Random Spanning Forests on Graphs for Fast Laplacian-Based Computations

Nicolas Tremblay

(joint work with P.-O. Amblard, L. Avena, S. Barthelmé, A. Gaudillière, Y. Pilavci)

CNRS, GIPSA-lab, Univ. Grenoble-Alpes, Grenoble-INP, France







Contents

A preamble: why am I here?

- Part I: Laplacian-based computations?
- Part II: A few basics on random spanning forests
- Part III: Inverse trace via random spanning forests
- ► Part IV: Extensions, ongoing work



Consider a connected graph with *n* nodes and *m* edges:



A spanning tree is a subset of edges that i/ are connected, ii/ do not contain cycles, iii/ span all nodes.



- A spanning tree is a subset of edges that i/ are connected, ii/ do not contain cycles, iii/ span all nodes.
- A graph of size *n* contains between 1 and n^{n-2} different spanning trees



- A spanning tree is a subset of edges that i/ are connected, ii/ do not contain cycles, iii/ span all nodes.
- A graph of size *n* contains between 1 and n^{n-2} different spanning trees
- A uniform spanning tree (UST) is a spanning tree, drawn uniformly at random from the set of all spanning trees.



- A spanning tree is a subset of edges that i/ are connected, ii/ do not contain cycles, iii/ span all nodes.
- A graph of size *n* contains between 1 and n^{n-2} different spanning trees
- A uniform spanning tree (UST) is a spanning tree, drawn uniformly at random from the set of all spanning trees.
- Classical result: A UST is a (projective) DPP defined over the set of edges

To give you the flavor of this talk

- DPPs can be understood as (weakly) repulsive summaries of a ground set of elements
- A general ML/SP question is: what information is kept by the sampling process? What can be estimated about the original set from this summary?

To give you the flavor of this talk

- DPPs can be understood as (weakly) repulsive summaries of a ground set of elements
- A general ML/SP question is: what information is kept by the sampling process? What can be estimated about the original set from this summary?
- Given that USTs (and their generalisation Random Spanning Forests) are DPPs, the question translates as:



Contents

Part I: Laplacian-based computations?

- Part II: A few basics on random spanning forests
- Part III: Inverse trace via random spanning forests
- ▶ Part IV: Extensions, ongoing work



































$$\boldsymbol{y} = [\boldsymbol{y}_1, \boldsymbol{y}_2, \boldsymbol{y}_3, \boldsymbol{y}_4]^\top \in \mathbb{R}^4$$

Different operations on graphs or signals defined on graphs

on graphs

- community detection, un-, semi-, supervised learning
- control of speed of diffusion (information, disease)
- graph simplification (sparsification, node merging, etc.)
- graph visualization

on signals defined on graphs

- Tikhonov/wavelet denoising
- signal compression / sampling
- inpainting / interpolation
- graph filtering (e.g. in GNNs)
- + many other examples from many different fields of research...
- \implies Some of these operations are based on the Laplacian matrix





Formally, we assume that:

- There is an underlying signal x(t)
- We have noisy observations for a discrete set of time samples:

 $y_i = x(t_i) + \epsilon_i$

• $\epsilon_i \sim N(0, \sigma^2)$ (white Gaussian noise)

We seek to estimate x from the noisy observations y

• Consider the following estimator (for q > 0)

$$\hat{x} = \underset{z}{\operatorname{argmin}} \quad q \; \frac{1}{2} \sum_{i=1}^{n} (z(t_i) - y(t_i))^2 + \frac{1}{2} \int_{\Omega} \left(\frac{d}{dt} z(t) \right)^2 dt$$

- $\sum_{i=1}^{n} (z(t_i) y(t_i))^2$ is a *data-fidelity* term
- $\int_{\Omega} (\frac{d}{dt} z(t))^2 dt$ is a smoothness penalty
- q sets the trade-off between data fidelity and smoothness

$$\hat{x} = \underset{z}{\operatorname{argmin}} \quad q \; \frac{1}{2} \sum_{i=1}^{n} (z(t_i) - y(t_i))^2 + \frac{1}{2} \int_{\Omega} \left(\frac{d}{dt} z(t) \right)^2 dt$$



