

Large-scale semi-supervised learning with online spectral graph sparsification

daniele.calandriello@inria.fr, alessandro.lazaric@inria.fr, michal.valko@inria.fr

Graph Learning


Draw $\mathcal{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ from \mathbb{R}^d ,
 Build the graph $\mathcal{G} = (\mathcal{X}, \mathcal{E})$ with $|\mathcal{E}| = m$,
 The weights $a_{e_{i,j}}$ encode the "distance" between nodes.

Transductive setting for Semi-Supervised Learning

There exists a label y_i for each node in \mathcal{G}
 l nodes are placed in \mathcal{S} the remaining $u = n - l$ in \mathcal{T}
 The algorithm receive the labels in \mathcal{S} and the graph \mathcal{G} and outputs a function $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}$.
 Measure \mathbf{f} prediction error over \mathcal{T}
 The graph \mathcal{G} never changes
 ↳ (e.g. in spam classification, our email corpus is fixed)
 The subset \mathcal{S} that is revealed to the algorithm is random
 ↳ (e.g. which emails the users classify as spam or ham)

Harmonic Function Solution (HFS)

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f} \in \mathbb{R}^n} \frac{1}{2}(\mathbf{f} - \mathbf{y})^\top I_S(\mathbf{f} - \mathbf{y}) + \gamma \mathbf{f}^\top L_G \mathbf{f}, \quad (1)$$



Algorithmic Stability

$$\mathcal{S} = (y_1, y_2, y_3, y_4, \dots, y_{l-1}, y_l, 0, 0, 0, 0, \dots) \rightarrow \mathbf{f}$$

$$\mathcal{S}' = (y_1, y_2, y_3, y_4, \dots, y_{l-1}, 0, y_{l+1}, 0, 0, 0, \dots) \rightarrow \mathbf{f}'$$

$$|(\mathbf{f}(\mathbf{x}) - \mathbf{y}(\mathbf{x}))^2 - (\mathbf{f}'(\mathbf{x}) - \mathbf{y}(\mathbf{x}))^2| \leq \beta.$$

Theoretical guarantees for stable transductive algorithms [1]

$$R(\tilde{\mathbf{f}}) \leq \hat{R}(\tilde{\mathbf{f}}) + \beta + \left(2\beta + \frac{c^2(l+u)}{lu}\right) \sqrt{\frac{\pi(l,u) \log(1/\delta)}{2}}$$

