N^{(\epsilon_N)} \leq N \cdot \overline{P}_N(t)$, where $t = \widehat{N} \ln(\widehat{N}^2/N)$.

Ideas from the Proof:

- Assume $|V_N| = N$ (for simplicity). Let:

$$\mathbb{E}_N\left[f(\overline{N}, \{\overline{\lambda}\})\right] := \frac{1}{N}\sum_{v\in V_N} f(\overline{N}(v), \{\overline{\lambda}(v)\})$$

where $\overline{N}(v) = \tilde{N}_j$ such that the j-th component with size $\tilde{N}_j$.

Ideas from the Proof:

- Assume $|V_N| = N$ (for simplicity). Let:

$$\mathbb{E}_N\left[f(\overline{N},\{\overline{\lambda}\})\right] := \frac{1}{N}\sum_{v\in V_N} f(\overline{N}(v),\{\overline{\lambda}(v)\})$$

where $\overline{N}(v) = \tilde{N}_j$ such that the j-th component with size $\tilde{N}_j$.

- Note: $N = \sum_{j=1}^{M_N} \tilde{N}_j$. Let $\tilde{p}_j = \tilde{N}_j / \sum_{i=1}^{M_N} \tilde{N}_i$,

Ideas from the Proof:

- Assume $|V_N| = N$ (for simplicity). Let:

$$\mathbb{E}_N\left[f(\overline{N},\{\overline{\lambda}\})\right] := \frac{1}{N}\sum_{v\in V_N} f(\overline{N}(v),\{\overline{\lambda}(v)\})$$

where $\overline{N}(v) = \tilde{N}_j$ such that the j-th component with size $\tilde{N}_j$.

- Note: $N = \sum_{j=1}^{M_N} \tilde{N}_j$. Let $\tilde{p}_j = \tilde{N}_j / \sum_{i=1}^{M_N} \tilde{N}_i$,
- then, with $\tilde{\mathbb{E}}_N\left[f(\tilde{N},\{\tilde{\lambda}\})\right] := \sum_{j=1}^{M_N} \tilde{p}_j f(\tilde{N}_j,\{\tilde{\lambda}_j\})$:

$$\mathbb{E}_N\left[f(\overline{N},\{\overline{\lambda}\})\right] = \tilde{\mathbb{E}}_N\left[f(\tilde{N},\{\tilde{\lambda}\})\right]$$

the expected value with respect to the size biased distribution over the space of labels of the connected components is the same as the expected value with respect to the uniform distribution over vertices in V_N.

Ideas from the Proof:

- Let $t = t_1$ if $\overline{P}_N(t) = M_N^{(\epsilon)} \;\Leftrightarrow:\; (*)$. In other words, at time t_1 the average return probability equals the # of connected or nearly separate components.

Ideas from the Proof:

- Let $t = t_1$ if $\overline{P}_N(t) \; = \; M_N^{(\epsilon)} \;\; \Leftrightarrow: \;\; (*)$. In other words, at time t_1 the average return probability equals the # of connected or nearly separate components.
- Note, $\mathcal{L}_N = \mathbb{I}_N - P_N$. Let $\prod_{<} = \mathbb{I}_N - \prod_{\geq}$ the projector onto

$$\mathrm{span}\{ \, v \in \mathbb{R}^N \mid \mathcal{L}_N \, v \; = \; \lambda v, \;\; \text{and } \lambda < \epsilon \, \}.$$

Ideas from the Proof:

- Let $t = t_1$ if $\overline{P}_N(t) = M_N^{(\epsilon)} \;\Leftrightarrow: \; (*)$. In other words, at time t_1 the average return probability equals the # of connected or nearly separate components.
- Note, $\mathcal{L}_N = \mathbb{I}_N - P_N$. Let $\prod_< = \mathbb{I}_N - \prod_\geq$ the projector onto

$$\mathrm{span}\{\, v \in \mathbb{R}^N \mid \mathcal{L}_N\, v \;=\; \lambda v, \;\text{ and } \lambda < \epsilon \,\}.$$

- Then $(*) \;\Leftrightarrow\; \mathrm{Tr}[\prod_\geq e^{-t\mathcal{L}_N}] \;=\; \mathrm{Tr}[\prod_<(\mathbb{I}_N \;-\; e^{-t\mathcal{L}_N})]$.