Observations y

N. Tremblay & co. Fast Graph Computations via Random Spanning Forests Lyon, June 2022 12 / 35

$$\hat{x} = \underset{z}{\operatorname{argmin}} \quad q \; \frac{1}{2} \sum_{i=1}^{n} (z(t_i) - y(t_i))^2 + \frac{1}{2} \int_{\Omega} \left(\frac{d}{dt} z(t) \right)^2 dt$$



For a given q > 0

$$\hat{x} = \underset{z}{\operatorname{argmin}} \quad q \; \frac{1}{2} \sum_{i=1}^{n} (z(t_i) - y(t_i))^2 + \frac{1}{2} \int_{\Omega} \left(\frac{d}{dt} z(t) \right)^2 dt$$



For a large q > 0

$$\hat{x} = \underset{z}{\operatorname{argmin}} \quad q \; \frac{1}{2} \sum_{i=1}^{n} (z(t_i) - y(t_i))^2 + \frac{1}{2} \int_{\Omega} \left(\frac{d}{dt} z(t) \right)^2 dt$$



For a small q > 0

$$\hat{x} = \underset{z}{\operatorname{argmin}} \quad q \; \frac{1}{2} \sum_{i=1}^{n} (z(t_i) - y(t_i))^2 + \frac{1}{2} \int_{\Omega} \left(\frac{d}{dt} z(t) \right)^2 dt$$



For $q
ightarrow 0^+$

There is a simple way of estimating the smoothness penalty term when the observations fall on a regular grid:

$$\int_{\Omega} (\frac{d}{dt}z(t))^2 dt \simeq \sum_{t=2}^n (z_t - z_{t-1})^2 = \frac{1}{2}z^t \mathsf{L}z$$

where $L \in \mathbb{R}^{n \times n}$ is the

- discrete Laplace operator.
- equivalently the Laplacian of the path graph:



There is a simple way of estimating the smoothness penalty term when the observations fall on a regular grid:

$$\int_{\Omega} (\frac{d}{dt}z(t))^2 dt \simeq \sum_{t=2}^n (z_t - z_{t-1})^2 = \frac{1}{2} z^t \mathsf{L} z$$

where
$$L \in \mathbb{R}^{n \times n}$$
 is the

- discrete Laplace operator.
- equivalently the Laplacian of the path graph:
- One obtains the so-called Tikhonov regularisation problem:

$$\hat{\mathbf{x}} = \underset{\mathbf{z}}{\operatorname{argmin}} q \frac{1}{2} \|\mathbf{z} - \mathbf{y}\|^2 + \frac{1}{2} \mathbf{z}^t \mathsf{L} \mathbf{z},$$

which exact solution reads

$$\hat{\pmb{x}} = \mathsf{K}\pmb{y}$$
 with $\mathsf{K} = (\pmb{q}\mathsf{I} + \mathsf{L})^{-1}\pmb{q}\mathsf{I} \in \mathbb{R}^{n imes n}$



Motivation I. Signal denoising: extension to arbitrary graphs



The graph Tikhonov regularization problem is similarly defined as

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{z}}{\operatorname{argmin}} \quad q \; \frac{1}{2} \, \|\boldsymbol{z} - \boldsymbol{y}\|^2 + \frac{1}{2} \boldsymbol{z}^t \boldsymbol{L} \boldsymbol{z}.$$

Note that the smoothness penalty term verifies $\frac{1}{2}z^t L z = \sum_{i,j} A_{ij}(z_i - z_j)^2$.

The analytic solution reads: $\hat{x} = Ky$ with $K = (qI + L)^{-1}qI \in \mathbb{R}^{n \times n}$.

