

Gaussian process regression for high dimensional graph inputs

Raphaël Carpintero Perez

Sébastien Da Veiga
Josselin Garnier
Brian Staber

12/10/2023



Introduction

Graph kernels

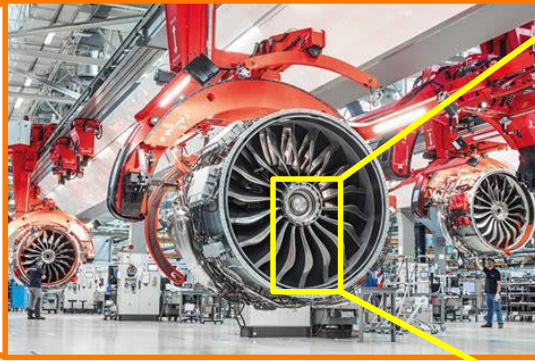
Sliced Wasserstein Weisfeiler Lehman (SWWL)

Conclusion and future work

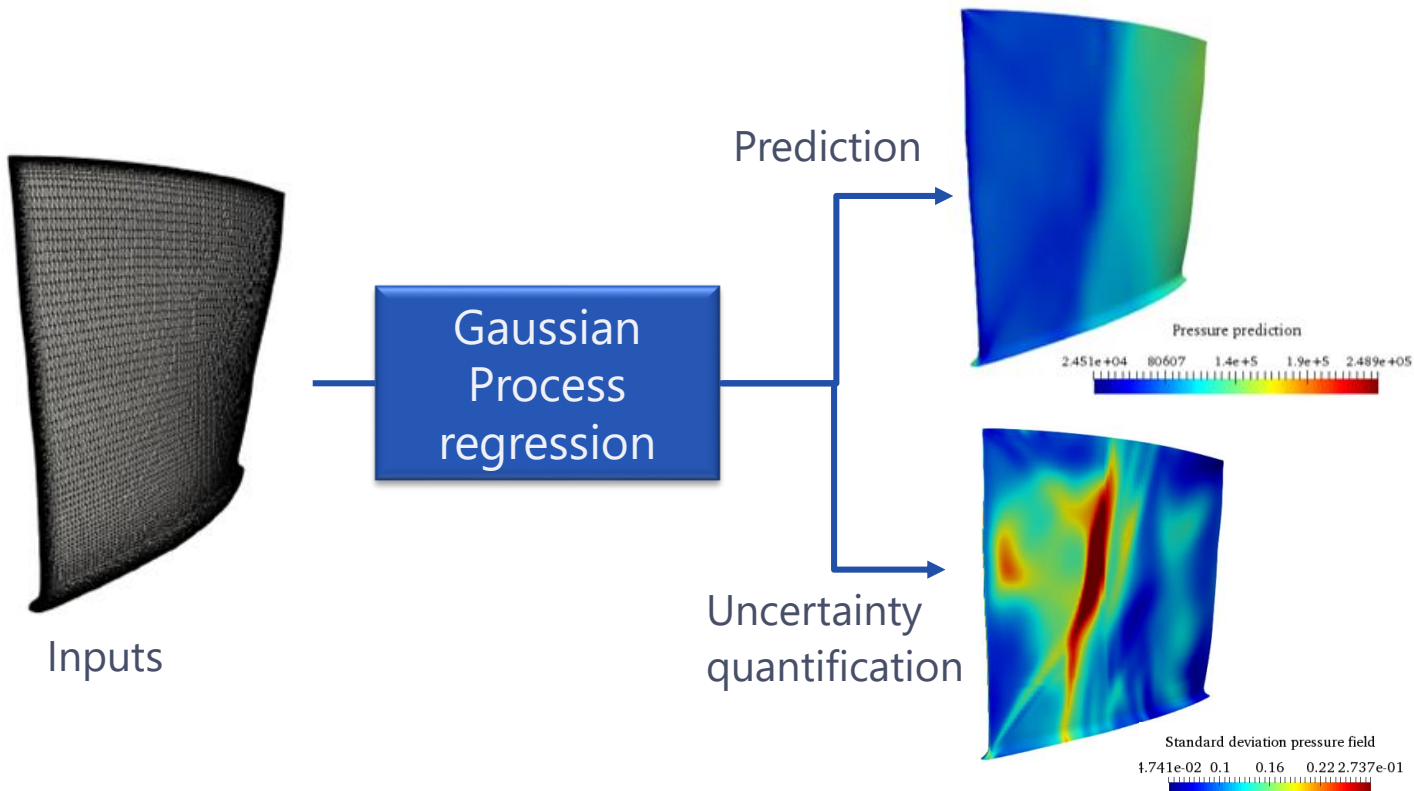
Introduction



Objectives



Objectives



Inputs and outputs

▪ Graph inputs

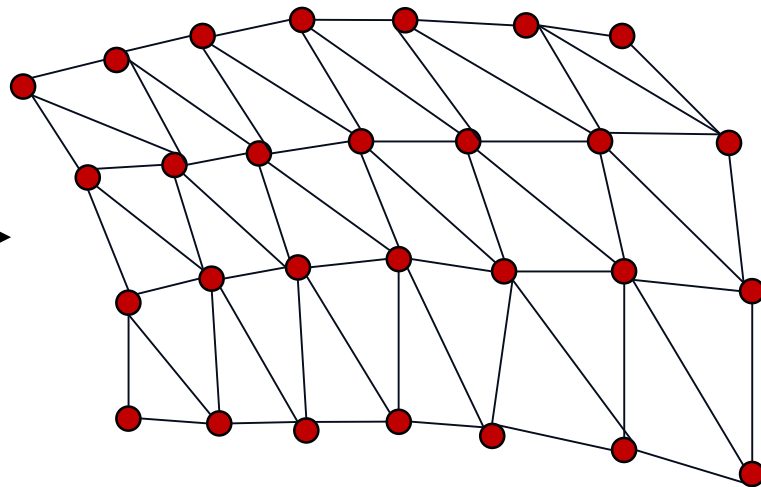
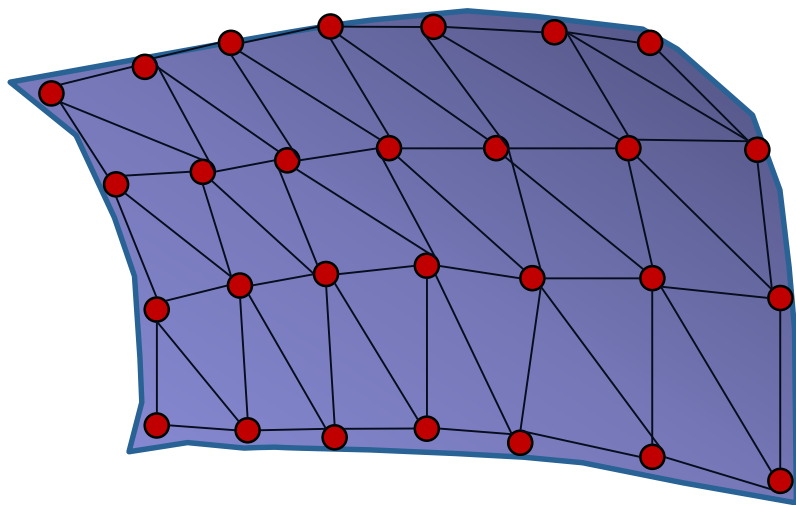
- Mesh → Graph structure
- 3D coordinates for all nodes

