

Graph-Aware Diffusion for Signal Generation

Vimal Kumarasamy Balasubramanian[†]

[†] Delft University of Technology, The Netherlands

IEEE ICASSP 2026

- ▶ **Generative modeling:** learn to draw new samples from an unknown distribution p_0 , given only samples $\{\mathbf{x}_0^{(i)}\}_{i=1}^M \sim p_0$
 - ⇒ huge success in images, audio, and *graph* generation

Generating signals on a graph

- ▶ **Generative modeling:** learn to draw new samples from an unknown distribution p_0 , given only samples $\{\mathbf{x}_0^{(i)}\}_{i=1}^M \sim p_0$
 - ⇒ huge success in images, audio, and *graph* generation

- ▶ Far less studied: generating **graph signals** on a **known, fixed graph** \mathcal{G}
 - ⇒ a signal $\mathbf{x} \in \mathbb{R}^N$ is a value per node (sensor reading, rating, speed, ...)
 - ⇒ most graph-diffusion work generates *the graph*; here the **graph is given**

Generating signals on a graph

- ▶ **Generative modeling:** learn to draw new samples from an unknown distribution p_0 , given only samples $\{\mathbf{x}_0^{(i)}\}_{i=1}^M \sim p_0$
 - ⇒ huge success in images, audio, and *graph* generation

- ▶ Far less studied: generating **graph signals** on a **known, fixed graph** \mathcal{G}
 - ⇒ a signal $\mathbf{x} \in \mathbb{R}^N$ is a value per node (sensor reading, rating, speed, ...)
 - ⇒ most graph-diffusion work generates *the graph*; here the **graph is given**

- ▶ **Why exploit a known graph?** It is a strong, free prior:
 - ⇒ bake the structure into **both** the forward *and* the backward process
 - ⇒ reuse the GSP toolbox – graph filters, smoothness priors, the spectrum
 - ⇒ generated signals respect the topology by construction

Generating signals on a graph

- ▶ **Generative modeling:** learn to draw new samples from an unknown distribution p_0 , given only samples $\{\mathbf{x}_0^{(i)}\}_{i=1}^M \sim p_0$
 - ⇒ huge success in images, audio, and *graph* generation
- ▶ Far less studied: generating **graph signals** on a **known, fixed graph** \mathcal{G}
 - ⇒ a signal $\mathbf{x} \in \mathbb{R}^N$ is a value per node (sensor reading, rating, speed, ...)
 - ⇒ most graph-diffusion work generates *the graph*; here the **graph is given**
- ▶ **Why exploit a known graph?** It is a strong, free prior:
 - ⇒ bake the structure into **both** the forward *and* the backward process
 - ⇒ reuse the GSP toolbox – graph filters, smoothness priors, the spectrum
 - ⇒ generated signals respect the topology by construction
- ▶ **Applications:** probabilistic forecasting, recommender systems, data imputation in sensor / traffic networks, wireless, 3D meshes ...

- ▶ A **forward** SDE gradually turns data into noise; a **backward** SDE turns noise back into data

Forward (data \rightarrow noise)

$$d\mathbf{x}_t = f(\mathbf{x}_t, t) dt + g(\mathbf{x}_t, t) d\mathbf{w}_t$$

- ▶ drift f , diffusion g , Brownian motion \mathbf{w}_t
 \Rightarrow converges to a simple distribution p_T

Backward (noise \rightarrow data)

$$d\mathbf{x}_t = [f(\mathbf{x}_t, t) - g(\mathbf{x}_t, t)^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)] dt + g d\mathbf{w}_t$$

- ▶ runs in reverse time from a sample of p_T
 \Rightarrow needs only the **score** $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$

- ▶ A **forward** SDE gradually turns data into noise; a **backward** SDE turns noise back into data

Forward (data \rightarrow noise)


$$d\mathbf{x}_t = f(\mathbf{x}_t, t) dt + g(\mathbf{x}_t, t) d\mathbf{w}_t$$

Backward (noise \rightarrow data)

$$d\mathbf{x}_t = [f(\mathbf{x}_t, t) - g(\mathbf{x}_t, t)^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)] dt + g d\mathbf{w}_t$$

- ▶ drift f , diffusion g , Brownian motion \mathbf{w}_t
 \Rightarrow converges to a simple distribution p_T

- ▶ runs in reverse time from a sample of p_T
 \Rightarrow needs only the **score** $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$