SSL: given a few pre-labeled nodes of a graph, infer communities



SSL: given a few pre-labeled nodes of a graph, infer communities



• Let y_{ℓ} encode the information *a priori* known:

$$y_{\ell}(i) = \begin{cases} 1 & \text{ if node } i \text{ is known to belong to class } \ell \\ 0 & \text{ if not.} \end{cases}$$

SSL: given a few pre-labeled nodes of a graph, infer communities



• Let y_{ℓ} encode the information *a priori* known:

$$y_\ell(i) = egin{cases} 1 & ext{ if node } i ext{ is known to belong to class } \ell \ 0 & ext{ if not.} \end{cases}$$

• A classical SSL algo looks for a labeling function f_{ℓ} that is both (with a trade-off parameter $\mu > 0$):

- smooth on the graph $(f_{\ell}^{\top} L f_{\ell}$ is small)
- \blacktriangleright close to y_{ℓ}

SSL: given a few pre-labeled nodes of a graph, infer communities



• Let y_{ℓ} encode the information *a priori* known:

$$y_{\ell}(i) = \begin{cases} 1 & \text{if node } i \text{ is known to belong to class } \ell \\ 0 & \text{if not.} \end{cases}$$

- A classical SSL algo looks for a labeling function f_{ℓ} that is both (with a trade-off parameter $\mu > 0$):
 - smooth on the graph $(f_{\ell}^{\top} L f_{\ell}$ is small)
 - \blacktriangleright close to y_{ℓ}
- the larger $f_{\ell}(i)$, the higher the chance that *i* belongs to class ℓ .

The solution may be written (in the usual SSL formalism) as

$$\mathbf{f}_{\ell} = \frac{\mu}{2+\mu} \left(\mathsf{I} - \frac{2}{2+\mu} \mathsf{D}^{-1} \mathsf{W} \right)^{-1} \mathbf{y}_{\ell}.$$

Can be re-written as:

$$f_{\ell} = K y_{\ell}$$

where $K = (Q + L)^{-1} Q$ and $Q = \frac{\mu}{2} D$ is diagonal and strictly positive.

This generalizes the previous case where we needed

$$\hat{\mathbf{x}} = \mathsf{K}\mathbf{y}$$
 with $\mathsf{K} = (q\mathsf{I} + \mathsf{L})^{-1}q\mathsf{I}$.

Motivation III. Degrees of freedom of the linear smoother $K = (Q + L)^{-1}Q$

- $\hat{x} = Ky$ is one instance of the more general class of *linear smoothers*
- a useful quantity in this context is its so-called *effective degrees of freedom* defined as

$$\delta = \operatorname{Tr}(\mathsf{K}) = \operatorname{Tr}((\mathsf{Q} + \mathsf{L})^{-1}\mathsf{Q})$$

this quantity is useful for instance for optimal choice of parameter q in Tikhonov's regularization (using AIC, GCV, etc.)
Definition Let L[†] be the Moore-Penrose pseudo-inverse of L. The effective resistance of any edge $e = (i, j) \in \mathcal{E}$ reads

$$r_e = \mathsf{L}_{ii}^\dagger + \mathsf{L}_{jj}^\dagger - 2\mathsf{L}_{ij}^\dagger > 0.$$

Definition Let L[†] be the Moore-Penrose pseudo-inverse of L. The effective resistance of any edge $e = (i, j) \in \mathcal{E}$ reads

$$r_e = \mathsf{L}_{ii}^\dagger + \mathsf{L}_{jj}^\dagger - 2\mathsf{L}_{ij}^\dagger > 0.$$

Property $\forall e \in \mathcal{E}, p_e \propto A_e r_e$ defines a probability distribution over the edges.

Definition Let L[†] be the Moore-Penrose pseudo-inverse of L. The effective resistance of any edge $e = (i, j) \in \mathcal{E}$ reads

$$r_e = \mathsf{L}_{ii}^\dagger + \mathsf{L}_{jj}^\dagger - 2\mathsf{L}_{ij}^\dagger > 0.$$

Property $\forall e \in \mathcal{E}, p_e \propto A_e r_e$ defines a probability distribution over the edges. *What for*? For instance: spectral sparsification.