▪ Scalar inputs

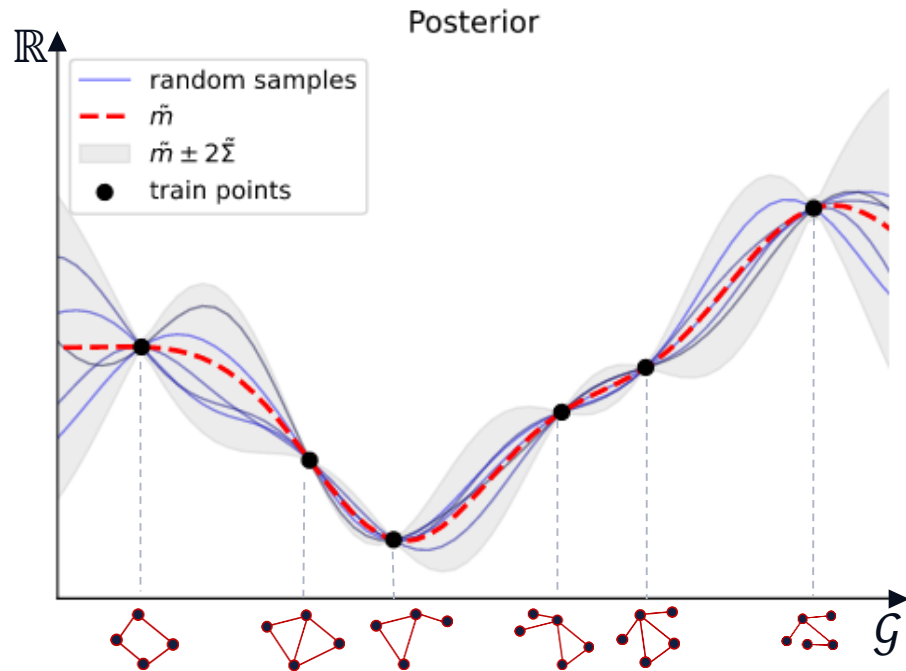
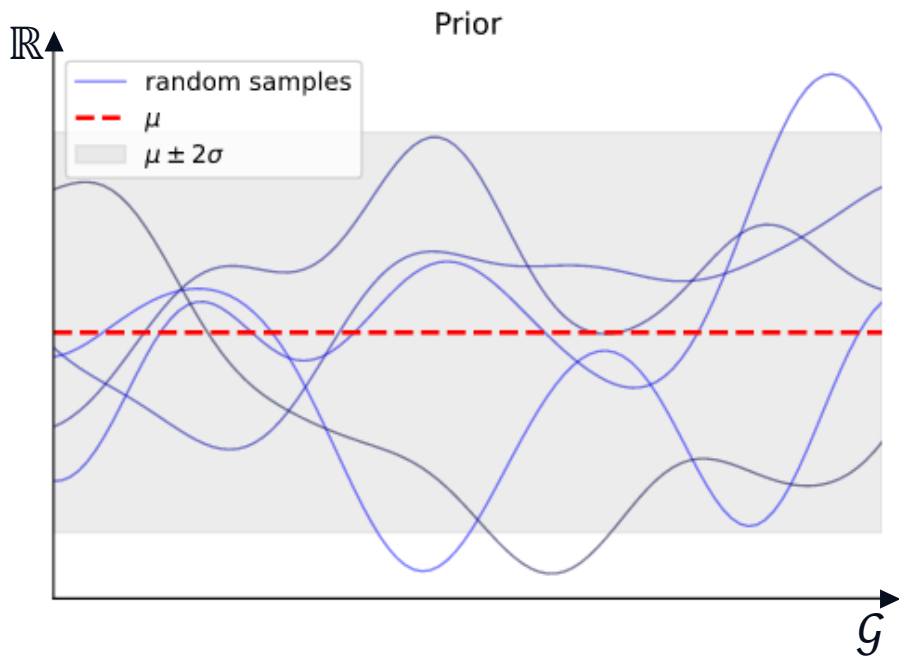
- Pressure
- Speed of rotation

▪ Scalar outputs

- Physical quantities

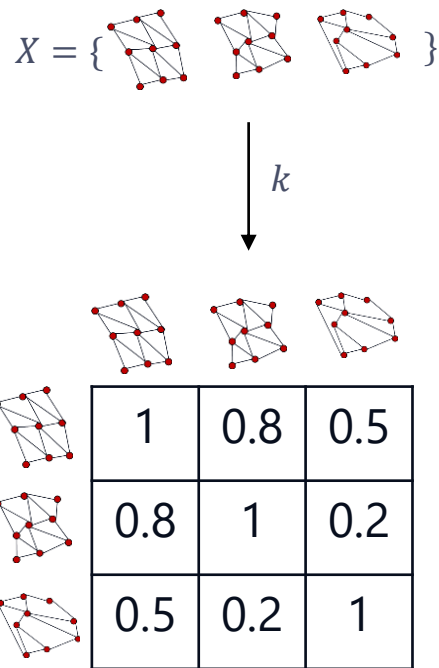


Gaussian process regression



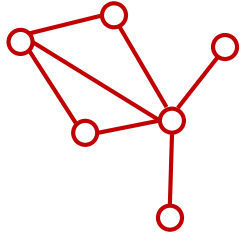
Gaussian process regression

- $X = (G_1, \dots, G_N)^T$ with $G_i \in \Gamma$ (**train input graphs**)
- $Y = (y_1, \dots, y_N)^T$, $y_i \in \mathbb{R}$ (scalar outputs)
- Observations: $y_i = f(G_i) + \epsilon_i$ where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$
 $f: \Gamma \rightarrow \mathbb{R}$
- $\bar{f} = (f(G_1), \dots, f(G_N))^T$
- **Gaussian prior** over functions: $\bar{f} | G_1, \dots, G_N \sim \mathcal{N}(0, K^{ff})$
- $K^{ff}: N \times N$ covariance matrix where $K_{ij}^{ff} = k(G_i, G_j)$
- and $k: \Gamma \times \Gamma \rightarrow \mathbb{R}$ is a **positive definite kernel**
- Question: how to choose k ?

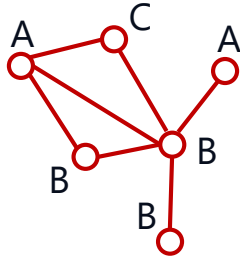


What is a graph ?

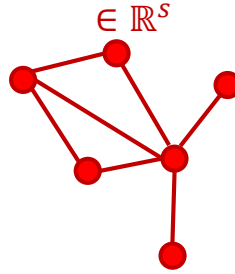
What is a graph ?



Case 1 :
Vertices + Edges

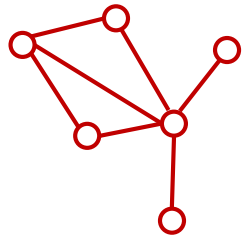


Case 2 :
Vertices + Edges
+ Node labels

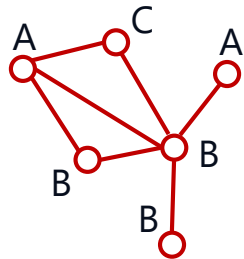


Case 3 :
Vertices + Edges
+ Node attributes

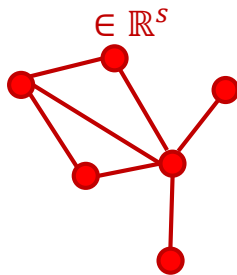
What is a graph ?



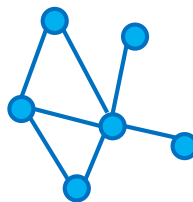
Case 1 :
Vertices + Edges



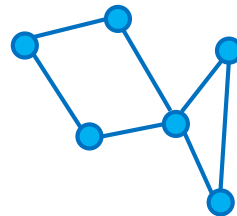
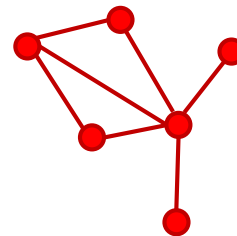
Case 2 :
Vertices + Edges
+ Node labels



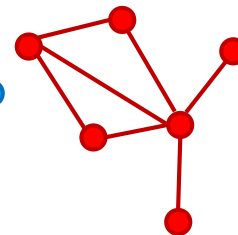
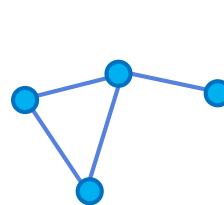
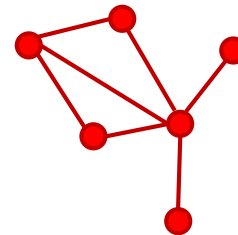
Case 3 :
Vertices + Edges
+ Node attributes



Case 3A: Fixed structure -> signal

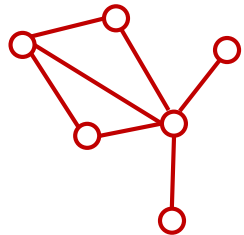


Case 3B: Fixed number of nodes

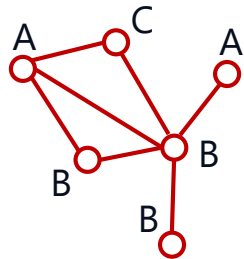


Case 3C: Varying number of
nodes + structure + attributes

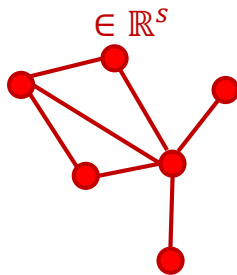
What is a graph ?