- ▶ **Our novelty:** particularize this framework to a known graph \Rightarrow graph-aware forward *and* backward

GAD: put the graph in the drift

- ▶ We use the [heat equation](#) on the Laplacian – a (time-inhomogeneous) Ornstein–Uhlenbeck SDE

$$d\mathbf{x}_t = -c_t \mathbf{L}_\gamma \mathbf{x}_t dt + \sqrt{2c_t} \sigma d\mathbf{w}_t, \quad \mathbf{L}_\gamma = \mathbf{L} + \gamma \mathbf{I}, \quad \gamma > 0.$$

Unlike most works, the **graph enters the drift** – not only the denoiser.

GAD: put the graph in the drift

- ▶ We use the [heat equation](#) on the Laplacian – a (time-inhomogeneous) Ornstein–Uhlenbeck SDE

$$d\mathbf{x}_t = -c_t \mathbf{L}_\gamma \mathbf{x}_t dt + \sqrt{2c_t} \sigma d\mathbf{w}_t, \quad \mathbf{L}_\gamma = \mathbf{L} + \gamma \mathbf{I}, \quad \gamma > 0.$$

Unlike most works, the [graph enters the drift](#) – not only the denoiser.

- ▶ The marginals are available in [closed form](#) (a GMRF):

$$\boldsymbol{\mu}_t = \mathbf{H}_t \mathbf{x}_0, \quad \boldsymbol{\Sigma}_t = \sigma^2 (\mathbf{I} - \mathbf{H}_t^2) \mathbf{L}_\gamma^{-1}, \quad \mathbf{H}_t = e^{-\bar{c}_t \mathbf{L}_\gamma}, \quad \bar{c}_t = \int_0^t c_s ds.$$

\mathbf{H}_t is a [low-pass graph filter](#) that attenuates the signal and injects graph-structured noise.

GAD: put the graph in the drift

- ▶ We use the **heat equation** on the Laplacian – a (time-inhomogeneous) Ornstein–Uhlenbeck SDE

$$d\mathbf{x}_t = -c_t \mathbf{L}_\gamma \mathbf{x}_t dt + \sqrt{2c_t} \sigma d\mathbf{w}_t, \quad \mathbf{L}_\gamma = \mathbf{L} + \gamma \mathbf{I}, \quad \gamma > 0.$$

Unlike most works, the **graph enters the drift** – not only the denoiser.

- ▶ The marginals are available in **closed form** (a GMRF):

$$\boldsymbol{\mu}_t = \mathbf{H}_t \mathbf{x}_0, \quad \boldsymbol{\Sigma}_t = \sigma^2 (\mathbf{I} - \mathbf{H}_t^2) \mathbf{L}_\gamma^{-1}, \quad \mathbf{H}_t = e^{-\bar{c}_t \mathbf{L}_\gamma}, \quad \bar{c}_t = \int_0^t c_s ds.$$

\mathbf{H}_t is a **low-pass graph filter** that attenuates the signal and injects graph-structured noise.

- ▶ **Convergence:** as $\bar{c}_t \rightarrow \infty$ the process reaches a graph-dependent stationary distribution

$$p_\infty = \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{L}_\gamma^{-1}) \quad \implies \quad \text{the simple end distribution is shaped by the graph.}$$

γ makes $-c_t \mathbf{L}_\gamma$ Hurwitz (guarantees convergence); too small \rightarrow slow modes, too large \rightarrow geometry lost.

- ▶ $\mathbf{H}_t = e^{-\bar{c}_t \mathbf{L} \gamma}$ decays **exponentially**. With a uniform linear scheduler, dominant graph modes are noised too early \Rightarrow poor reverse generation.

- ▶ $\mathbf{H}_t = e^{-\bar{c}_t \mathbf{L} \gamma}$ decays **exponentially**. With a uniform linear scheduler, dominant graph modes are noised too early \Rightarrow poor reverse generation.
- ▶ We use a **floor-constrained polynomial scheduler** (FCPS) to slow down the early decay:

$$c_t = c_{\min} + k u^\alpha, \quad \bar{c}_t = c_{\min} t + (c_0 - c_{\min} T) u^{\alpha+1}, \quad u = t/T.$$

k is chosen such that $\bar{c}_T = c_0$. For $\alpha \gg 1$, the term $u^{\alpha+1}$ grows slowly near $t = 0$, so graph modes are noised more gradually.