Definition Let L[†] be the Moore-Penrose pseudo-inverse of L. The effective resistance of any edge $e = (i, j) \in \mathcal{E}$ reads

$$r_e = \mathsf{L}_{ii}^\dagger + \mathsf{L}_{jj}^\dagger - 2\mathsf{L}_{ij}^\dagger > 0.$$

Property $\forall e \in \mathcal{E}, p_e \propto A_e r_e$ defines a probability distribution over the edges.

What for? For instance: spectral sparsification.

Definition Let $\epsilon \in [0, 1[$. Let \mathcal{G} be a graph and L its Laplacian. Let \mathcal{H} be a subgraph of \mathcal{G} , of same size but hopefully much sparser, and X its Laplacian. Then \mathcal{H} is a ϵ -spectral sparsifier of \mathcal{G} if

$$(1-\epsilon)\mathsf{L} \preceq \mathsf{X} \preceq (1+\epsilon)\mathsf{L}.$$

Definition Let L[†] be the Moore-Penrose pseudo-inverse of L. The effective resistance of any edge $e = (i, j) \in \mathcal{E}$ reads

$$r_e = \mathsf{L}_{ii}^\dagger + \mathsf{L}_{jj}^\dagger - 2\mathsf{L}_{ij}^\dagger > 0.$$

Property $\forall e \in \mathcal{E}, p_e \propto A_e r_e$ defines a probability distribution over the edges.

What for? For instance: spectral sparsification.

Definition Let $\epsilon \in [0, 1[$. Let \mathcal{G} be a graph and L its Laplacian. Let \mathcal{H} be a subgraph of \mathcal{G} , of same size but hopefully much sparser, and X its Laplacian. Then \mathcal{H} is a ϵ -spectral sparsifier of \mathcal{G} if

$$(1-\epsilon)\mathsf{L} \preceq \mathsf{X} \preceq (1+\epsilon)\mathsf{L}.$$

Theorem Let $\epsilon \in]0, 1[$. Sample *m* edges iid with replacement according to *p*. The obtained subgraph (with properly re-weighted edges) is a spectral sparsifier with high probability provided

$$m \geq \mathcal{O}\left(\epsilon^{-2}n\log n\right).$$

N. Tremblay & co. Fast Graph Computations via Random Spanning Forests Lyon, June 2022 18 / 35

A (partial) list of Laplacian-based computation problems

Problem 1 Compute the regularized inverse

$$\hat{x} = Ky$$
 with

$$\begin{cases}
K = (Q + L)^{-1}Q \\
Q \text{ diagonal and positive}
\end{cases}$$

Problem 2 Compute the degrees of freedom of the linear smoother K:

$$\delta = \operatorname{Tr}\left((\mathsf{Q} + \mathsf{L})^{-1}\mathsf{Q}\right)$$

Problem 3 Compute the effective resistance $r_e = L_{ii}^{\dagger} + L_{ii}^{\dagger} - 2L_{ii}^{\dagger}$ of all edges

Many other L-based problems partial spectral information for spectral clustering, graph cuts, etc.

N. Tremblay & co. Fast Graph Computations via Random Spanning Forests Lyon, June 2022 19 / 35

In this talk, we will discuss:

Problem Compute the degrees of freedom of the linear smoother K:

$$\delta = \mathsf{Tr}\left((\mathsf{Q} + \mathsf{L})^{-1}\mathsf{Q}\right)$$

- ▶ to simplify, in this talk, we only consider the case where L is the combinatorial Laplacian of a connected graph and Q = qI (q > 0)
- however, keep in mind that the techniques developed can be extended to:
 - any positive diagonal matrix Q
 - any choice of Laplacian, e.g., the degree-normalized Laplacian matrix L:

$$\hat{\mathbf{x}} = (\mathbf{Q} + \mathcal{L})^{-1} \mathbf{Q} \mathbf{y} = \left(\mathbf{Q} + \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}\right)^{-1} \mathbf{Q} \mathbf{y}$$
$$= \mathbf{D}^{1/2} \left(\mathbf{Q}' + \mathbf{L}\right)^{-1} \mathbf{Q}' \mathbf{D}^{-1/2} \mathbf{y}$$