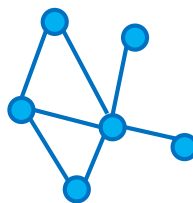
Case 1 :
Vertices + Edges



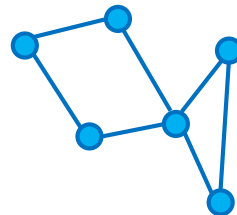
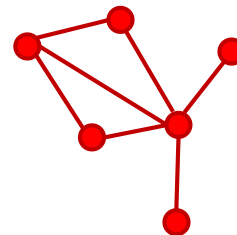
Case 2 :
Vertices + Edges
+ Node labels



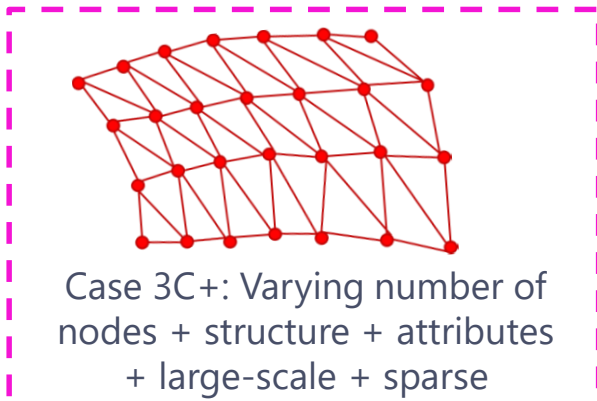
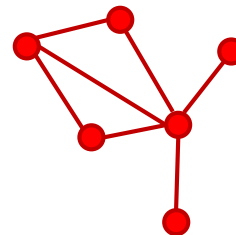
Case 3 :
Vertices + Edges
+ Node attributes



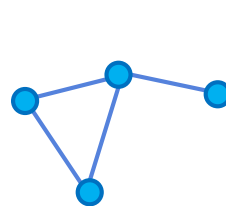
Case 3A: Fixed structure -> signal



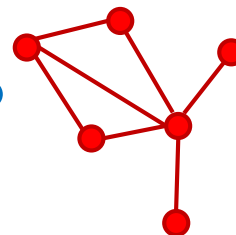
Case 3B: Fixed number of nodes



Case 3C+: Varying number of
nodes + structure + attributes
+ large-scale + sparse



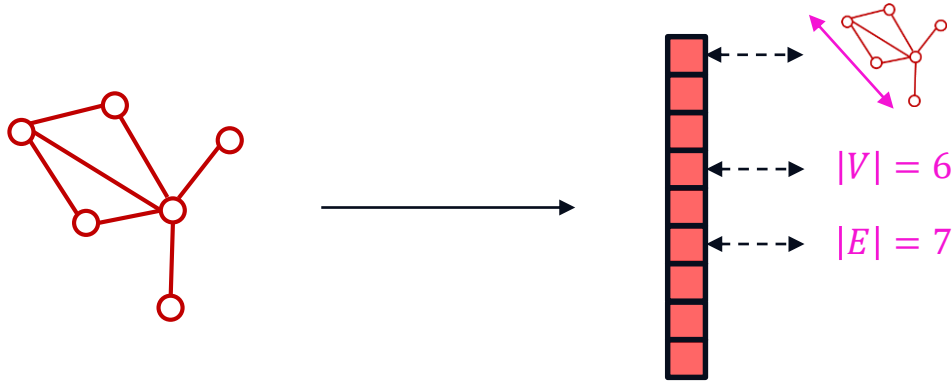
Case 3C: Varying number of
nodes + structure + attributes



Graph kernels

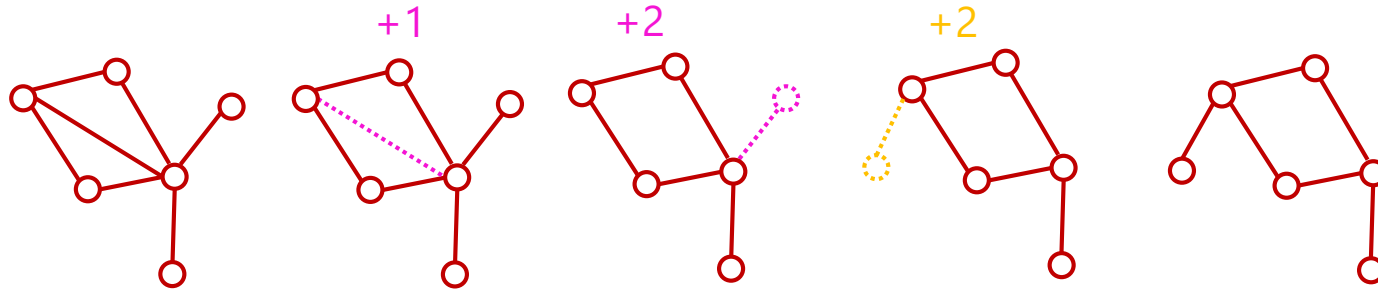


Invariants / Topological descriptors



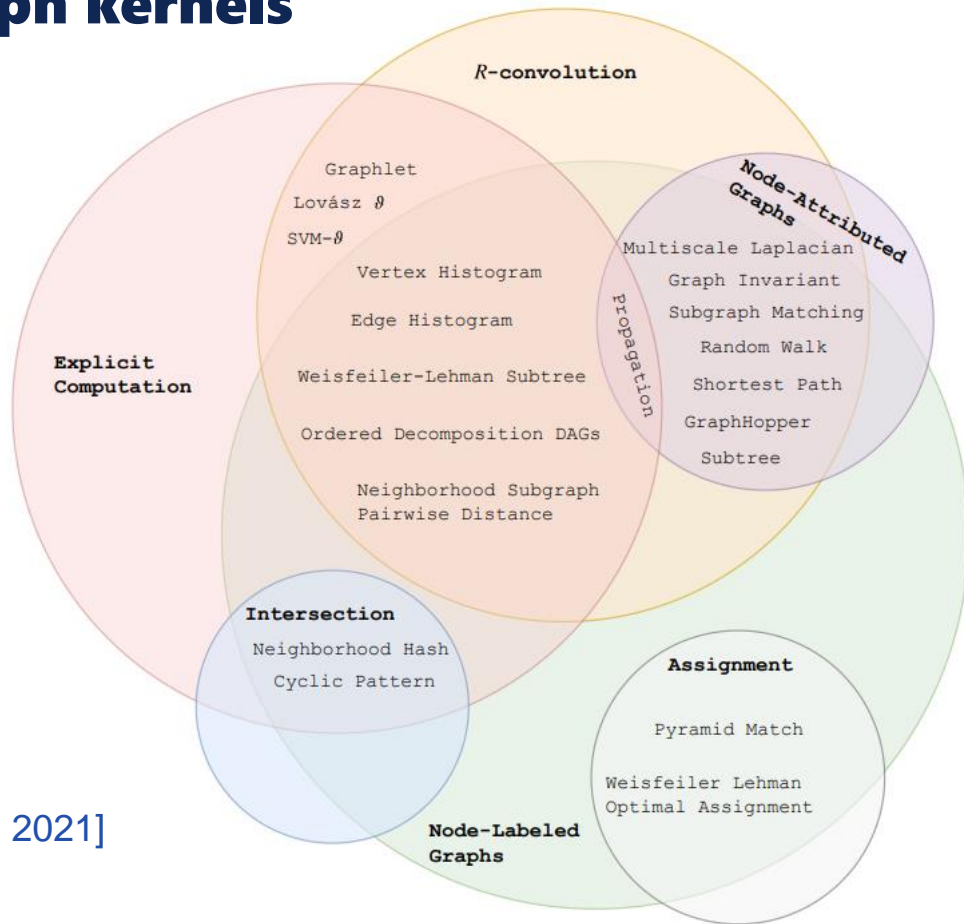
- Map the graph to a vectorial representation
- Invariants: do not change under graph isomorphism (diameter, average clustering coefficient, ...)
- Complete invariants require exponential time

Graph edit distance



- $d(G_1, G_2)$ = minimal number of operations to transform G_1 in G_2 (adding/removing an edge/vertex, node relabeling)
- NP-complete
- Not suited for node-attributed graphs...

