# **Riemannian Flow Matching for InSAR**

G. Le Bellier, D. El Hajjar, A. Breloy, N. Audebert



Séminaire TéSA, 2025

# In this talk

- Interferometric SAR (InSAR)
- Diffusion Models and flow matching
- Riemannian Flow matching for InSAR
  - Interferometric phase denoising
  - Sampling realistic interferograms



### Interferometric SAR

Diffusion models and flow matching

Riemannian flow matching for InSAR

# **Motivations**



## Monitoring ground subsidence

- Large scale and long term
- Ground-based methods (GPS, topography)
  - · Costly
  - · Limited spatial coverage

# Interferometry with SAR (InSAR)

- Covers large areas
- Systematic acquisitions
- · Unaffected by weather
- Sub-centimeter accuracy

## SAR produces complex valued images





modulus

phase

#### Two images with same position at different time

#### Phase difference contains information about surface displacement

$$\Delta \Phi = \Delta \Phi_{displ} + \underbrace{\Delta \Phi_{topo} + \Delta \Phi_{orb} + \Delta \Phi_{atm}}_{\bullet} + \Delta \Phi_{noise}$$

removed by pre-processing



# **Multi-temporal InSAR**

#### 2-Pass InSAR

date 1



- · Only 2 SAR images
- · Loss of coherence over time
- · Noisy phase difference

#### **Multi-temporal InSAR**



- · SAR image time series
- · Leverages redundancy
- · Requires models & optimization

# Multi-temporal InSAR covariance modelling



### **DS-InSAR assumptions**

Zero-mean

$$\mathbb{E}[x^{\ell}] = 0, \forall \ell \in \llbracket 1, p \rrbracket$$

Multidimensional pixels (p images)

$$\mathbf{x} = [x^1, \cdots, x^p] \in \mathbb{C}^p$$

Spatial window (n pixels)

$$\{\mathbf{x}_i\}_{i=1}^n \in (\mathbb{C}^p)^n$$

• Structured correlations w.r.t. time

$$\mathbb{E}[x^q(x^\ell)^*] = \varphi_{q\ell} \quad \underbrace{e^{j \ (\theta_q - \theta_\ell)}}_{\text{phase diff.}})$$

 $\Rightarrow$  not any local covariance matrix!

### Structured modulus-argument decomposition

$$\mathbb{E}\left[\mathbf{x}\mathbf{x}^{H}\right] = \mathbf{\Sigma} = \operatorname{mod}(\mathbf{\Sigma}) \odot \operatorname{arg}(\mathbf{\Sigma}) \stackrel{\Delta}{=} \mathbf{\Psi} \odot (\mathbf{w}\mathbf{w}^{H})$$

with

- $\Psi$  : coherence and variance
- +  $\mathbf{w} = [\mathrm{e}^{\mathrm{j}\theta_1},...,\mathrm{e}^{\mathrm{j}\theta_\mathrm{p}}]^\top$  : phase at each time

This structure implies the phase-closure property :

$$\arg(\mathbf{\Sigma}_{q\ell}) + \arg(\mathbf{\Sigma}_{\ell j}) + \arg(\mathbf{\Sigma}_{jq}) = 0$$

# Interferometric phase linking problem

From pixel patch  $\{\mathbf{x}_i\}_{i=1}^n$ , estimate  $\mathbf{w}$ 

## Two main IPL approaches

Maximum likelihood estimation

assume  $\mathbf{x} \sim \mathcal{CN}(\mathbf{0}, \boldsymbol{\Psi} \odot \mathbf{w} \mathbf{w}^{\mathit{H}})$  and roll

Covariance matrix fitting

force phase closure on plug-in estimate

$$\min_{\mathbf{w}\in\mathbb{T}_p} \quad d^2(\hat{\boldsymbol{\Sigma}}, \ \mathrm{mod}(\hat{\boldsymbol{\Sigma}})\odot\mathbf{w}\mathbf{w}^H)$$

# Generative models for InSAR : two goals

# Sampling



#### Denoising







### Interferometric SAR

# Diffusion models and flow matching

Riemannian flow matching for InSAR

Goal : generate new samples similar to existing ones

Probabilistic view : data was sampled from an underlying density

#### What we want :

- data probability density function  $p_{\text{data}}(\mathbf{x})$
- a way to sample data as  $\mathbf{x} \sim p_{\mathrm{data}}(\mathbf{x})$