- in fact, L may even be any diagonally dominant matrix, via a standard trick involving two graphs
- It so happens that Yigit Pilavci will discuss some other problems (graph Tikhonov regularisation) this afternoon at the MLSP seminar

N. Tremblay & co. Fast Graph Computations via Random Spanning Forests Lyon, June 2022 20 / 35

Computational challenges

Goal Compute $\delta = \text{Tr}((ql + L)^{-1}q)$ for large graphs (say $n > 10^6$)

 $\begin{array}{c} {\rm Difficulty}\\ {\rm Need \ to \ compute \ an \ inverse, \ with \ (naive) \ cost \ } {\mathcal O}(n^3) \end{array}$

State-of-the-art

For sparse graphs, a combination of Hutchinson's estimator and iterative methods such as Preconditioned Conjugate Gradients (effective to approximate efficiently $(qI + L)^{-1}qy$ given a vector y) is state-of-the-art

In this work

We formulate (very different) Monte Carlo estimators based on random spanning forests

Contents

- Part I: Laplacian-based computations?
- Part II: A few basics on random spanning forests
- Part III: Inverse trace via random spanning forests
- ▶ Part IV: Extensions, ongoing work







A spanning tree





A spanning tree







A spanning tree





A spanning forest





A spanning tree





A spanning forest



N. Tremblay & co. Fast Graph Computations via Random Spanning Forests Lyon, June 2022 23 / 35



A spanning tree





A spanning forest

A rooted spanning forest



Another rooted spanning forest

A few notations



Given a rooted spanning forest ϕ , we define:

- $\rho(\phi)$ its set of roots (the red nodes)
- the root function: $r_{\phi}(i)$ is the root of the tree containing node *i*

Random forests?

• Consider the following distribution on forests (q > 0):

$$\mathbb{P}(\Phi_q=\phi) \propto q^{|
ho(\phi)|} \prod_{(ij)\in \phi} \mathsf{A}_{ij}$$

where A_{ij} is the weight of the link between nodes *i* and *j*.

- Importantly, sampling from this distribution is efficient (in time roughly $\mathcal{O}(|E|/q))$ via a variant of Wilson's algorithm.
- The forest obtained, Φ_q , is called a Random Spanning Forest (RSF)



1. add Δ and connect it with all nodes with weight q



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node


- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.
- 8. Repeat 5-7 until all nodes have been visited



Output: a random spanning forest Φ_{q}

- 1. add Δ and connect it with all nodes with weight q
- 2. Start from any node and propagate a random walk until it reaches $\boldsymbol{\Delta}$
- 3. Erase its eventual loops
- 4. Keep the selected edges and add the last node before Δ to the set of roots
- 5. Start a new random walk from any node not yet visited and stop when:
 - it reaches Δ
 - it reaches any already visited node
- 6. Erase its eventual loops and keep the remaining edges
- 7. If it "died" in $\Delta,$ then add the last visited node before Δ to the set of roots.

Output: a random spanning forest Φ_{α}

8. Repeat 5-7 until all nodes have been visited



 $\mathbb{E} \text{ (number of steps)} = \text{Tr} \left[(L + ql)^{-1} (D + ql) \right]$ $\leq n + 2 \frac{|\mathcal{E}|}{q}$

Two key properties of RSFs

The matrix $K = (L + qI)^{-1} qI$ plays a central role in the theory of random forests. Among other things:

- 1. the root process $\rho(\Phi_q)$ is a determinantal point process (DPP) with marginal kernel K, implying for instance:
 - K_{ii} is the probability that node i is a root
 - Tr (K) is the expected number of roots
- 2. the probability that node *i* is rooted at root *j* in a random forest Φ_q equals

$$\mathbb{P}\left(r_{\Phi_q}(i)=j\right)=\mathsf{K}_{ij}$$

Contents

- Part I: Laplacian-based computations?
- Part II: A few basics on random spanning forests
- Part III: Inverse trace via random spanning forests
- ▶ Part IV: Extensions, ongoing work