Taxonomy of graph kernels



Taxonomy of graph kernels.

Figure from [Nikolentzos et al., 2021]

\mathcal{R} -convolution kernels

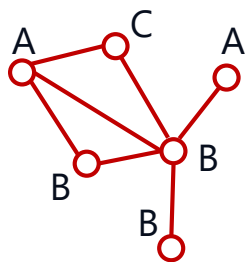
- $\mathcal{R}(g_1, \dots, g_d, G) : \mathcal{R}$ -decomposition where g_i is a 'part' of G (relationship)
- $\mathcal{R}^{-1}(G) = \{g := (g_1, \dots, g_d) \mid \mathcal{R}(g_1, \dots, g_d, G)\}$: pre-image of the relation
- Let k_i a base kernel based on a subset of the parts denoted G_i .
- The \mathcal{R} -convolution kernel between G and G' is defined as

$$k_{\mathcal{R}}(G, G') := \sum_{g \in \mathcal{R}^{-1}(G)} \sum_{g' \in \mathcal{R}^{-1}(G')} \prod_{i=1}^d k_i(g_i, g'_i)$$

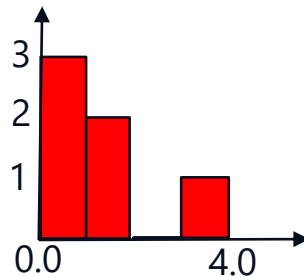
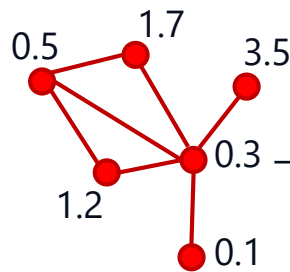
All node-pairs kernel / node histogram kernel

$k_N(G, G') := \sum_{v \in V} \sum_{v' \in V'} k_{node}(v, v')$ where k_{node} is a positive definite kernel between node attributes/labels \rightarrow feature map ϕ_{node}

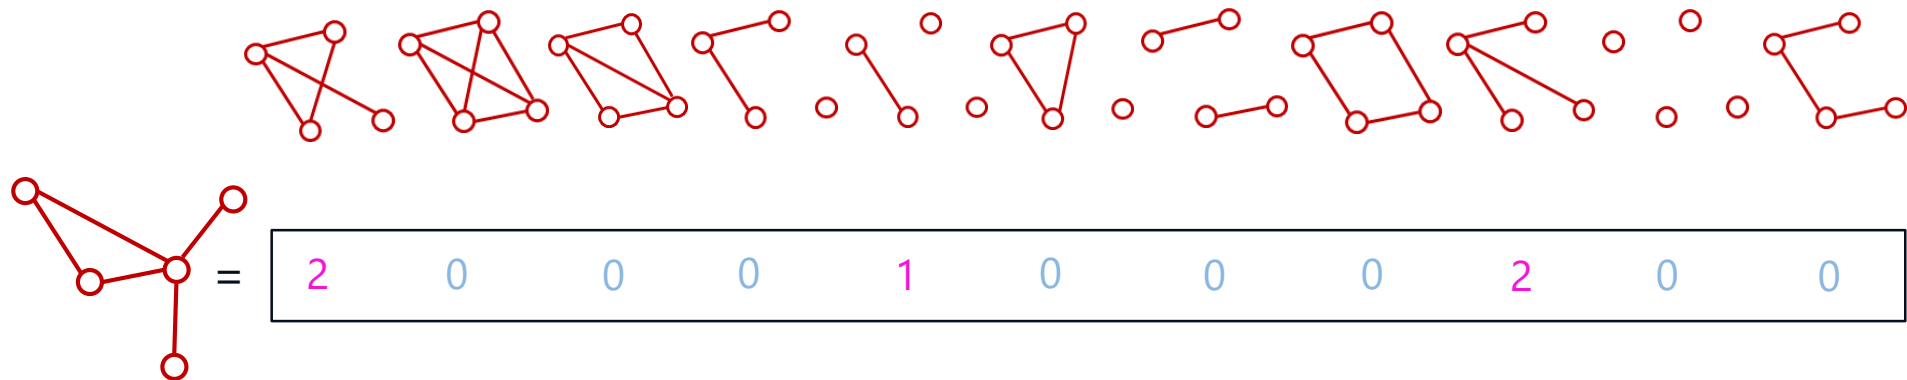
- $k_N(G, G') = \langle \phi_N(G), \phi_N(G') \rangle_{\mathcal{H}}$ where $\phi_N(G) := \sum_{v \in V} \phi_{node}(v)$
- When $\phi_{node}(v) = e_{l(v)}$ (k_{node} is a Dirac kernel on node labels), ϕ_N is an unnormalized histogram that counts occurrences of node labels



A	B	C
2	3	1



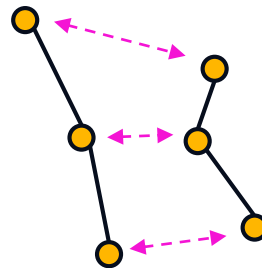
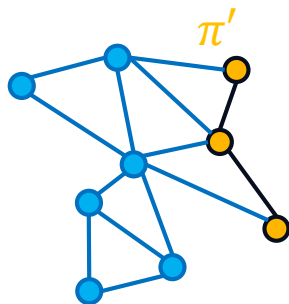
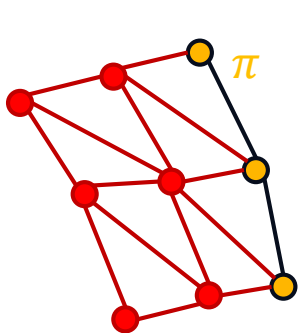
Graphlet kernel



- Set of k -graphlets of size N_k , $k \geq 3$
- k -spectrum of G : vector $\phi_{GL}(G)$ of the frequencies of all graphlets in G
- $k_{GL}(G, G') := \phi_{GL}(G)\phi_{GL}(G')^T$
- Issue: does not take into account labels or attributes

Graph Hopper

[Feragen et al., 2013]

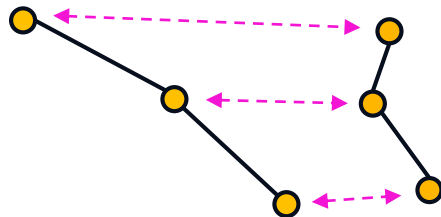
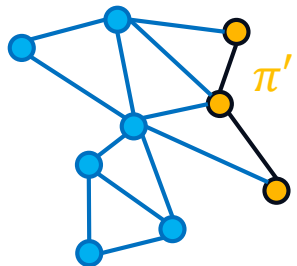
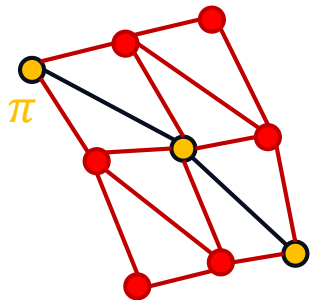


$$k(G, G') := \sum_{\pi \in \mathcal{P}, \pi' \in \mathcal{P}'} k_p(\pi, \pi') \quad \text{with } k_p(\pi, \pi') := \begin{cases} \sum_{j=1}^{|\pi|} \text{RBF}(\pi_j, \pi'_j) & \text{if } |\pi| = |\pi'| \\ 0 & \text{otherwise} \end{cases}$$