- ▶ $\mathbf{H}_t = e^{-\bar{c}_t \mathbf{L} \gamma}$ decays **exponentially**. With a uniform linear scheduler, dominant graph modes are noised too early \Rightarrow poor reverse generation.
- ▶ We use a **floor-constrained polynomial scheduler** (FCPS) to slow down the early decay:

$$c_t = c_{\min} + ku^\alpha, \quad \bar{c}_t = c_{\min}t + (c_0 - c_{\min}T)u^{\alpha+1}, \quad u = t/T.$$

k is chosen such that $\bar{c}_T = c_0$. For $\alpha \gg 1$, the term $u^{\alpha+1}$ grows slowly near $t = 0$, so graph modes are noised more gradually.




Figure. FCPS delays mean decay and covariance growth compared with ULS.

Backward process: the score is a denoiser

- ▶ The corresponding **backward** SDE keeps the graph in the drift:

$$d\mathbf{x}_t = -c_t \mathbf{L}_\gamma \mathbf{x}_t dt - 2c_t \sigma^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) dt + \sqrt{2c_t} \sigma d\mathbf{w}_t.$$

Backward process: the score is a denoiser

- ▶ The corresponding **backward** SDE keeps the graph in the drift:

$$d\mathbf{x}_t = -c_t \mathbf{L}_\gamma \mathbf{x}_t dt - 2c_t \sigma^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) dt + \sqrt{2c_t} \sigma d\mathbf{w}_t.$$

- ▶ **Tweedie's formula** turns score estimation into denoising:

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) = \Sigma_t^{-1} (\mathbf{H}_t \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] - \mathbf{x}_t),$$

$$\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = g^*(\mathbf{x}_t), \quad g^* \in \underset{g}{\operatorname{argmin}} \mathbb{E}[\|\mathbf{x}_0 - g(\mathbf{x}_t)\|^2].$$

Since \mathbf{H}_t and Σ_t are known, the hard part is learning the clean-signal predictor g^* .

Backward process: the score is a denoiser

- ▶ The corresponding **backward** SDE keeps the graph in the drift:

$$d\mathbf{x}_t = -c_t \mathbf{L}_\gamma \mathbf{x}_t dt - 2c_t \sigma^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) dt + \sqrt{2c_t} \sigma d\mathbf{w}_t.$$

- ▶ **Tweedie's formula** turns score estimation into denoising:

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) = \Sigma_t^{-1} (\mathbf{H}_t \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] - \mathbf{x}_t),$$

$$\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = g^*(\mathbf{x}_t), \quad g^* \in \underset{g}{\operatorname{argmin}} \mathbb{E}[\|\mathbf{x}_0 - g(\mathbf{x}_t)\|^2].$$

Since \mathbf{H}_t and Σ_t are known, the hard part is learning the clean-signal predictor g^* .

- ▶ We use a **GCNN denoiser**: a cascade of learnable graph filters, conditioned on t :

$$\hat{\mathbf{x}}_0 = g_{\Theta}(\mathbf{x}_t, \mathbf{L}, t), \quad \Theta^* \in \underset{\Theta}{\operatorname{argmin}} \mathbb{E}[\|\mathbf{x}_0 - g_{\Theta}(\mathbf{x}_t, \mathbf{L}, t)\|^2].$$

Graph filters mix information through neighborhoods; conditioning on t lets the denoiser adapt from coarse recovery to fine-detail reconstruction.

Algorithm 1: Training of GAD

Input: \mathbf{L} , data \mathcal{D} , T , c_t , σ , lr η .

Output: trained denoiser g_{Θ^*} .

1. Sample a clean signal and time:

$$\mathbf{x}_0 \sim \mathcal{D}, \quad t \sim \mathcal{U}(0, T).$$

2. Compute the forward marginal:

$$\begin{aligned} \mathbf{L}_\gamma &= \mathbf{L} + \gamma \mathbf{I}, & \mathbf{H}_t &= e^{-\bar{c}_t \mathbf{L}_\gamma}, \\ \boldsymbol{\mu}_t &= \mathbf{H}_t \mathbf{x}_0, & \boldsymbol{\Sigma}_t &= \sigma^2 (\mathbf{I} - \mathbf{H}_t^2) \mathbf{L}_\gamma^{-1}. \end{aligned}$$

3. Corrupt and denoise:

$$\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t), \quad \hat{\mathbf{x}}_0 = g_{\Theta}(\mathbf{x}_t, \mathbf{L}, t).$$

4. Update the GCNN denoiser:

$$\Theta \leftarrow \Theta - \eta \nabla_{\Theta} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|_2^2.$$

Algorithm 1: Training of GAD

Input: \mathbf{L} , data \mathcal{D} , T , c_t , σ , lr η .

