

Learning to Identify Sources of Network Diffusion

Gonzalo Mateos

Dept. of ECE and Goergen Institute for Data Science University of Rochester gmateosb@ece.rochester.edu http://www.hajim.rochester.edu/ece/sites/gmateos

Co-author: Chang Ye Acknowledgment: NSF Awards CCF-1750428, CCF-1934962 and ECCS-1809356

> EUSIPCO 2022, Belgrade, Serbia August 30, 2022





- Network as undirected graph $G = (\mathcal{V}, \mathcal{E})$: encode pairwise relationships
- Desiderata: Process, analyze and learn from network data [Kolaczyk'09]
 - \Rightarrow Study graph signals, data associated with N nodes in $\mathcal V$
- ► Ex: Opinion profile, buffer congestion levels, neural activity, epidemic

э.

(日)



• Graph signals mappings $x : \mathcal{V} \to \mathbb{R}$, represented as vectors $\mathbf{x} \in \mathbb{R}^N$

 \Rightarrow As.: Signal properties related to topology of G

- ► To process graph signals \Rightarrow Graph-shift operator $\mathbf{S} \in \mathbb{R}^{N \times N}$
 - \Rightarrow Local $S_{ij} = 0$ for $i \neq j$ and $(i, j) \notin \mathcal{E} \Rightarrow \mathsf{Ex}$: **A** or $\mathbf{L} = \mathbf{D} \mathbf{A}$
 - \Rightarrow Spectrum of symmetric $\mathbf{S} = \mathbf{V} \wedge \mathbf{V}^{\top}$

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● のへで



• Graph signals mappings $x : \mathcal{V} \to \mathbb{R}$, represented as vectors $\mathbf{x} \in \mathbb{R}^N$

 \Rightarrow As.: Signal properties related to topology of G

- ► To process graph signals \Rightarrow Graph-shift operator $\mathbf{S} \in \mathbb{R}^{N \times N}$
 - \Rightarrow Local $S_{ij} = 0$ for $i \neq j$ and $(i, j) \notin \mathcal{E} \Rightarrow \mathsf{Ex}$: **A** or $\mathbf{L} = \mathbf{D} \mathbf{A}$
 - \Rightarrow Spectrum of symmetric $\mathbf{S} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{\top}$
- Graph Fourier Transform (GFT) for signals: $\tilde{\mathbf{x}} = \mathbf{V}^{\top} \mathbf{x}$
- Graph filters $\mathbf{H} : \mathbb{R}^N \to \mathbb{R}^N$ are maps between graph signals
 - \Rightarrow Polynomial in **S** with coefficients $h \in \mathbb{R}^{L} \Rightarrow H := \sum_{l=0}^{L-1} h_{l} S^{l}$
 - \Rightarrow Orthogonal frequency operator: $\textbf{H} = \textbf{V} \text{diag}(\tilde{\textbf{h}}) \textbf{V}^{\top}$
 - \Rightarrow Freq. response (GFT for filters): $\tilde{\mathbf{h}} = \Psi \mathbf{h}$ and $[\Psi]_{k,l} = \lambda_k^{l-1}$

・ロット (日)・ (日)・ (日)・ (日)・

Diffusion processes as graph filter outputs

- Q: Upon observing a graph signal y, how was this signal generated?
- Postulate y is the response of linear diffusion to a sparse input x

$$\mathbf{y} = \alpha_0 \prod_{l=1}^{\infty} (\mathbf{I} - \alpha_l \mathbf{S}) \mathbf{x} = \sum_{l=0}^{\infty} \beta_l \mathbf{S}' \mathbf{x}$$

 \Rightarrow Common generative model, e.g., heat diffusion, consensus

▲□▶▲□▶▲≡▶▲≡▶ ≡ のへで



- Q: Upon observing a graph signal y, how was this signal generated?
- Postulate y is the response of linear diffusion to a sparse input x

$$\mathbf{y} = \alpha_0 \prod_{l=1}^{\infty} (\mathbf{I} - \alpha_l \mathbf{S}) \mathbf{x} = \sum_{l=0}^{\infty} \beta_l \mathbf{S}' \mathbf{x}$$