- \mathcal{P} : set of all shortest paths in G , $|\pi|$: discrete length of the path $\pi = (\pi_1, \dots, \pi_{|\pi|})$
- Complexity: $O(n^2(|E| + \log n))$

Graph Hopper

[Feragen et al., 2013]

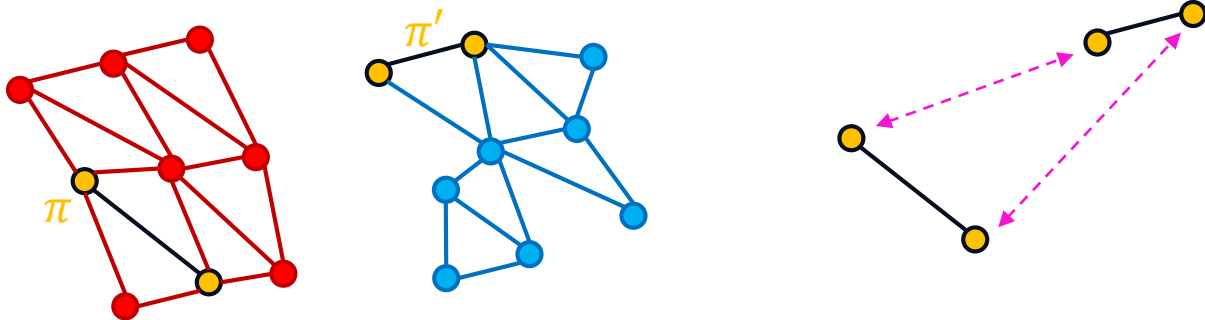


$$k(G, G') := \sum_{\pi \in \mathcal{P}, \pi' \in \mathcal{P}'} k_p(\pi, \pi') \quad \text{with } k_p(\pi, \pi') := \begin{cases} \sum_{j=1}^{|\pi|} \text{RBF}(\pi_j, \pi'_j) & \text{if } |\pi| = |\pi'| \\ 0 & \text{otherwise} \end{cases}$$

- \mathcal{P} : set of all shortest paths in G ,
 - Complexity: $O(n^2(|E| + \log n))$
- $|\pi|$: discrete length of the path $\pi = (\pi_1, \dots, \pi_{|\pi|})$

Graph Hopper

[Feragen et al., 2013]



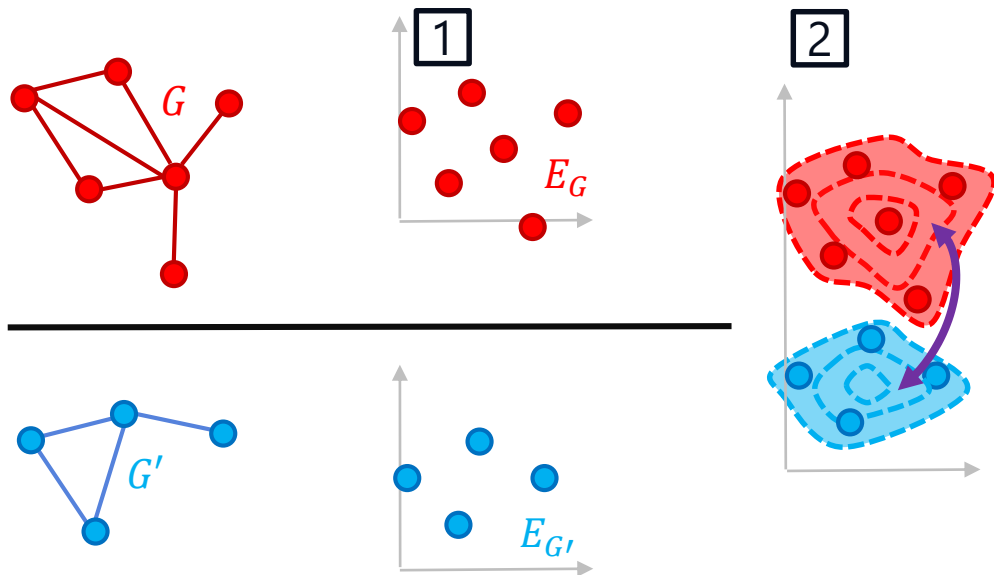
$$k(G, G') := \sum_{\pi \in \mathcal{P}, \pi' \in \mathcal{P}'} k_p(\pi, \pi') \quad \text{with } k_p(\pi, \pi') := \begin{cases} \sum_{j=1}^{|\pi|} \text{RBF}(\pi_j, \pi'_j) & \text{if } |\pi| = |\pi'| \\ 0 & \text{otherwise} \end{cases}$$

- \mathcal{P} : set of all shortest paths in G ,
 - Complexity: $O(n^2(|E| + \log n))$
- $|\pi|$: discrete length of the path $\pi = (\pi_1, \dots, \pi_{|\pi|})$

Sliced Wasserstein Weisfeiler Lehman (SWWL)



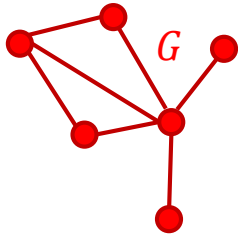
Node embeddings + Optimal transport approaches



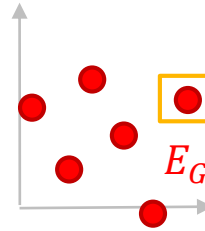
Wasserstein Weisfeiler-Lehman Graph kernel (step 1)

[Togninalli et al., 2019]

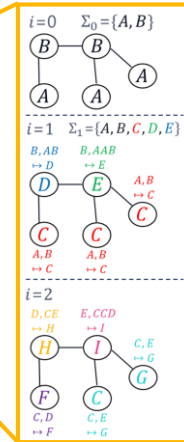
1



ϕ
Node embedding



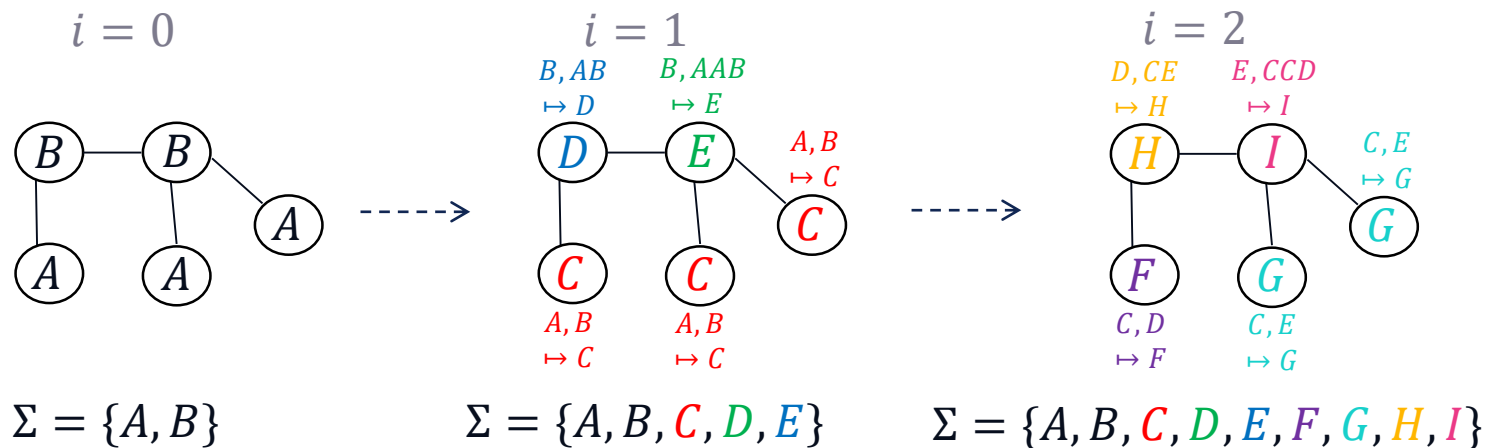
E_G



Weisfeiler-Lehman embeddings

Figure From [Kriege et al., 2020]