Output: trained denoiser g_{Θ^*} .

1. Sample a clean signal and time:

$$\mathbf{x}_0 \sim \mathcal{D}, \quad t \sim \mathcal{U}(0, T).$$

2. Compute the forward marginal:

$$\begin{aligned} \mathbf{L}_\gamma &= \mathbf{L} + \gamma \mathbf{I}, & \mathbf{H}_t &= e^{-\bar{c}_t \mathbf{L}_\gamma}, \\ \boldsymbol{\mu}_t &= \mathbf{H}_t \mathbf{x}_0, & \boldsymbol{\Sigma}_t &= \sigma^2 (\mathbf{I} - \mathbf{H}_t^2) \mathbf{L}_\gamma^{-1}. \end{aligned}$$

3. Corrupt and denoise:

$$\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t), \quad \hat{\mathbf{x}}_0 = g_{\Theta}(\mathbf{x}_t, \mathbf{L}, t).$$

4. Update the GCNN denoiser:

$$\Theta \leftarrow \Theta - \eta \nabla_{\Theta} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|_2^2.$$

Algorithm 2: Sampling of GAD

Input: \mathbf{L}_γ , σ , grid $0 = t_0 < \dots < t_N = T$, g_{Θ^*} .

Output: generated signal \mathbf{x}_{t_0} .

1. Initialize from the terminal prior:

$$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{L}_\gamma^{-1}).$$

2. For $i = N, \dots, 1$, set

$$t = t_i, \quad t^- = t_{i-1}, \quad \Delta t = t - t^-.$$

3. Estimate \mathbf{x}_0 and the score:

$$\hat{\mathbf{x}}_0 = g_{\Theta^*}(\mathbf{x}_t, \mathbf{L}, t),$$

$$\hat{\mathbf{s}}_t = \boldsymbol{\Sigma}_t^{-1} (\mathbf{H}_t \hat{\mathbf{x}}_0 - \mathbf{x}_t).$$

4. Reverse Euler–Maruyama step:

$$\begin{aligned} \mathbf{x}_{t^-} &\leftarrow \mathbf{x}_t + \left(c_t \mathbf{L}_\gamma \mathbf{x}_t + 2c_t \sigma^2 \hat{\mathbf{s}}_t \right) \Delta t \\ &\quad + \sqrt{2c_t} \sigma \sqrt{\Delta t} \mathbf{z}_i, \quad \mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \end{aligned}$$

5. Return \mathbf{x}_{t_0} .
-

A GSP view: the optimal denoiser is a graph filter

- ▶ To build intuition, look at the denoising problem for a **single** signal at a **fixed** t . The noise model is fully known:

$$\mathbf{x}_t = \mathbf{H}_t \mathbf{x}_0 + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \mathbf{H}_t, \boldsymbol{\Sigma}_t \text{ known.}$$

A GSP view: the optimal denoiser is a graph filter

- ▶ To build intuition, look at the denoising problem for a **single** signal at a **fixed** t . The noise model is fully known:

$$\mathbf{x}_t = \mathbf{H}_t \mathbf{x}_0 + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \mathbf{H}_t, \boldsymbol{\Sigma}_t \text{ known.}$$

- ▶ Many graph signals are **smooth**, so we add a Laplacian smoothing prior:

$$\hat{\mathbf{x}}_0 = g(\mathbf{x}_t) = \underset{\mathbf{x}_0}{\operatorname{argmin}} \|\mathbf{x}_t - \mathbf{H}_t \mathbf{x}_0\|_{\boldsymbol{\Sigma}_t^{-1}}^2 + \alpha \|\mathbf{x}_0\|_{\mathbf{L}}^2.$$

A GSP view: the optimal denoiser is a graph filter

- ▶ To build intuition, look at the denoising problem for a **single** signal at a **fixed** t . The noise model is fully known:

$$\mathbf{x}_t = \mathbf{H}_t \mathbf{x}_0 + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \mathbf{H}_t, \boldsymbol{\Sigma}_t \text{ known.}$$