 \Rightarrow Common generative model, e.g., heat diffusion, consensus

• Cayley-Hamilton asserts we can write diffusion as $(L \le N)$

$$\mathbf{y} = \left(\sum_{l=0}^{L-1} h_l \mathbf{S}^l\right) \mathbf{x} := \mathbf{H} \mathbf{x}$$

- Model: Observed network process as output of a graph filter
 - \Rightarrow View few elements in supp(**x**) =: { $i : x_i \neq 0$ } as sources

・ロット (日)・ (日)・ (日)・ (日)・



Motivation and source localization problem



- Ex: Global opinion/belief profile formed by spreading a rumor
 - \Rightarrow What was the rumor? Who started it?



- **Problem:** Blind identification of graph filters with multiple sparse inputs
 - \Rightarrow Suppose we observe *P* output signals $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_P] \in \mathbb{R}^{N \times P}$

3

イロト イヨト イヨト

Motivation and source localization problem



- Ex: Global opinion/belief profile formed by spreading a rumor
 - \Rightarrow What was the rumor? Who started it?



- **Problem:** Blind identification of graph filters with multiple sparse inputs
 - \Rightarrow Suppose we observe *P* output signals $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_P] \in \mathbb{R}^{N \times P}$
- **Q**: Given **S**, can we find sparse **X** and the filter coeffs. h from $\mathbf{Y} = \mathbf{HX}$?
 - \Rightarrow Extends classical blind deconvolution to graphs
 - \Rightarrow Localization of sources that diffuse on the network

= nac

イロト 人間 ト イヨト イヨト



Leverage frequency response of graph filters

 $\mathbf{Y} = \mathbf{H}\mathbf{X} \Rightarrow \mathbf{Y} = \mathbf{V} \operatorname{diag}(\mathbf{\Psi}\mathbf{h}) \mathbf{V}^{\top} \mathbf{X}$

 \Rightarrow **Y** is a bilinear function of the unknowns **h** and **X**

▶ III-posed problem \Rightarrow *L* + *NP* unknowns and *NP* observations

 \Rightarrow As.: X has S-sparse columns i.e., $\|X\|_0 := |supp(X)| \le PS$

▲□▶▲□▶▲≡▶▲≡▶ ≡ のへで



Leverage frequency response of graph filters

 $\mathbf{Y} = \mathbf{H}\mathbf{X} \Rightarrow \mathbf{Y} = \mathbf{V} \operatorname{diag}(\mathbf{\Psi}\mathbf{h}) \mathbf{V}^{\top} \mathbf{X}$

 \Rightarrow **Y** is a bilinear function of the unknowns **h** and **X**

▶ III-posed problem \Rightarrow *L* + *NP* unknowns and *NP* observations

 \Rightarrow As.: X has S-sparse columns i.e., $\|X\|_0 := |supp(X)| \le PS$

▶ Blind graph filter identification ⇒ Non-convex feasibility problem

find $\{\mathbf{h}, \mathbf{X}\}$, s. to $\mathbf{Y} = \mathbf{V} \text{diag}(\mathbf{\Psi}\mathbf{h}) \mathbf{V}^{\top} \mathbf{X}$, $\|\mathbf{X}\|_{0} \leq PS$

 \Rightarrow Identifiability for Bernoulli-Gaussian model on X [Li et al'17]

S. Segarra, G. Mateos, A. G. Marques and A. Ribeiro, "Blind identification of graph filters," *IEEE TSP*, 2017

▲□▶▲□▶▲≡▶▲≡▶ ≡ のへで

Supervised learning setting



• Given independent training samples $\mathcal{T} := \{\mathbf{X}_i, \mathbf{Y}_i\}_{i=1}^{|\mathcal{T}|}$