#### What we have :

- samples  $\{\mathbf{x}_i\}_{i=1}^n$
- problems

# Diffusion models, overview

Diffusion models : "long sequence of small denoisers"



#### The score matching view of diffusion

- Step 1 : Learning models with score matching
- Step 2 : Sampling with Langevin dynamics

next slides mostly inspired (if not taken) from Stanford-CS236

# **Score function**

### Definition

The score  $s: \mathbb{R}^d \to \mathbb{R}^d$  of a distribution p is

 $s(\mathbf{x}) = \nabla_{\mathbf{x}} \log p(\mathbf{x})$ 

"assigns to each x the steepest direction to increase likelihood"





# Score matching

- Samples  $\{\mathbf{x}_i\}_{i=1}^n \sim p_{\text{data}}(\mathbf{x})$
- Parametric function  $s_{ heta}: \mathbb{R}^d 
  ightarrow \mathbb{R}^d$  (neural netork)
- Find  $\theta$  such that  $s_{\theta}(\mathbf{x}) \simeq \nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x})$

# Score matching with Fisher divergence

$$\mathcal{L}(\theta) = \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \left[ ||\nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x}) - s_{\theta}(\mathbf{x})||_{2}^{2} \right]$$
$$= \frac{1}{2} \sum_{i=1}^{n} ||\nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x}) - s_{\theta}(\mathbf{x})||_{2}^{2}$$

It depends on  $p_{\text{data}}$  :(

#### Theorem

(Hyvärinen 2005)

Under regularity assumptions

$$\mathcal{L}(\theta) = \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \left[ || \nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x}) - s_{\theta}(\mathbf{x}) ||_{2}^{2} \right]$$

$$= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \left[ \frac{1}{2} || s_{\theta}(\mathbf{x}) ||_{2}^{2} + \text{tr}(\nabla_{\mathbf{x}} s_{\theta}(\mathbf{x})) \right] + \text{const.}$$

where  $abla_{\mathbf{x}} s_{ heta}(\mathbf{x})$  is the Jacobian of  $s_{ heta}(\mathbf{x})$ 

- Does not depend on  $p_{data}\,$  :)
- Still, evaluationg  $abla_{x} s_{\theta}(x)$  is costly  $\mathcal{O}(\mathit{d})$  backpropagations
- Can be well approximated by denoising score matching

# **Denoising scroe matching**

$$\begin{array}{c|c} & p_{\text{data}}(\mathbf{x}) \\ \mathbf{X} & & \\ q_{\sigma}(\tilde{\mathbf{x}} \mid \mathbf{x}) \\ q_{\sigma}(\tilde{\mathbf{x} \mid \mathbf{x}) \\ q_{\sigma}(\tilde$$

 $\tilde{\mathbf{x}}$ 

# Score matching on $q_{\sigma}$

(Vincent 2011)

$$\begin{split} \tilde{\mathcal{L}}(\theta) &= \frac{1}{2} \mathbb{E}_{\tilde{\mathbf{x}} \sim q_{\sigma}} \left[ || \nabla_{\tilde{\mathbf{x}}} \log q_{\sigma}(\tilde{\mathbf{x}}) - s_{\theta}(\tilde{\mathbf{x}}) ||_{2}^{2} \right] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}, \tilde{\mathbf{x}} \sim q_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x})} \left[ || s_{\theta}(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log q_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}) ||_{2}^{2} \right] + \text{const.} \end{split}$$

If  $q_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}) = \mathcal{N}(\mathbf{x}, \sigma^2 \mathbf{I})$ , everything gets simple

$$abla_{ ilde{\mathbf{x}}} \log q_{\sigma}( ilde{\mathbf{x}}|\mathbf{x}) = rac{ ilde{\mathbf{x}} - \mathbf{x}}{\sigma^2}$$



The score model simply tries to predict the added noise!