- ▶ Many graph signals are **smooth**, so we add a Laplacian smoothing prior:

$$\hat{\mathbf{x}}_0 = g(\mathbf{x}_t) = \underset{\mathbf{x}_0}{\operatorname{argmin}} \|\mathbf{x}_t - \mathbf{H}_t \mathbf{x}_0\|_{\boldsymbol{\Sigma}_t^{-1}}^2 + \alpha \|\mathbf{x}_0\|_{\mathbf{L}}^2.$$

Proposition: the solution is a graph filter

The minimizer is $\hat{\mathbf{x}}_0 = h(\mathbf{L})\mathbf{x}_t$, an **ARMA graph filter**.

$$h(\mathbf{L}) = (\mathbf{H}_t^\top \boldsymbol{\Sigma}_t^{-1} \mathbf{H}_t + \alpha \mathbf{L})^{-1} \mathbf{H}_t^\top \boldsymbol{\Sigma}_t^{-1},$$

$$h(\lambda) = \frac{e^{(\lambda+\gamma)\bar{c}_t}}{1 + \alpha\sigma^2 \frac{\lambda}{\lambda+\gamma} (e^{2(\lambda+\gamma)\bar{c}_t} - 1)}.$$

A GSP view: the optimal denoiser is a graph filter

- ▶ To build intuition, look at the denoising problem for a **single** signal at a **fixed** t . The noise model is fully known:

$$\mathbf{x}_t = \mathbf{H}_t \mathbf{x}_0 + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \mathbf{H}_t, \boldsymbol{\Sigma}_t \text{ known.}$$

- ▶ Many graph signals are **smooth**, so we add a Laplacian smoothing prior:

$$\hat{\mathbf{x}}_0 = g(\mathbf{x}_t) = \underset{\mathbf{x}_0}{\operatorname{argmin}} \|\mathbf{x}_t - \mathbf{H}_t \mathbf{x}_0\|_{\boldsymbol{\Sigma}_t^{-1}}^2 + \alpha \|\mathbf{x}_0\|_{\mathbf{L}}^2.$$

Proposition: the solution is a graph filter


The minimizer is $\hat{\mathbf{x}}_0 = h(\mathbf{L})\mathbf{x}_t$, an **ARMA graph filter**.

$$h(\mathbf{L}) = (\mathbf{H}_t^\top \boldsymbol{\Sigma}_t^{-1} \mathbf{H}_t + \alpha \mathbf{L})^{-1} \mathbf{H}_t^\top \boldsymbol{\Sigma}_t^{-1},$$

$$h(\lambda) = \frac{e^{(\lambda+\gamma)\bar{c}_t}}{1 + \alpha\sigma^2 \frac{\lambda}{\lambda+\gamma} (e^{2(\lambda+\gamma)\bar{c}_t} - 1)}.$$

- ▶ The backward process is therefore a **sequence of graph-signal denoising problems** – a classical, well-understood task in GSP.

- ▶ $h(\lambda)$ is a tension between two terms:
 - ⇒ an exponential numerator that replenishes high frequencies
 - ⇒ a doubly-exponential denominator that suppresses them



$$h(\lambda) = \frac{e^{(\lambda+\gamma)\bar{c}_t}}{1 + \alpha\sigma^2 \frac{\lambda}{\lambda+\gamma} (e^{2(\lambda+\gamma)\bar{c}_t} - 1)}$$

Figure. Frequency response $h(\lambda)$ for varying $\alpha\sigma^2$ (left) and \bar{c}_t (right).

- ▶ $h(\lambda)$ is a tension between two terms:
 - ⇒ an exponential numerator that replenishes high frequencies
 - ⇒ a doubly-exponential denominator that suppresses them




Figure. Frequency response $h(\lambda)$ for varying $\alpha\sigma^2$ (left) and \bar{c}_t (right).

$$h(\lambda) = \frac{e^{(\lambda+\gamma)\bar{c}_t}}{1 + \alpha\sigma^2 \frac{\lambda}{\lambda+\gamma} (e^{2(\lambda+\gamma)\bar{c}_t} - 1)}$$

- ▶ **Noise level $\alpha\sigma^2$:**
 - ⇒ small ⇒ numerator wins ⇒ high-pass
 - ⇒ large ⇒ denominator wins ⇒ low-pass
- ▶ **Time \bar{c}_t :**
 - ⇒ large \bar{c}_t ⇒ denominator dominates ⇒ low-pass
 - ⇒ small \bar{c}_t ⇒ near-flat / high-pass