 \Rightarrow Drawn from joint distribution of filters and sparse sources



Goal: learn the parametric mapping $\hat{\mathbf{X}} = \Phi(\mathbf{Y}; \mathbf{\Theta})$ by minimizing a loss

$$L(\mathbf{\Theta}) := rac{1}{|\mathcal{T}|} \sum_{i \in \mathcal{T}} \ell(\mathbf{X}_i, \Phi(\mathbf{Y}_i; \mathbf{\Theta}))$$

э

(日)



Source localization on graphs

- Maximum-likelihood estimator optimal for trees [Pinto et al'12]
- Scalable under restrictive dependency assumptions [Feizi el al'16]
- Non-convex estimators of sparse sources [Pena et al'16], [Hu et al'16]
- Blind identification of graph filters [Segarra et al'17]
 - Matrix lifting can hinder applicability to large graphs
- Blind identification of invertible graph filters [Ye et al'18]
 - Convex formulation amenable to e.g., ADMM solvers
 - Multi-signal case with arbitrary supports

イロト 不得 トイヨト イヨト ニヨー



Source localization on graphs

- Maximum-likelihood estimator optimal for trees [Pinto et al'12]
- Scalable under restrictive dependency assumptions [Feizi el al'16]
- Non-convex estimators of sparse sources [Pena et al'16], [Hu et al'16]
- Blind identification of graph filters [Segarra et al'17]
 - Matrix lifting can hinder applicability to large graphs
- Blind identification of invertible graph filters [Ye et al'18]
 - Convex formulation amenable to e.g., ADMM solvers
 - Multi-signal case with arbitrary supports
- Our contribution: data-driven deep learning solution rooted in GSP
 - \Rightarrow Unroll and truncate the model-based ADMM iterations
 - \Rightarrow Trainable parametric architecture is interpretable
 - \Rightarrow Parameter efficient and offers controllable complexity

3

イロト イヨト イヨト イヨト



• Inverse filter $\mathbf{G} = \mathbf{H}^{-1}$ is also a graph filter on *G* [Sandryhaila-Moura'13]

$$\Rightarrow$$
 Requires $\tilde{h}_i = \sum_{l=0}^{L-1} h_l \lambda_i^l \neq 0$, for all $i = 1, ..., N$

 \Rightarrow Inverse-filter coefficients $\mathbf{g} \in \mathbb{R}^{\textit{N}}$, frequency response $\tilde{\mathbf{g}} = \boldsymbol{\Psi} \mathbf{g}$

Recast as linear inverse problem [Wang-Chi'16], [Ye et al'18]

$$\min_{\{\tilde{\mathbf{g}},\mathsf{X}\}} \|\mathsf{X}\|_0, \quad \text{s. to} \quad \mathsf{X} = \mathsf{V}\mathsf{diag}(\tilde{\mathbf{g}})\mathsf{V}^{\top}\mathsf{Y}, \ \mathsf{X} \neq \mathbf{0}$$

► Still NP hard.



• Inverse filter $\mathbf{G} = \mathbf{H}^{-1}$ is also a graph filter on *G* [Sandryhaila-Moura'13]

$$\Rightarrow$$
 Requires $\tilde{h}_i = \sum_{l=0}^{L-1} h_l \lambda_i^l \neq 0$, for all $i = 1, ..., N$

 \Rightarrow Inverse-filter coefficients $\mathbf{g} \in \mathbb{R}^{\textit{N}}$, frequency response $\tilde{\mathbf{g}} = \boldsymbol{\Psi} \mathbf{g}$

Recast as linear inverse problem [Wang-Chi'16], [Ye et al'18]

$$\min_{[\frac{g}{2},\mathsf{X}\}} \|\mathsf{X}\|_0, \quad \text{s. to} \quad \mathsf{X} = \mathsf{V}\mathsf{diag}(\tilde{\mathsf{g}})\mathsf{V}^{\top}\mathsf{Y}, \ \mathsf{X} \neq \mathbf{0}$$