#### Score matching w.r.t. noise level

Overload model : noise level  $\sigma$  is an input

$$s_{\theta}(\mathbf{x},\sigma): \mathbb{R}^d \times \mathbb{R}^+ \longrightarrow \mathbb{R}^d$$

Multi-scale loss

$$\begin{aligned} \mathcal{L}(\theta) &= \frac{1}{L} \sum_{\ell=1}^{L} \lambda(\sigma_{\ell}) \mathbb{E}_{q_{\sigma_{\ell}}(\tilde{\mathbf{x}})} \left[ ||\nabla_{\tilde{\mathbf{x}}} \log q_{\sigma_{\ell}}(\tilde{\mathbf{x}}) - s_{\theta}(\tilde{\mathbf{x}}, \sigma_{\ell})||_{2}^{2} \right] \\ &= \frac{1}{L} \sum_{\ell=1}^{L} \lambda(\sigma_{\ell}) \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}, \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[ ||s_{\theta}(\mathbf{x} + \sigma_{\ell} \mathbf{z}, \sigma_{\ell}) - \mathbf{z}/\sigma_{\ell}||_{2}^{2} \right] \end{aligned}$$

#### At this point, step 1 is done!

**Step 2** : how to sample data from learned  $s_{\theta}(\mathbf{x}, \sigma)$ ?

# Sampling with Langevin dynamics

### Langevin dynamics

- $\mathbf{x}_0 \sim \pi(\mathbf{x}) \leftarrow \text{random sampling}$
- Follow T iterations of

$$\mathbf{x}_t \leftarrow \mathbf{x}_t + \frac{\epsilon}{2} \nabla_{\mathbf{x}} \log p(\mathbf{x}_{t-1}) + \epsilon \mathbf{z}_t$$

if 
$$\epsilon << 0$$
 and  $T 
ightarrow \infty$ , then  $\mathbf{x}_{T} \sim p$ 



#### Two steps :

- Learn score  $s_{\theta}$  on  $\{\mathbf{x}_i\}_{i=1}^n$  with denoising score matching
- Uste the learned score within Langevin dynamics

$$\mathbf{x}_t \leftarrow \mathbf{x}_t + \frac{\epsilon}{2} s_{\theta}(\mathbf{x}) + \epsilon \mathbf{z}_t, \quad t \in \llbracket 1, T \rrbracket$$

• If  $s_{\theta}(\mathbf{x}) \simeq \nabla_{\mathbf{x}} \log p_{\mathrm{data}}(\mathbf{x})$ , then we approximate  $\mathbf{x}_{T} \sim p_{\mathrm{data}}$ !

#### In practice

- We learn a score conditioned to the scale  $s_{\theta}(\mathbf{x}, \sigma)$
- Dynamics is **annealed** within a loop over  $\sigma \downarrow$

# Algorithm 1 Annealed Langevin dynamics.

```
Require: \{\sigma_i\}_{i=1}^L, \epsilon, T.
   1: Initialize \tilde{\mathbf{x}}_0
   2: for i \leftarrow 1 to L do
  3: \alpha_i \leftarrow \epsilon \cdot \sigma_i^2 / \sigma_I^2 \qquad \triangleright \alpha_i is the step size.
  4: for t \leftarrow 1 to T do
                           Draw \mathbf{z}_t \sim \mathcal{N}(0, I)
\tilde{\mathbf{x}}_t \leftarrow \tilde{\mathbf{x}}_{t-1} + \frac{\alpha_i}{2} \mathbf{s}_{\boldsymbol{\theta}}(\tilde{\mathbf{x}}_{t-1}, \sigma_i) + \sqrt{\alpha_i} \mathbf{z}_t
   5:
  6:
  7:
                   end for
         \tilde{\mathbf{x}}_0 \leftarrow \tilde{\mathbf{x}}_T
   8:
   9: end for
          return \tilde{\mathbf{x}}_T
```



Sampling  $\simeq$  inverting a stochastic differential equation (SDE)

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

From there, what is **flow matching**?

- Learn another transport from  $\pi()$  to  $p_{data}$
- Find an ordinary differential equation (ODE) deterministic transformation !