- WL relabeling (discrete case)



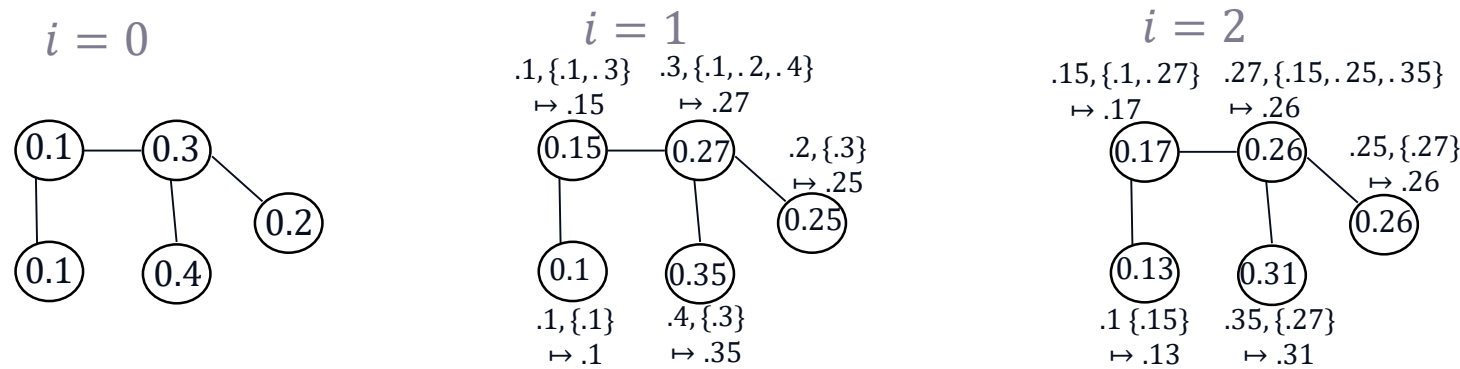
$$l^{(i+1)}(v) = \text{Hash}(l^i(v), \{l^i(u), u \in \mathcal{N}(v)\})$$

$$X_G^{(i)} = [l^{(i)}(v), v \in V_G] \quad X_G = \text{Concatenate}(X_G^{(0)}, \dots, X_G^{(H)})$$

Continuous Weisfeiler-Lehman embeddings

[Togninalli et al., 2019]

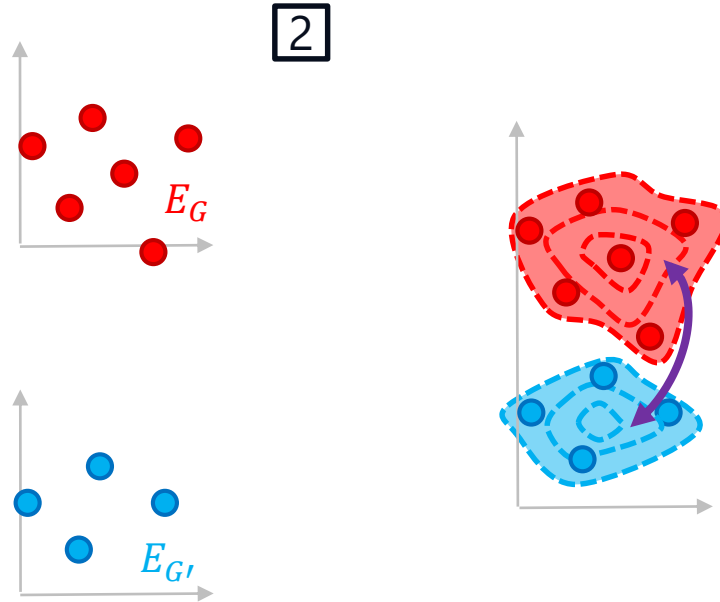
- WL relabeling (continuous case)



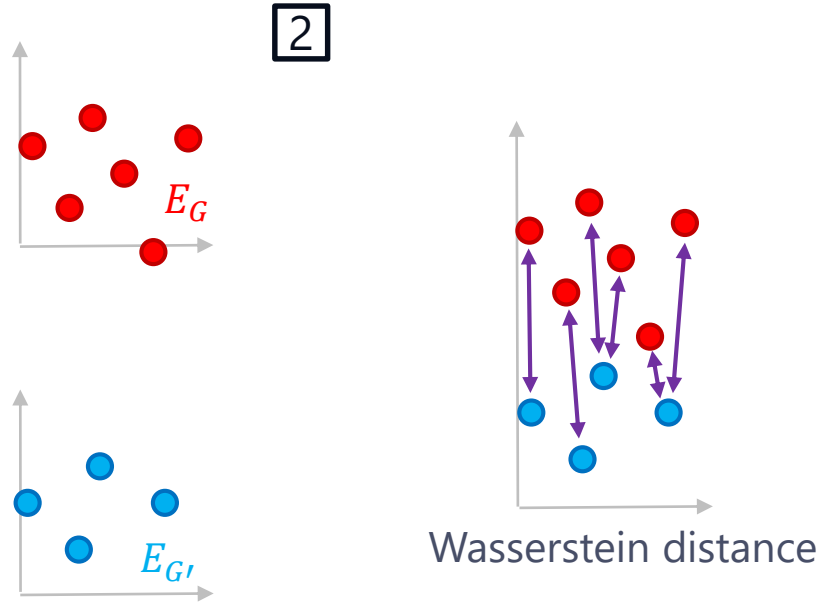
$$a^{(i+1)}(v) = \frac{1}{2} (a^{(i)}(v) + \frac{1}{\deg(v)} \sum_{u \in \mathcal{N}(v)} w(v, u) a^{(i)}(u))$$

$$X_G^{(i)} = [a^{(i)}(v), v \in V_G] \quad X_G = \text{Concatenate}(X_G^{(0)}, \dots, X_G^{(H)})$$

Wasserstein Weisfeiler-Lehman graph kernel (step 2)



Wasserstein Weisfeiler-Lehman graph kernel (step 2)



Wasserstein distance

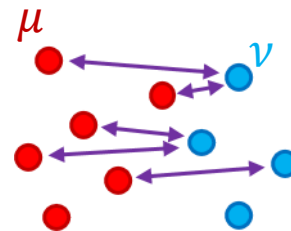
- $\forall r \in [1, +\infty)$, $\mathcal{P}_r(\mathbb{R}^S)$: probability measures on \mathbb{R}^S with finite moments of order r .

$$\forall \mu, \nu \in \mathcal{P}_r(\mathbb{R}^S), \mathcal{W}_r^r(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^S \times \mathbb{R}^S} \|x - y\|^r d\pi(x, y)$$

where:

- $\|\cdot\|$ denotes the Euclidean norm,
- $\Pi(\mu, \nu)$ the set of probability measures on $\mathbb{R}^S \times \mathbb{R}^S$ whose marginals w.r.t. the 1st/2nd variable are resp. μ and ν

- Discrete case: $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ $\nu = \frac{1}{n'} \sum_{i=1}^{n'} \delta_{y_i}$