▶ Still NP hard. Relax! and minimize convex **||X**||₁

$$\hat{\tilde{\mathbf{g}}} = \underset{\tilde{\mathbf{g}}}{\operatorname{argmin}} \| (\mathbf{Y}^{\top} \mathbf{V} \odot \mathbf{V}) \tilde{\mathbf{g}} \|_{1}, \quad \text{s. to} \quad \mathbf{1}^{\top} \tilde{\mathbf{g}} = c$$

 \Rightarrow Constraint fixes the scale and avoids all-zero solution

 \Rightarrow $\ell_1\text{-synthesis}$ problem, can be solved via ADMM

イロト イヨト イヨト



▶ Let $\mathbf{Z} = \mathbf{Y}^{\top} \mathbf{V} \odot \mathbf{V}$ and $\mathbf{x} = \operatorname{vec}(\mathbf{X})$, rewrite problem as

 $\min_{\{\tilde{\mathbf{g}},\mathbf{x}\}} \|\mathbf{x}\|_1, \quad \text{s. to} \quad \mathbf{1}^\top \tilde{\mathbf{g}} = c, \ \mathbf{Z}\tilde{\mathbf{g}} - \mathbf{x} = \mathbf{0}$

• ADMM solver ($\Gamma := \rho_{\lambda} \mathbf{Z}^{\top} \mathbf{Z} + \rho_{\mu} \mathbf{1}_{N} \mathbf{1}_{N}^{\top}$, dual variables $\{\lambda, \mu\}$)

$$\begin{split} \tilde{\mathbf{g}}[k+1] &= \mathbf{\Gamma}^{-1} \left[\mathbf{Z}^{\top} (\rho_{\lambda} \mathbf{x}[k] - \lambda[k]) + (\rho_{\mu} c - \mu[k]) \mathbf{1}_{N} \right], \\ \mathbf{x}[k+1] &= \mathcal{S}_{\rho_{\lambda}^{-1}} (\mathbf{Z} \tilde{\mathbf{g}}[k+1] + \lambda[k] / \rho_{\lambda}), \\ \lambda[k+1] &= \lambda[k] + \rho_{\lambda} (\mathbf{Z} \tilde{\mathbf{g}}[k+1] - \mathbf{x}[k+1]), \\ \mu[k+1] &= \mu[k] + \rho_{\mu} (\mathbf{1}_{N}^{\top} \tilde{\mathbf{g}}[k+1] - c) \end{split}$$

Limitations

- Step-sizes $\rho_{\lambda}, \rho_{\mu}$ need to be tuned
- Hundreds or thousands of iterations until convergence
- Matrix inversion may hinder scalability to large graphs
- **Idea:** unroll the iterations and learn the parameters from dataset \mathcal{T}



- Map ADMM updates as sub-layers within a layer of the deep net
 - \Rightarrow Stack K of the resulting layers to form $\Phi(\mathbf{Y}; \mathbf{\Theta})$
 - \Rightarrow End-to-end learning of Θ using mini-batch SGD

Filter sub-layer \mathcal{G}_k

$$\begin{split} \tilde{\mathbf{g}}[k+1] &= (\mathbf{\Gamma}^{(k)})^{-1} [\mathbf{Z}^{\top} (\mathbf{x}[k] - \rho_1^{(k)} \lambda[k]) + (\rho_2^{(k)} c - \rho_1^{(k)} \mu[k]) \mathbf{1}_N] \\ \mathbf{\Gamma}^{(k)} &= \mathbf{Z}^{\top} \mathbf{Z} + \rho_2^{(k)} \mathbf{1}_N \mathbf{1}_N^{\top} \\ \Rightarrow \text{Learn } \{\rho_1^{(k)}, \rho_2^{(k)}\}_{k=1}^K, \text{ where } \rho_1^{(k)}, \rho_2^{(k)} \ge 0, \text{ for } k = 1, \dots, K \end{split}$$

э.