From random forests to the inverse trace

- ▶ Recall our objective: compute $Tr(K) = Tr[(qI + L)^{-1}qI]$
- Recall that Tr(K) is the expected number of roots of a RSF Φ_q
- A trivial RSF-based Monte-Carlo estimator of Tr (K) is thus:
 - draw N independent RSFs
 - compute the average of their number of roots
 - ▶ it is unbiased and its per-sample variance is $\sum_{i=1}^{n} \frac{q\lambda_i}{(q+\lambda_i)^2}$

From random forests to the inverse trace

- ▶ Recall our objective: compute $Tr(K) = Tr[(qI + L)^{-1}qI]$
- Recall that Tr(K) is the expected number of roots of a RSF Φ_q
- A trivial RSF-based Monte-Carlo estimator of Tr (K) is thus:
 - draw N independent RSFs
 - compute the average of their number of roots
 - ▶ it is unbiased and its per-sample variance is $\sum_{i=1}^{n} \frac{q\lambda_i}{(q+\lambda_i)^2}$
- State-of-the-art: Hutchinson's (aka Girard's) estimator:
 - draw *N* independent random vectors *x* verifying $\mathbb{E}(xx^{\top}) = I_n$
 - compute the average $\frac{1}{N} \sum_{i=1}^{N} x^{\top} K x$ (via sparse Cholesky, PCG or poly approx to compute K x)
 - ► unbiased: $\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N} \mathbf{x}^{\top}\mathbf{K}\mathbf{x}\right] = \mathbb{E}\left(\mathbf{x}^{\top}\mathbf{K}\mathbf{x}\right) = \mathbb{E}\left(\mathsf{Tr}\left(\mathsf{K}\mathbf{x}\mathbf{x}^{\top}\right)\right) = \mathsf{Tr}(\mathsf{K})$

▶ per-sample variance in case of Gaussian entries is $\sum_{i=1}^{n} \frac{2q^2}{(q+\lambda_i)^2}$

From random forests to the inverse trace

- Recall our objective: compute $Tr(K) = Tr[(qI + L)^{-1}qI]$
- Recall that Tr(K) is the expected number of roots of a RSF Φ_q
- A trivial RSF-based Monte-Carlo estimator of Tr (K) is thus:
 - draw N independent RSFs
 - compute the average of their number of roots
 - ▶ it is unbiased and its per-sample variance is $\sum_{i=1}^{n} \frac{q\lambda_i}{(q+\lambda_i)^2}$
- State-of-the-art: Hutchinson's (aka Girard's) estimator:
 - draw N independent random vectors x verifying $\mathbb{E}(xx^{\top}) = I_n$
 - compute the average $\frac{1}{N} \sum_{i=1}^{N} x^{\top} K x$ (via sparse Cholesky, PCG or poly approx to compute K x)
 - unbiased: $\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N} \mathbf{x}^{\top} \mathbf{K} \mathbf{x}\right] = \mathbb{E}\left(\mathbf{x}^{\top} \mathbf{K} \mathbf{x}\right) = \mathbb{E}\left(\mathsf{Tr}\left(\mathsf{K} \mathbf{x} \mathbf{x}^{\top}\right)\right) = \mathsf{Tr}\left(\mathsf{K}\right)$

▶ per-sample variance in case of Gaussian entries is $\sum_{i=1}^{n} \frac{2q^2}{(q+\lambda_i)^2}$

- for large/small enough q, RSF's per-sample variance is better. The precise interval is spectrum-dependent
- ▶ both SOTA and RSF-based estimator run in time $\mathcal{O}(|\mathcal{E}|)$

Inverse trace estimation: RSF-based vs. Hutchinson/Girard, pros and cons

Notable advantages of RSF-estimator:

- ▶ very easy to implement (~ 20 lines in Julia)
- minimal memory footprint
- no preprocessing of the graph needed (in fact: no centralized knowledge of L is needed): only need the ability of running a random walk

Notable disadvantages of RSF-estimator:

- only for diagonally-dominant matrices L and q > 0
- even if theoretical complexities are comparable, the RSF-estimator cannot take advantage of over-optimized matrix-vector implementations.