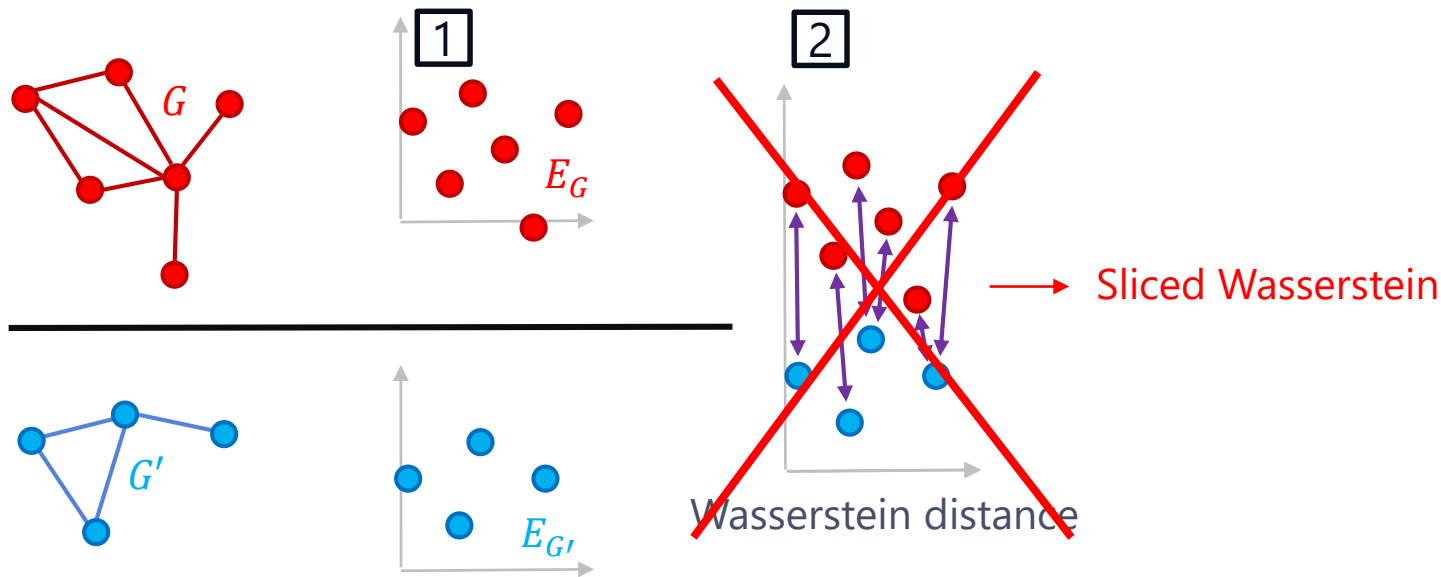
Wasserstein distance: issues

- ✗ Impossible to build a positive definite kernel (*in dimension ≥ 2 *) [Peyré, Cuturi, 2019]
- ✗ Computationally expensive : $O(n^3 \log(n))$
- Use case: 1000 graphs with 30 000 vertices
→ **400 days** to build the Gram matrix...



Sliced Wasserstein Weisfeiler Lehman graph kernel

[Us]



Idea: replace Wasserstein by **sliced Wasserstein** !

→ ✓ $O(n \log(n))$ and ✓ **positive definite** substitution kernels

[Meunier et al., 2022]

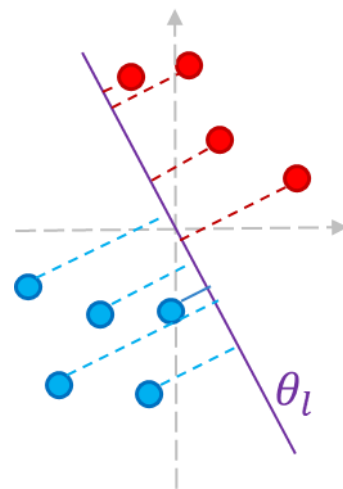
Sliced Wasserstein distance

- The **sliced Wasserstein** distance is defined as:

$$\mathcal{SW}_r^r(\mu, \nu) = \int_{\mathbb{S}^{S-1}} \mathcal{W}_r^r(\theta_{\#}^* \mu, \theta_{\#}^* \nu) d\sigma(\theta)$$

where

- \mathbb{S}^d : d -dimensional unit sphere, σ : uniform distribution on \mathbb{S}^d
- $\theta_{\#}^* \mu$: push-forward measure of $\mu \in \mathcal{P}_r(\mathbb{R}^S)$ by $\theta^* \left(\begin{array}{l} \mathbb{R}^S \rightarrow \mathbb{R} \\ x \mapsto \langle \theta, x \rangle \end{array} \right)$



$$\mathcal{W}_r^r(\mu, \nu) = \int_0^1 |F^{-1}(\mu) - F^{-1}(\nu)|^r dt$$

Quantile
function

1-d Wasserstein distances between

μ and ν

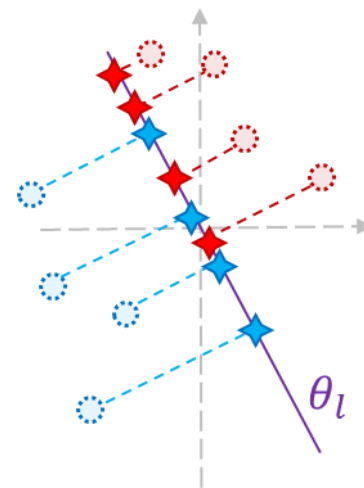
Sliced Wasserstein distance

- The **sliced Wasserstein** distance is defined as:

$$\mathcal{SW}_r^r(\mu, \nu) = \int_{\mathbb{S}^{S-1}} \mathcal{W}_r^r(\theta_{\#}\mu, \theta_{\#}\nu) d\sigma(\theta)$$

where

- \mathbb{S}^d : d -dimensional unit sphere, σ : uniform distribution on \mathbb{S}^d
- $\theta_{\#}\mu$: push-forward measure of $\mu \in \mathcal{P}_r(\mathbb{R}^S)$ by $\theta^* \left(\begin{array}{c} \mathbb{R}^S \rightarrow \mathbb{R} \\ x \mapsto \langle \theta, x \rangle \end{array} \right)$



$$\mathcal{W}_r^r(\mu, \nu) = \frac{1}{n} \sum_{i=1}^n |x_{(i)} - y_{(i)}|^r$$

1-d Wasserstein distances between

$$\mu = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \quad \text{and} \quad \nu = \frac{1}{n} \sum_{i=1}^n \delta_{y_i}$$

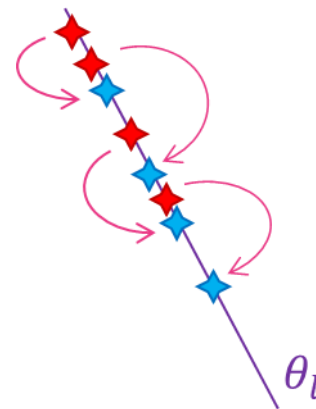
Sliced Wasserstein distance

- The **sliced Wasserstein** distance is defined as:

$$\mathcal{SW}_r^r(\mu, \nu) = \int_{\mathbb{S}^{S-1}} \mathcal{W}_r^r(\theta_{\#}^* \mu, \theta_{\#}^* \nu) d\sigma(\theta)$$

where

- \mathbb{S}^d : d -dimensional unit sphere, σ : uniform distribution on \mathbb{S}^d
- $\theta_{\#}^* \mu$: push-forward measure of $\mu \in \mathcal{P}_r(\mathbb{R}^S)$ by $\theta^* \left(\begin{array}{c} \mathbb{R}^S \rightarrow \mathbb{R} \\ x \mapsto \langle \theta, x \rangle \end{array} \right)$



$$\mathcal{W}_r^r(\mu, \nu) = \frac{1}{n} \sum_{i=1}^n |x_{(i)} - y_{(i)}|^r$$

1-d Wasserstein distances between

$$\mu = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \quad \text{and} \quad \nu = \frac{1}{n} \sum_{i=1}^n \delta_{y_i}$$

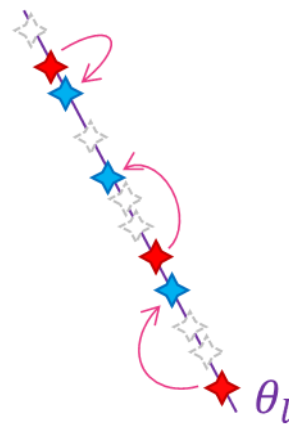
Sliced Wasserstein distance

- The **sliced Wasserstein** distance is defined as:

$$\mathcal{S}\mathcal{W}_r^r(\mu, \nu) = \int_{\mathbb{S}^{S-1}} \mathcal{W}_r^r(\theta_{\#}^* \mu, \theta_{\#}^* \nu) d\sigma(\theta)$$

where

- \mathbb{S}^d : d -dimensional unit sphere, σ : uniform distribution on \mathbb{S}^d
- $\theta_{\#}^* \mu$: push-forward measure of $\mu \in \mathcal{P}_r(\mathbb{R}^S)$ by $\theta^* \left(\begin{array}{c} \mathbb{R}^S \rightarrow \mathbb{R} \\ x \mapsto \langle \theta, x \rangle \end{array} \right)$



$$\widehat{\mathcal{W}}_r^r(\mu, \nu) = \frac{1}{Q} \sum_{q=1}^Q |x_{(q)} - y_{(q)}|^r$$