・ロト ・日ト ・ヨト ・ヨト



- Map ADMM updates as sub-layers within a layer of the deep net
 - \Rightarrow Stack K of the resulting layers to form $\Phi(\mathbf{Y}; \mathbf{\Theta})$
 - \Rightarrow End-to-end learning of Θ using mini-batch SGD

Sources sub-layer \mathcal{X}_k

$$\begin{split} \mathbf{x}[k+1] &= \mathcal{S}_{\tau^{(k)}}(\alpha_1^{(k)} \mathbf{Z} \mathbf{\tilde{g}}[k+1] + \alpha_2^{(k)} \lambda[k]) \\ &\Rightarrow \text{Learn } \{\alpha_1^{(k)}, \alpha_2^{(k)}, \tau^{(k)}\}, \text{ where } \tau^{(k)} \geq 0, \text{ for } k = 1, \dots, K \end{split}$$

2

・ロト ・日 ・ ・ ヨ ・ ・ ヨ ・



- Map ADMM updates as sub-layers within a layer of the deep net
 - \Rightarrow Stack K of the resulting layers to form $\Phi(\mathbf{Y}; \mathbf{\Theta})$
 - \Rightarrow End-to-end learning of Θ using mini-batch SGD

Multiplier sub-layer \mathcal{M}_k

$$\begin{split} \lambda[k+1] &= \beta_1^{(k)} \lambda[k] + \beta_2^{(k)} \mathbf{Z} \tilde{\mathbf{g}}[k+1] + \beta_3^{(k)} \mathbf{x}[k+1] \\ \mu[k+1] &= \gamma_1^{(k)} \mu[k] + \gamma_2^{(k)} \mathbf{1}_N^\top \tilde{\mathbf{g}}[k+1] + \gamma_3^{(k)} c \\ &\Rightarrow \text{Learn } \{\beta_i^{(k)}\}_{i=1}^3, \{\gamma_i^{(k)}\}_{i=1}^3, \text{ for } k = 1, \dots, K \end{split}$$

Design considerations

- Additional parameters to broaden the model's expressive power
- Forgo the parameter sharing constraint imposed by the unrolling

Source location predictions:
$$\Phi(\mathbf{Y}_{test}, \hat{\boldsymbol{\Theta}}) = \operatorname{unvec}[(\mathbf{Y}_{test}^{\top} \mathbf{V} \odot \mathbf{V}) \mathbf{\tilde{g}}[\mathbf{k}]]$$

(日)







- ► Consider undirected random graphs with $\mathbf{S} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$ ⇒ Erdős-Rénvi with N = 20 and edge prob. p = 0.3
- Sources $X_{train} \in \mathbb{R}^{N \times |\mathcal{T}|}$ from a Bernoulli-Gaussian model

 \Rightarrow $|\mathcal{T}|$ = 64000 with Bernoulli parameter heta = 0.2

- Filter $\mathbf{h}_{\text{train}} = (\mathbf{e}_1 + \alpha \mathbf{b})/\|\mathbf{e}_1 + \alpha \mathbf{b}\|_1$ as in [Wang-Chi'16] $\Rightarrow \mathbf{e}_1 = [1, 0, \dots, 0]^\top \in \mathbb{R}^L$ and $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - \Rightarrow Recovery performance increases while $\alpha \geq$ 0 decreases

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●



- ► Consider undirected random graphs with $\mathbf{S} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$ ⇒ Erdős-Rényi with N = 20 and edge prob. p = 0.3
- Sources $\mathbf{X}_{\text{train}} \in \mathbb{R}^{N \times |\mathcal{T}|}$ from a Bernoulli-Gaussian model