- Five methods to compare...
- rf RSF-based estimator
- direct uses Julia's backslash operator (calls CHOLMOD)
- amg Algebraic Multigrid with Ruge-Stuben coarsening
- cg Conjugate Gradient with diagonal preconditioning
- cg-amg Conjugate Gradient with AMG preconditioning

rf RSF-based estimator

Five methods to compare...

- direct uses Julia's backslash operator (calls CHOLMOD)
- amg Algebraic Multigrid with Ruge-Stuben coarsening
- cg Conjugate Gradient with diagonal preconditioning
- cg-amg Conjugate Gradient with AMG preconditioning

... on five different graphs: all have around 10^4 nodes and less than 10^5 edges

rf RSF-based estimator

Five methods to compare...

- direct uses Julia's backslash operator (calls CHOLMOD)
- amg Algebraic Multigrid with Ruge-Stuben coarsening
- cg Conjugate Gradient with diagonal preconditioning
- cg-amg Conjugate Gradient with AMG preconditioning

... on five different graphs: all have around 10^4 nodes and less than 10^5 edges For fair comparison, we plot the effective runtime = average time needed per iteration × the number of iterations needed to reach a relative error of 2%

rf RSF-based estimator

Five methods to compare...

- direct uses Julia's backslash operator (calls CHOLMOD)
- amg Algebraic Multigrid with Ruge-Stuben coarsening
- cg Conjugate Gradient with diagonal preconditioning
- cg-amg Conjugate Gradient with AMG preconditioning

... on five different graphs: all have around 10^4 nodes and less than 10^5 edges For fair comparison, we plot the effective runtime = average time needed per iteration × the number of iterations needed to reach a relative error of 2%



Tr(K) / n

N. Tremblay & co. Fast Graph Computations via Random Spanning Forests Lyon, June 2022 31 / 35

Contents

- Part I: Laplacian-based computations?
- Part II: A few basics on random spanning forests
- Part III: Inverse trace via random spanning forests
- Part IV: Extensions, ongoing work

Extensions

- Variance reduction techniques (control variates, stratified sampling, etc.) can be used to further improve the RSF-based estimator
- more importantly, Wilson's algorithm can be modified to sample not only one forest at a given q for a cost of order

$$n+\frac{m}{q}$$

but the whole trajectory of (coupled) forests between q_{\min} and q_{\max} ; for a cost of order

$$n + \frac{m}{q_{\min}} + n \log \frac{d_{\max}}{q_{\min}}$$

This is very powerful if one wishes to estimate the function $q \rightarrow \text{Tr}((ql+L)^{-1}q)$ on a given range $[q_{\min}, q_{\max}]$. A precise comparison with state-of-the-art will be the object of a future publication.

Conclusion

- State-of-the-art performance for inverse trace estimation via RSFs
- Computational cost: O (|E| / q) will probably require multiscale approximations in small q to become competitive in practice in this range
- Very simple and "natural" algorithm on graphs, VERY low memory footprint, no need of centralized knowledge of the graph, no preprocessing
- Similar works (some published, some in progress in ANR JCJC GRANOLA):
 - RSFs for graph Tikhonov regularization, for effective resistance estimation, for partial spectrum estimation, etc.
 Main research question: what is the extent of the information one can retrieve from a few RSEs?
 - ii/ other similar DPPs over edges and/or nodes of graphs (hopefully "Wilson-able") for similar objectives?

More details?

A journal paper is here (IEEE TSIPN 2021):

```
https://arxiv.org/pdf/2011.10450
```

- Conference papers are here:
 - ► (GRETSI 2019):

https://arxiv.org/pdf/1905.02086

(ICASSP 2020):

https://arxiv.org/pdf/1910.07963

(EUSIPCO 2022):

https://arxiv.org/pdf/2110.07894

(GRETSI 2022):

not online yet

- Julia code publicly available
- These slides can be found on my website:

gipsa-lab.fr/~nicolas.tremblay