- ▶ $h(\lambda)$ is a tension between two terms:
 - ⇒ an exponential numerator that replenishes high frequencies
 - ⇒ a doubly-exponential denominator that suppresses them




Figure. Frequency response $h(\lambda)$ for varying $\alpha\sigma^2$ (left) and \bar{c}_t (right).

$$h(\lambda) = \frac{e^{(\lambda+\gamma)\bar{c}_t}}{1 + \alpha\sigma^2 \frac{\lambda}{\lambda+\gamma} (e^{2(\lambda+\gamma)\bar{c}_t} - 1)}$$

- ▶ **Noise level $\alpha\sigma^2$:**
 - ⇒ small ⇒ numerator wins ⇒ high-pass
 - ⇒ large ⇒ denominator wins ⇒ low-pass
- ▶ **Time \bar{c}_t :**
 - ⇒ large \bar{c}_t ⇒ denominator dominates ⇒ low-pass
 - ⇒ small \bar{c}_t ⇒ near-flat / high-pass
- ▶ **Takeaway:** generation unfolds coarse-to-fine: recover low-frequency structure first, then reintroduce high-frequency detail; this motivates a learnable GCNN that adapts across t .

- ▶ **Datasets** (one fixed graph each):
 - ⇒ **Synthetic SBM** – two communities of 10 nodes, smooth signals
 - ⇒ **METR-LA** – traffic speeds at 207 sensors on the LA road network
 - ⇒ **Molene** – daily temperatures at 37 weather stations in France
- ▶ **Baselines**: variance-preserving (VPD) and variance-exploding (VED) diffusion
 - ⇒ both are **graph-agnostic**; under the SDE view they are directly comparable to GAD

- ▶ **Evaluation** (following graph-generation practice): three signal statistics

$$QV(\mathbf{x}) = \mathbf{x}^\top \mathbf{L} \mathbf{x}, \quad SC(\mathbf{x}) = \frac{\sum_i \lambda_i |\tilde{x}_i|^2}{\sum_i |\tilde{x}_i|^2}, \quad DC(\mathbf{x}) = \text{corr}(\mathbf{x}, \mathbf{d}),$$

smoothness (QV), spectrum (SC), degree alignment (DC) \Rightarrow MMD per metric \Rightarrow averaged into **aMMD** (lower is better).

- ▶ **Sampling**: Euler–Maruyama; each step is one GCNN pass, so we report quality vs. the **number of steps**




Figure. aMMD (lower is better) across sampling budgets.




Figure. aMMD (lower is better) across sampling budgets.

- ▶ **Main takeaway:** GAD reaches low aMMD with far fewer reverse-SDE steps than graph-agnostic VPD/VED.

The gain is strongest in the low-step regime, where each saved step avoids one GCNN pass.

- ▶ **GAD:** a **graph-aware** diffusion model for generating signals on a known graph
- ▶ **Forward:** heat equation with the graph in the **drift**; closed-form GMRF marginals; converges to a **graph-dependent** stationary distribution; FCPS controls the noising rate
- ▶ **Backward:** via Tweedie, score estimation reduces to a **sequence of graph-signal denoising** problems; spectral analysis reveals a **coarse-to-fine** filter \Rightarrow a learnable GCNN denoiser
- ▶ **Empirics:** consistent gains over graph-agnostic baselines, especially with **few sampling steps**

- ▶ **GAD:** a **graph-aware** diffusion model for generating signals on a known graph
- ▶ **Forward:** heat equation with the graph in the **drift**; closed-form GMRF marginals; converges to a **graph-dependent** stationary distribution; FCPS controls the noising rate
- ▶ **Backward:** via Tweedie, score estimation reduces to a **sequence of graph-signal denoising** problems; spectral analysis reveals a **coarse-to-fine** filter \Rightarrow a learnable GCNN denoiser
- ▶ **Empirics:** consistent gains over graph-agnostic baselines, especially with **few sampling steps**
- ▶ **Where next:** multi-graph / transfer to larger graphs, per-eigenvalue schedulers, advanced SDE solvers

Thank you!

Questions?

► **Graph-Aware Diffusion for Signal Generation**

⇒ code: github.com/vimalkb7/gad

⇒ contact: vkumarasamybal@tudelft.nl



scan for the paper

Backup material

Backup: where is graph-signal generation useful?