$$\pi(l,u) = \frac{lu}{l+u - 0.52 \max\{l,u\} - 1}$$

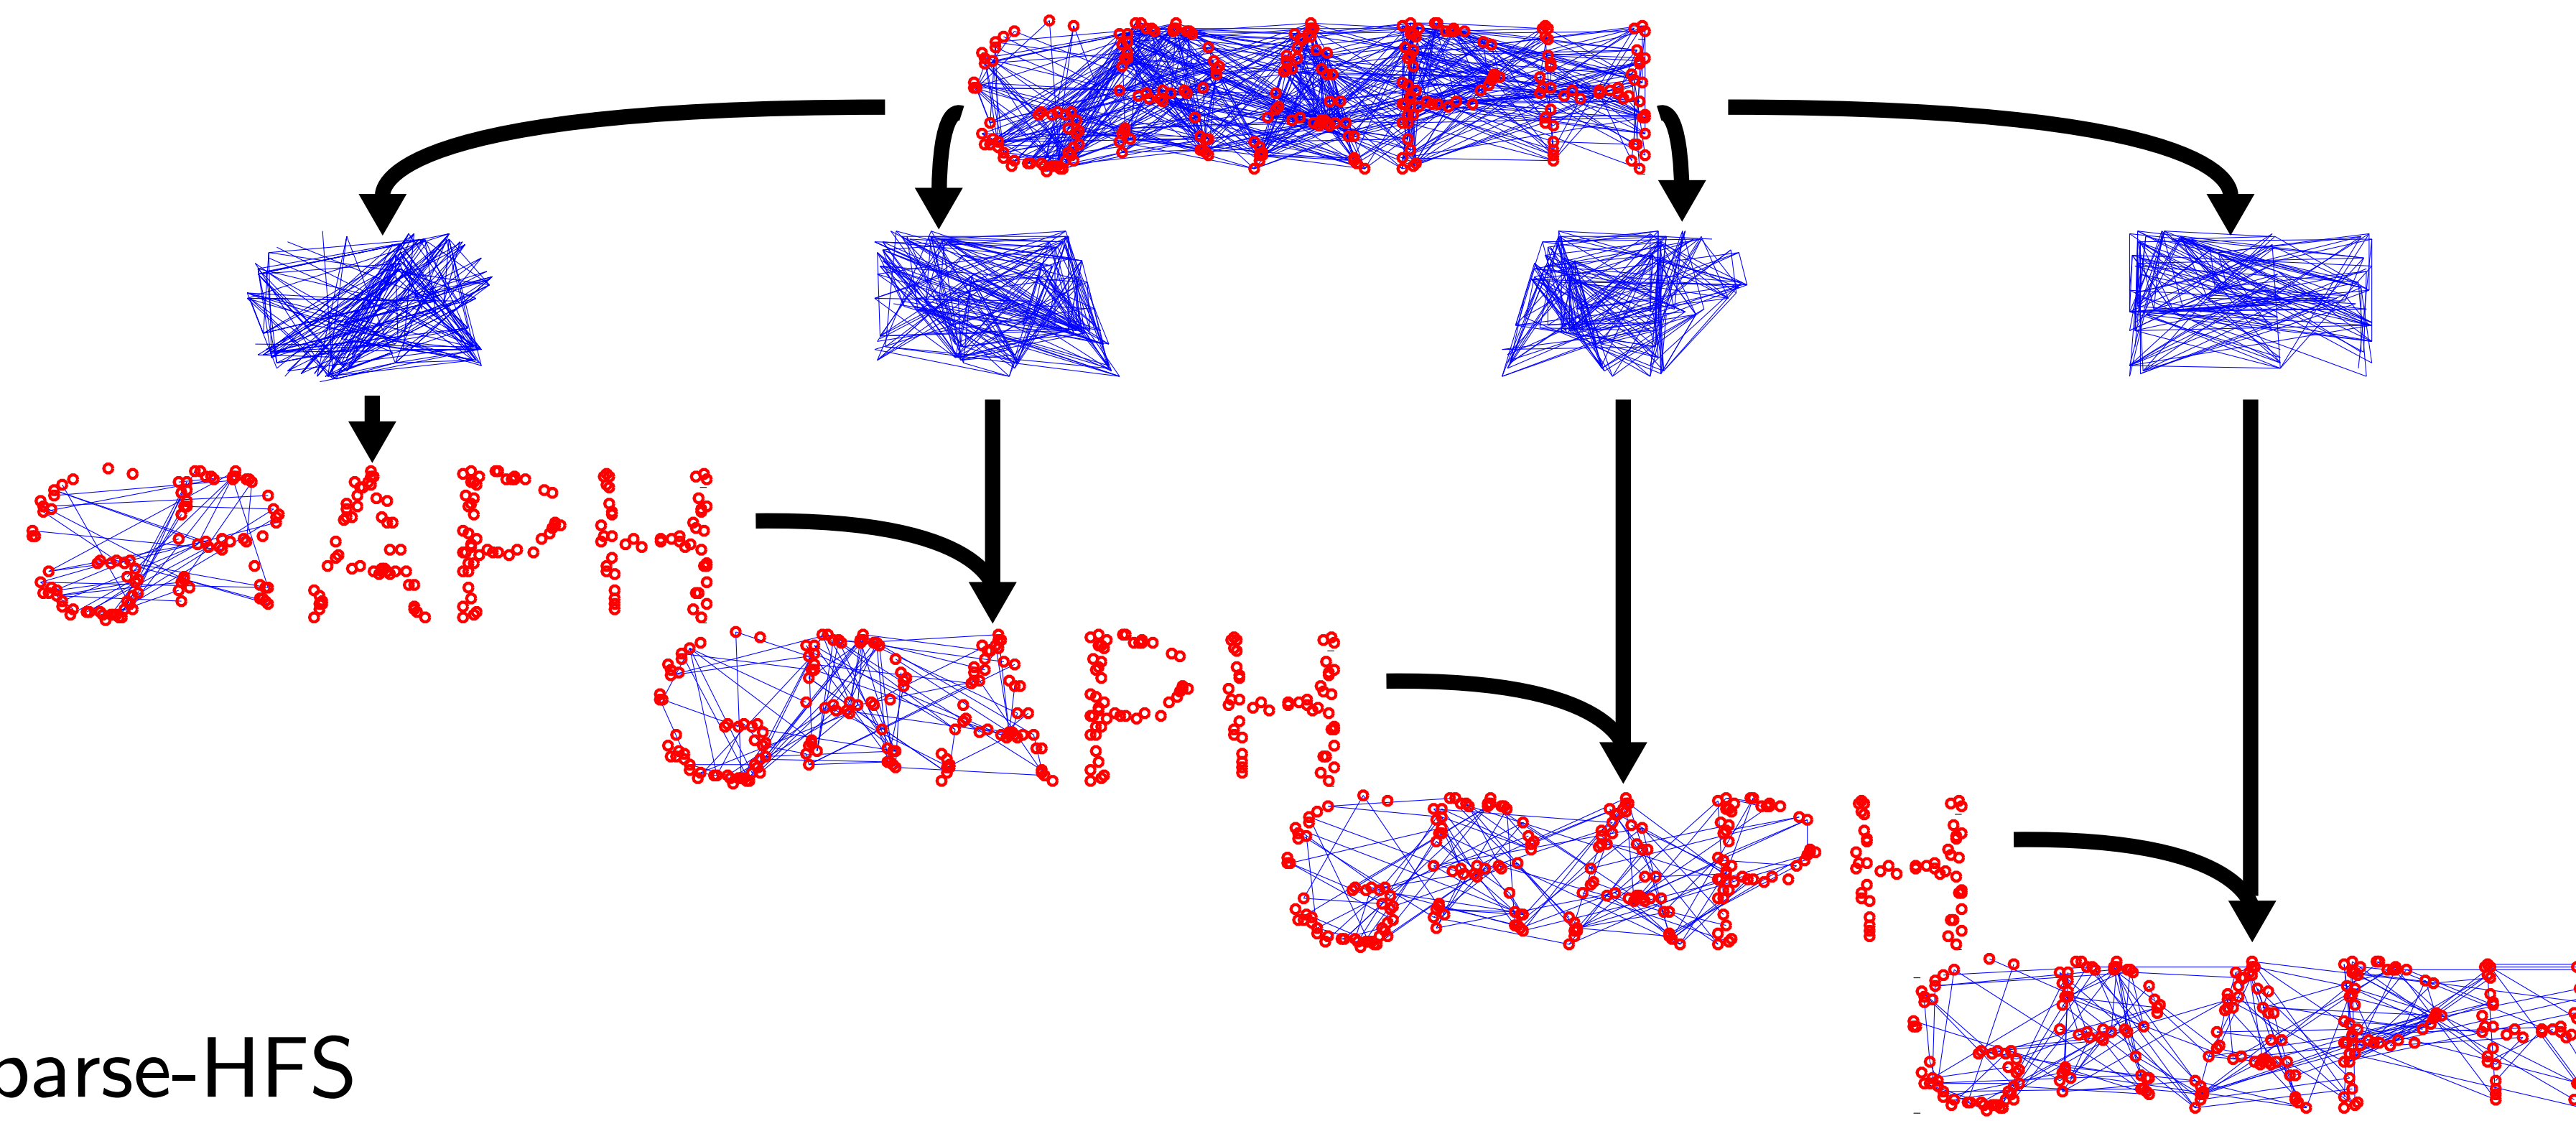
Stable-HFS [2]: $\mu = ((\gamma l L_G + I_S)^+ \mathbf{y}_S)^\top \mathbf{1} / ((\gamma l L_G + I_S)^+ \mathbf{1})^\top \mathbf{1}$, $\hat{\mathbf{f}} = (\gamma l L_G + I_S)^+ (\tilde{\mathbf{y}}_S - \mu \mathbf{1}) = (P_{\mathcal{F}}(\gamma l L_G + I_S))^+ \tilde{\mathbf{y}}_S$