 \Rightarrow $|\mathcal{T}| =$ 64000 with Bernoulli parameter heta = 0.2

- Filter $\mathbf{h}_{\text{train}} = (\mathbf{e}_1 + \alpha \mathbf{b}) / \|\mathbf{e}_1 + \alpha \mathbf{b}\|_1$ as in [Wang-Chi'16] $\Rightarrow \mathbf{e}_1 = [1, 0, \dots, 0]^\top \in \mathbb{R}^L$ and $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - \Rightarrow Recovery performance increases while $\alpha \ge 0$ decreases
- Figure of merit: Relative recovery error of X and ğ ⇒ RMSE_X = ||X̂ - X_{test}||_F/||X_{test}||_F, RMSE_ğ = ||ĝ - g̃_{test}||₂/||g̃_{test}||₂
- **Figure of merit:** Source localization accuracy

$$\Rightarrow ACC = |\hat{\mathcal{I}} \cap \mathcal{I}_{\mathsf{test}}| / |\mathcal{I}_{\mathsf{test}}|$$

 $\Rightarrow \text{Support estimate } \hat{\mathcal{I}} = \text{supp}_{\kappa}(\hat{\mathbf{X}}): \text{ if } |\hat{\mathbf{X}}_{ip}| \geq \kappa, \text{ then } (i, p) \in \hat{\mathcal{I}}$

・ロット (日)・ (日)・ (日)・ (日)・



- Each epoch: randomly split X_{train} into Q = 1600 mini-batches $\{X_q\}_{q=1}^Q$
- ▶ Then sample Q graph filter coefficients $\{h_q\}_{q=1}^Q$, $L = 3, \alpha = 1$
 - \Rightarrow Generate the observations $\mathbf{Y}_q = \mathbf{V} \text{diag}(\mathbf{\Psi}_L \mathbf{h}_q) \mathbf{V}^\top \mathbf{X}_q$
- Loss function accounting for sign ambiguity

$$L(\boldsymbol{\Theta}) = \sum_{q=1}^{Q} \min\left(\frac{\|\boldsymbol{\Phi}(\mathbf{Y}_{q};\boldsymbol{\Theta}) - \mathbf{X}_{q}\|_{F}}{\|\mathbf{X}_{q}\|_{F}}, \frac{\|\boldsymbol{\Phi}(\mathbf{Y}_{q};\boldsymbol{\Theta}) + \mathbf{X}_{q}\|_{F}}{\|\mathbf{X}_{q}\|_{F}}\right)$$

- Parameter initialization
 - $\{\rho_1^{(k)}, \rho_2^{(k)}, \tau^{(k)}\}_{k=1}^K$ are i.i.d. uniformly distributed in [0, 1]
 - All other parameters are drawn from a standard Gaussian

▲□▶▲□▶▲≡▶▲≡▶ ≡ のへで



▶ Visualization of recovery performance for SLoG-Net with K = 5 layers



 \blacktriangleright Recovery errors are $\mathrm{RMSE}_{\bm{X}}=0.090$ and $\mathrm{RMSE}_{\bm{\tilde{g}}}=0.086$

æ

★ B + ★ B +



- Comparisons with ADMM ($N = 20, P_{\text{test}} = 100$, over 50 trials)
 - ► Top: RMSE_X for ADMM and SLoG-Net ⇒ The shaded region indicates the standard deviation
 - ► Bottom: ACC of support recovery, $\kappa = 0.1$ \Rightarrow Support estimate $\hat{\mathcal{I}} := \operatorname{supp}_{\kappa}(\Phi(\mathbf{Y}_{\text{test}}; \hat{\mathbf{\Theta}})),$
 - \Rightarrow Ground truth, $\mathcal{I}_{test} := supp_{\kappa}(\mathbf{X}_{test}).$



Mean elapsed time: 0.005s for SLoG-Net, 0.067s for ADMM



Learning to identify sources of network diffusion

- \Rightarrow Data-driven trainable neural network architecture
- \Rightarrow Key: leverage inductive biases of the GSP model-based solution

Unrolling the ADMM iterations

- \Rightarrow The SLoG-Net architecture is interpretable
- \Rightarrow Parameter and time efficient compared with iterative ADMM

Ongoing work

- \Rightarrow Recovery of noise-corrupted signals
- \Rightarrow Effective recovery of more general graph filters
- \Rightarrow Evaluate performance on real data sets

3

イロン 人間 とくほ とくほど