- ▶ **Sensor / monitoring networks:** generate plausible temperature, traffic-speed, or air-quality fields over a fixed station graph – data augmentation, missing-value imputation, scenario simulation
- ▶ **Probabilistic forecasting:** sample multiple future graph signals to capture uncertainty in traffic, weather, energy demand, or sensor dynamics
- ▶ **Recommender systems:** synthesize user–item rating signals over a known user/item graph
- ▶ **Wireless / power networks:** simulate signals over a fixed infrastructure topology
- ▶ **Privacy & stress-testing:** produce realistic-yet-synthetic graph data when sharing the real signals is restricted
- ▶ The common thread: the **graph is fixed and known**, and we want new *signals* that respect it

Backup: the centering parameter γ

- ▶ \mathbf{L} has a **zero eigenvalue**, so $-c_t \mathbf{L}$ does not converge. Adding $\gamma \mathbf{I}$ shifts the spectrum:
$$\mathbf{L}_\gamma = \mathbf{L} + \gamma \mathbf{I} \Rightarrow -c_t \mathbf{L}_\gamma \text{ is Hurwitz} \Rightarrow \text{convergence to } \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{L}_\gamma^{-1}).$$

Backup: the centering parameter γ

- ▶ \mathbf{L} has a **zero eigenvalue**, so $-c_t \mathbf{L}$ does not converge. Adding $\gamma \mathbf{I}$ shifts the spectrum:

$$\mathbf{L}_\gamma = \mathbf{L} + \gamma \mathbf{I} \Rightarrow -c_t \mathbf{L}_\gamma \text{ is Hurwitz} \Rightarrow \text{convergence to } \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{L}_\gamma^{-1}).$$

- ▶ A genuine **trade-off** in γ :
 - \Rightarrow **too small** – low modes decay slowly; \mathbf{L}_γ ill-conditioned; high-variance Σ_T
 - \Rightarrow **too large** – the geometry induced by the Laplacian is washed out

Backup: the centering parameter γ

- ▶ \mathbf{L} has a **zero eigenvalue**, so $-c_t \mathbf{L}$ does not converge. Adding $\gamma \mathbf{I}$ shifts the spectrum:

$$\mathbf{L}_\gamma = \mathbf{L} + \gamma \mathbf{I} \Rightarrow -c_t \mathbf{L}_\gamma \text{ is Hurwitz} \Rightarrow \text{convergence to } \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{L}_\gamma^{-1}).$$
- ▶ A genuine **trade-off** in γ :
 - \Rightarrow **too small** – low modes decay slowly; \mathbf{L}_γ ill-conditioned; high-variance Σ_T
 - \Rightarrow **too large** – the geometry induced by the Laplacian is washed out

Remark (numerical stability)

We never form the matrix exponential directly; we use the eigendecomposition $\mathbf{H}_t = \mathbf{V} e^{-\bar{c}_t (\mathbf{L} + \gamma \mathbf{I})} \mathbf{V}^\top$. This also hints at defining \bar{c}_t **per eigenvalue** for independent mode control – a promising direction.

- ▶ The closed-form filter $h(\mathbf{L})$ is optimal only for a **single** signal at a **single** t
- ▶ In practice the denoiser must handle:
 - ⇒ \mathbf{x}_0 **in expectation** over p_0 (the MMSE objective), not one signal
 - ⇒ **every** noise level t , where the optimal filter sweeps low-pass \rightarrow high-pass

Backup: why a GCNN denoiser?

- ▶ The closed-form filter $h(\mathbf{L})$ is optimal only for a **single** signal at a **single** t
- ▶ In practice the denoiser must handle:
 - ⇒ \mathbf{x}_0 **in expectation** over p_0 (the MMSE objective), not one signal
 - ⇒ **every** noise level t , where the optimal filter sweeps low-pass \rightarrow high-pass
- ▶ A **GCNN** – a cascade of graph filters with nonlinearities – is the natural learnable generalization of $h(\mathbf{L})$
 - ⇒ expressive enough to start from pure noise and adapt across t
 - ⇒ **permutation-equivariant**, stable to perturbations, and transferable across graphs and timesteps

- ▶ Each Euler–Maruyama step calls the GCNN once, converts its prediction into the Tweedie score, and takes one reverse-SDE step; iterating from $\mathbf{x}_T \sim p_\infty$ yields a sample \mathbf{x}_0