Fast SDD Linear Solver [3] + Stable-HFS



Laplacian spectral ε -sparsifier: $(1 - \varepsilon) \mathbf{x}^\top L_G \mathbf{x} \leq \mathbf{x}^\top L_{\mathcal{H}} \mathbf{x} \leq (1 + \varepsilon) \mathbf{x}^\top L_G \mathbf{x}$.



Sparse-HFS

input Graph $\mathcal{G} = (\mathcal{X}, \mathcal{E})$, labels \mathbf{y}_S , accuracy ε
output Solution $\hat{\mathbf{f}}$, sparsified graph \mathcal{H}
 Let $\alpha = 1/(1 - \varepsilon)$ and $N = \alpha^2 n \log^2(n)/\varepsilon^2$
 Partition \mathcal{E} in $\tau = \lceil m/N \rceil$ blocks $\Delta_1, \dots, \Delta_\tau$
 Initialize $\mathcal{H} = \emptyset$
for $t = 1, \dots, \tau$ **do**
 Load Δ_t in memory
 Compute $\mathcal{H}_t = \text{SPARSIFY}(\mathcal{H}_{t-1}, \Delta_t, N, \alpha)$
end for
 Center the labels $\tilde{\mathbf{y}}_S$
 Compute $\hat{\mathbf{f}}$ with STABLE-HFS with $\tilde{\mathbf{y}}_S$ using a suitable SDD solver

Figure 1: SPARSE-HFS

input A sparsifier \mathcal{H} , block Δ , number of edges N , effective resistance accuracy α
output A sparsifier \mathcal{H}' , probabilities $\{\tilde{p}_e : e \in \mathcal{H}'\}$.
 Compute estimates of \tilde{R}_e for any edge in $\mathcal{H} + \Delta$ such that $1/\alpha \leq \tilde{R}_e/R_e \leq \alpha$ with an SDD solver [13]
 Compute probabilities $\tilde{p}_e = (a_e \tilde{R}_e) / (\alpha(n-1))$ and weights $w_e = a_e / (N \tilde{p}_e)$
 For all edges $e \in \mathcal{H}$ compute $\tilde{p}_e \leftarrow \min\{\tilde{p}_e, \tilde{p}_e\}$ and initialize $\mathcal{H}' = \emptyset$
for all edges $e \in \mathcal{H}$ **do**
 Add edge e to \mathcal{H}' with weight w_e with probability $\tilde{p}_e / \tilde{p}_e$
end for
for all edges $e \in \Delta$ **do**
 for $i = 1$ to N **do**
 Add edge e to \mathcal{H}' with weight w_e with probability \tilde{p}_e
 end for
end for