# Flow matching - goal



We seek a velocity field  $u_t$  such that

$$\frac{d}{dt}\varphi_t(x) = u_t(\varphi_t(x)), \quad \forall t \in [0, 1]$$
(1)

for which  $\varphi_0 = \pi$  and  $\varphi_0 = p_{data}$ 

Learn a function  $v_{\theta}(t, x)$  with the algorithm

- Sample  $z_0 \sim \mathcal{N}(0, \mathbf{I})$ ,  $z_1 \sim p_{data}$ , and  $t \sim \mathcal{U}(0, 1)$
- Stochastic gradient descent on  $\theta$  with loss

$$\left\| v_{\theta}(t, x_t) - \frac{z_1 - z_0}{1 - t} \right\|^2$$

Looks like score matching, but we predict a direction rather than noise

#### What does it do on average?

$$\mathcal{L}_{\mathsf{FM}}(\theta) = \mathbb{E}_{\substack{x_0, x_1 \sim p(x_0, x_1) \\ t \sim U([0, 1])}} \left\| v_{\theta}(t, x_t) - u_t(x_t \mid x_0, x_1) \right\|^2$$

Visualize with the following [blog-post] or [playground]

### Sampling from the model

$$\hat{x}_1 = \varphi_{t=1}(x_0) = \mathsf{EDO}^{v_\theta}(x_0, 0 \to 1)$$

where EDO<sup> $v_{\theta}$ </sup>( $\cdot, t_0 \rightarrow t_1$ ) solves (1) from  $t_0$  to  $t_1$  with  $v_{\theta}(t, \cdot)$  instead of  $u_t$ 



Interferometric SAR

Diffusion models and flow matching

Riemannian flow matching for InSAR

### An image of InSAR phases

- $\Leftrightarrow d \text{ pixels that are } \mathrm{mod} 2\pi$ 
  - $\Leftrightarrow d \text{ complex phases}$ 
    - $\Leftrightarrow d \text{ circles}$ 
      - $\Leftrightarrow$  a *d*-torus



## A d-torus with the Euclidean metric is a Riemannian Manifold

Flow matching for distributions on Riemannian manifolds exists!

RFM : equations are more intricate but conceptually the same !

Velocities lie on the tangent space



Straight lines between two points are geodesics

$$\mathbf{w}_t = \exp_{\mathbf{w}_1} \left( \kappa(t) \log_{\mathbf{w}_1}(\mathbf{w}_0) \right)$$

With simple operators for the torus 
$$\begin{cases} \exp_w(u) = (w+u) \pmod{2\pi} \\ \log_{wa}(w_b) = \arctan2(\sin(w_b - w_a), \cos(w_b - w_a)) \end{cases}$$

## Learning

$$\mathcal{L}_{\mathsf{RFM}}(\theta) = \mathbb{E}_{\mathbf{w}_{0}, \mathbf{w}_{1} \sim p(\mathbf{w}_{0}, \mathbf{w}_{1})} \left\| \mathbf{v}_{\theta}(t, \mathbf{w}_{t}) - \mathbf{u}_{t}(\mathbf{w}_{t} \mid \mathbf{w}_{0}, \mathbf{w}_{1}) \right\|_{g}^{2}$$

$$t \sim U([0,1])$$

### Sampling (inference)

$$\hat{\mathbf{w}}_1 = \varphi_{t=1}(\mathbf{w}_0) = \mathsf{EDO}_{\mathbb{T}_d}^{v_\theta}(\mathbf{w}_0, 0 \to 1)$$
(2)

Dataset : Mexico city seen from Sentinel-1

- · 2 cubes of 40 SAR image phases
  - Raw (2p-InSAR)  $\sim p_{data}$
  - Clean (MT-InSAR)  $\sim p_{data}^{clean}$
- Every 12 day between 14/08/19 and 6/12/20
- $\simeq$  3.6k × 16k pixels  $\longrightarrow$  many patches
- 40(39)/2 = 780 possible interferograms

# Flow matching tasks

- Generation : map  $\mathcal{U}(\mathbb{T}_d) \longrightarrow p_{data}$
- **Denoising** : map  $p_{data} \longrightarrow p_{data}^{clean}$





# Preliminary results - generation (1/2)



# Preliminary results - generation (2/2)



# Raw $\mathbf{w}_0$ Denoised $\varphi_1(\mathbf{w}_0)$ Clean $\mathbf{w}_1$