Ideas from the Proof:

- Let $t = t_1$ if $\overline{P}_N(t) = M_N^{(\epsilon)} \;\Leftrightarrow: \; (*)$. In other words, at time t_1 the average return probability equals the # of connected or nearly separate components.
- Note, $\mathcal{L}_N = \mathbb{I}_N - P_N$. Let $\prod_< = \mathbb{I}_N - \prod_\geq$ the projector onto

$$\mathrm{span}\{\, v \in \mathbb{R}^N \mid \mathcal{L}_N\, v \;=\; \lambda v, \;\text{ and } \lambda < \epsilon \,\}.$$

- Then $(*) \;\Leftrightarrow\; \mathrm{Tr}[\prod_\geq e^{-t\mathcal{L}_N}] \;=\; \mathrm{Tr}[\prod_<(\mathbb{I}_N \;-\; e^{-t\mathcal{L}_N})]$.
- The function on the left is decreasing in t, the one on the right increasing. Replacing the LHS with an upper bound and the RHS with a lower bound will lead to an equation which will only be solved for a greater t, yielding an upper bound for t_1.

Ideas from the Proof:

- The following bounds for the LHS and RHS are available: Let $\epsilon = \epsilon_N := \frac{c}{\widehat{N} \ln \widehat{N}}$ for some $c > 0$, and $\widehat{N} := \max_j \tilde{N}_j)$

Bound (↓) - Side(→)	LHS	RHS
Upper	$e^{-t\epsilon}$	$1 - e^{-t\epsilon}$
Lower	$\frac{\widehat{N}}{N} e^{-t/\widehat{N}}$	$1/\widehat{N}$

Ideas from the Proof:

- The following bounds for the LHS and RHS are available:
  Let $\epsilon = \epsilon_N := \frac{c}{\widehat{N} \ln \widehat{N}}$ for some $c > 0$, and $\widehat{N} := \max_j \tilde{N}_j)$

Bound (↓) - Side(→)	LHS	RHS
Upper	$e^{-t\epsilon}$	$1 - e^{-t\epsilon}$
Lower	$\frac{\widehat{N}}{N} e^{-t/\widehat{N}}$	$1/\widehat{N}$

- So: $\exists_{t>t_1} \quad e^{-t\epsilon} = 1 - e^{-t\epsilon} \Rightarrow t_1 < \frac{\ln 2}{\epsilon} \leq \frac{\ln 2}{c} N_1 \ln N_1$

Ideas from the Proof:

- The following bounds for the LHS and RHS are available: Let $\epsilon = \epsilon_N := \frac{c}{\widehat{N} \ln \widehat{N}}$ for some $c > 0$, and $\widehat{N} := \max_j \tilde{N}_j)$

Bound (↓) - Side(→)	LHS	RHS
Upper	$e^{-t\epsilon}$	$1 - e^{-t\epsilon}$
Lower	$\frac{\widehat{N}}{N} e^{-t/\widehat{N}}$	$1/\widehat{N}$

- So: $\exists_{t > t_1} \quad e^{-t\epsilon} = 1 - e^{-t\epsilon} \Rightarrow t_1 < \frac{\ln 2}{\epsilon} \leq \frac{\ln 2}{c} N_1 \ln N_1$
- Similarily, a lower bound for t_1 can be obtained by using the monotonicity of LHS and RHS in t in the reversed way: $\exists_{t < t_1} \quad \frac{\widehat{N}}{N} e^{-t/\widehat{N}} = \frac{1}{\widehat{N}} \Rightarrow t_1 > \widehat{N} \ln(\widehat{N}^2/N)$

Discussion

- When are these bounds tight?

Discussion

- When are these bounds tight?
- Note $N_1 \leq \widehat{N}$.

Discussion

- When are these bounds tight?
- Note $N_1 \leq \widehat{N}$.
- Also: $\widehat{N}\ln(\widehat{N}^2/N) \;=\; \widehat{N}(\ln\widehat{N} \;-\; \ln(N/\widehat{N}))$
  Therefore, criterion for tightness:

$$\ln(N/\widehat{N}) \;=\; o(\ln\widehat{N})$$

Discussion

- When are these bounds tight?
- Note $N_1 \leq \widehat{N}$.
- Also: $\widehat{N} \ln(\widehat{N}^2/N) \;=\; \widehat{N}(\ln \widehat{N} \;-\; \ln(N/\widehat{N}))$
  Therefore, criterion for tightness:

$$\ln(N/\widehat{N}) \;=\; o(\ln \widehat{N})$$

- Example: $\widehat{N} \;\sim\; \sqrt{N} \ln N$
  i.e. the largest component must be 'slightly larger' than $\sqrt{N}$.

Thank you for a great conference!

Figure: From Saguaro National Park. See more pictures here: http://web.cs.du.edu/~sobieczk/Tucson