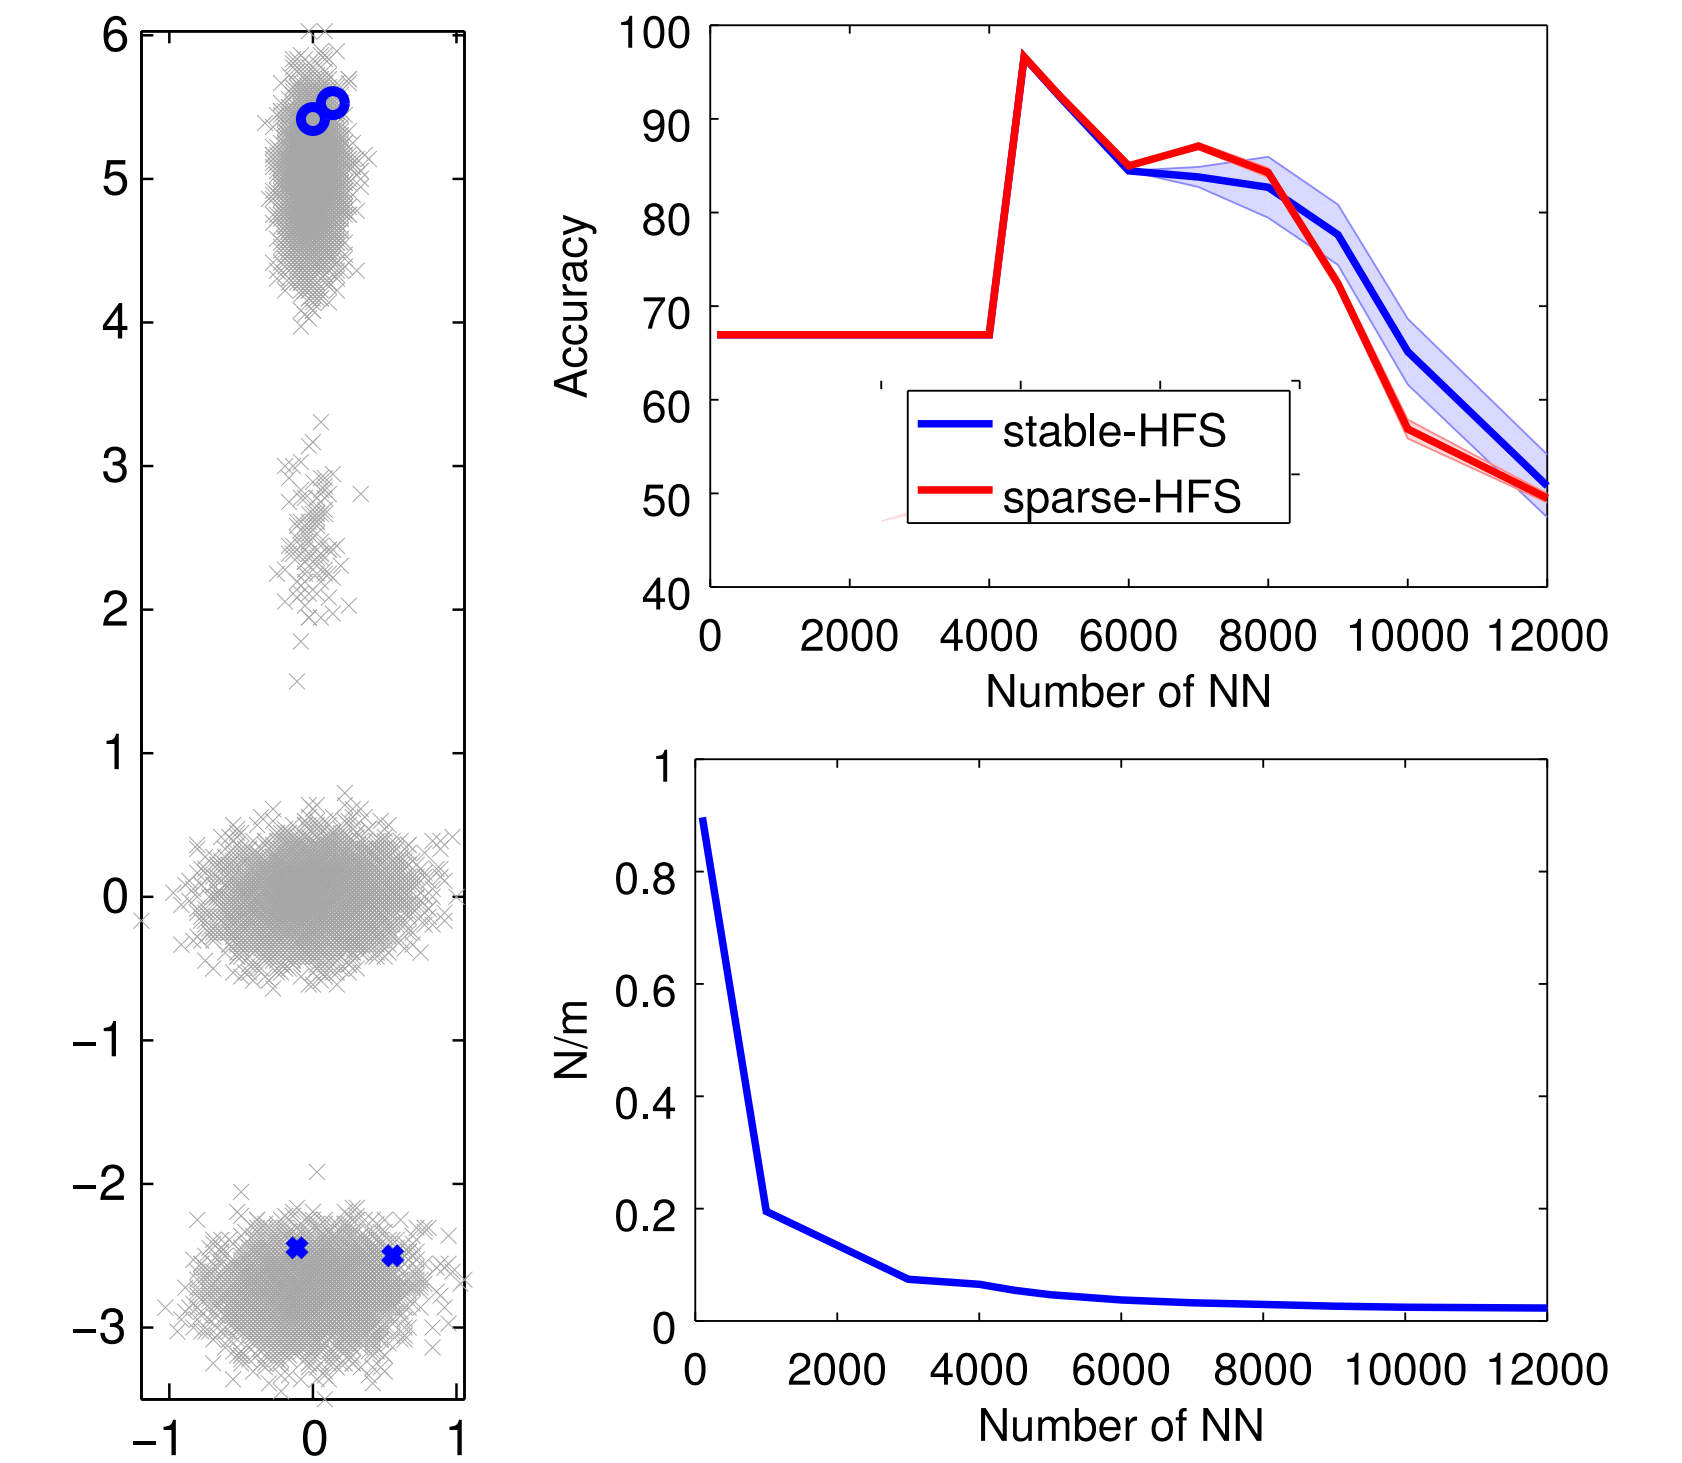
Figure 2: Kelner-Levin Sparsification Algorithm [12]



$$R(\tilde{\mathbf{f}}) \leq \hat{R}(\tilde{\mathbf{f}}) + \frac{l^2 \gamma^2 \lambda_n(\mathcal{G})^2 M^2 \varepsilon^2}{(l\gamma(1 - \varepsilon)\lambda_2(\mathcal{G}) - 1)^4} + \beta + \left(2\beta + \frac{c^2(l+u)}{lu}\right) \sqrt{\frac{\pi(l,u) \ln \frac{1}{\delta}}{2}}$$

$$\beta \leq \frac{1.5M\sqrt{l}}{(l\gamma(1 - \varepsilon)\lambda_2(\mathcal{G}) - 1)^2} + \frac{4M}{l\gamma(1 - \varepsilon)\lambda_2(\mathcal{G}) - 1}$$

Toy example



Spam Classification (TREC07)

	Guarantees	Space	Preprocessing Time	Solving Time
SparseHFS	✓	$\mathcal{O}(n \text{ polylog}(n))$	$\mathcal{O}(m \text{ polylog}(n))$	$\mathcal{O}(n \text{ polylog}(n))$
StableHFS	✓	$\mathcal{O}(n^2)$	$\mathcal{O}(m)$	$\mathcal{O}(m \text{ polylog}(n))$
Fergus	✗	$\mathcal{O}(nd + nk + b^2)$	$\mathcal{O}(kb^3 + db^3)$	$\mathcal{O}(k^2 \text{ polylog}(k) + nk)$
SimpleHFS	✓	$\mathcal{O}(n^2)$	$\mathcal{O}(n^3)$	$\mathcal{O}(k^3)$
SubSample	✗	$\mathcal{O}(k^2)$	$\mathcal{O}(m)$	$\mathcal{O}(k^2 \text{ polylog}(k) + n)$

✓ (under assumptions)


$n = 75419$ raw emails

$y = \text{HAM or SPAM}$

$d = 68697$ features, TF-IDF extracted from the text

$l = \{100, 1000\}$

$\varepsilon = 0.8$



ExtraLearn

We would like to thank Ioannis Koutis for many useful discussions.