(Approximation with $Q \ll \max(n, n')$
quantiles)

1-d Wasserstein distances between

$$\mu = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \quad \text{and} \quad \nu = \frac{1}{n'} \sum_{i=1}^{n'} \delta_{y_i}$$

Sliced Wasserstein distance

- The (estimated) **sliced Wasserstein** distance is defined as:

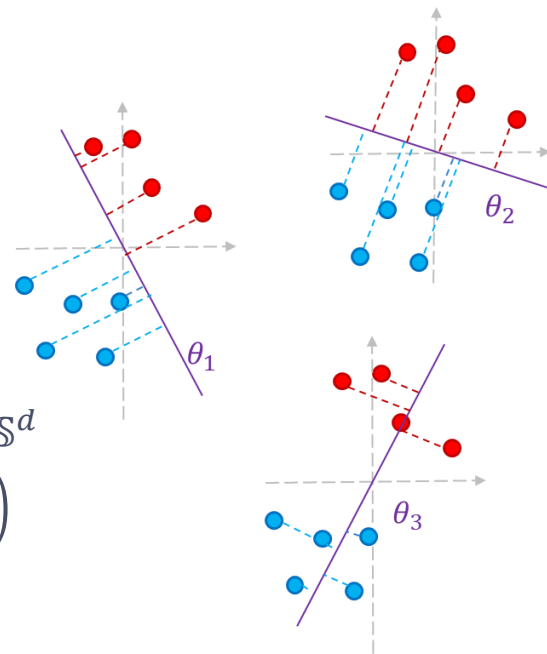
$$\widehat{\mathcal{S}\mathcal{W}}_r^r(\mu, \nu) = \frac{1}{P} \sum_{p=1}^P \widehat{\mathcal{W}}_r^r((\theta_p^*)_{\#}\mu, (\theta_p^*)_{\#}\nu)$$

where

- \mathbb{S}^d : d -dimensional unit sphere, σ : uniform distribution on \mathbb{S}^d
- $\theta_{\#}\mu$: push-forward measure of $\mu \in \mathcal{P}_r(\mathbb{R}^s)$ by $\theta^* \left(\begin{array}{c} \mathbb{R}^s \rightarrow \mathbb{R} \\ x \mapsto \langle \theta, x \rangle \end{array} \right)$

$$\widehat{\mathcal{W}}_r^r(\mu, \nu) = \frac{1}{Q} \sum_{q=1}^Q |x_{(q)} - y_{(q)}|^r$$

(Approximation with $Q \ll \max(n, n')$ quantiles)

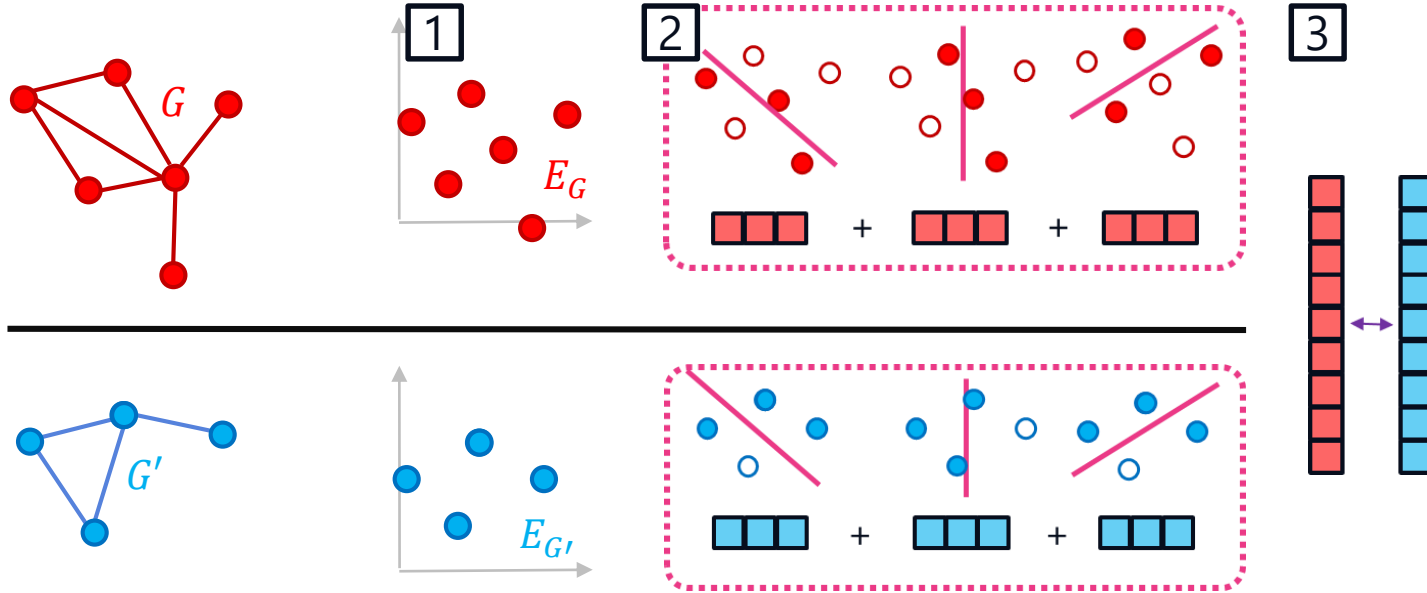


1-d Wasserstein distances between

$$\mu = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \quad \text{and} \quad \nu = \frac{1}{n'} \sum_{i=1}^{n'} \delta_{y_i}$$

Sliced Wasserstein Weisfeiler Lehman (SWWL)

[Us]



Sliced Wasserstein Weisfeiler Lehman (SWWL)

[Us]

- $\phi: G \mapsto X_G \in \mathbb{R}^{|V_G| \times d(H+1)}$: WL embeddings after H iterations
- $k_{SWWL}(G, G') = e^{-\lambda \widehat{\mathcal{SW}}_2^2(\phi(G), \phi(G'))}$ (* considering by abuse $\phi(G), \phi(G')$ as empirical measures *)

with

$$\widehat{\mathcal{SW}}_2^2(\mu, \nu) = \frac{1}{PQ} \sum_{p=1}^P \sum_{q=1}^Q |u_q^{\theta_p} - u'_q{}^{\theta_p}|^2 = \|E_{\phi(G)} - E_{\phi(G')}\|_2^2$$

→ Precomputed embeddings $E_{\phi(G)}, E_{\phi(G')} \in \mathbb{R}^{PQ}$ where $u_q^{\theta_p} = \langle \theta_p, \phi(G) \rangle_{(q)}$

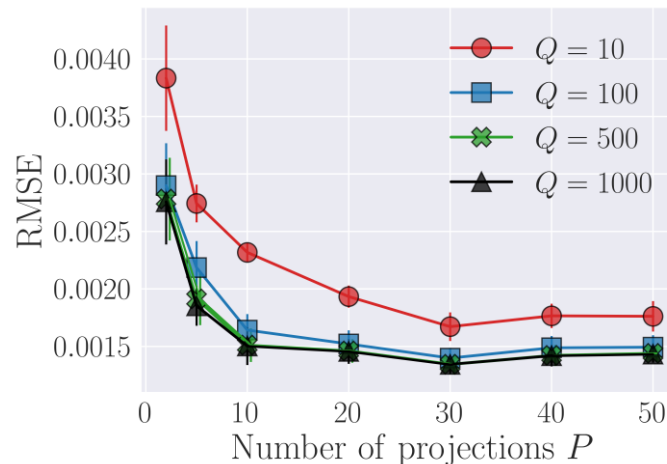
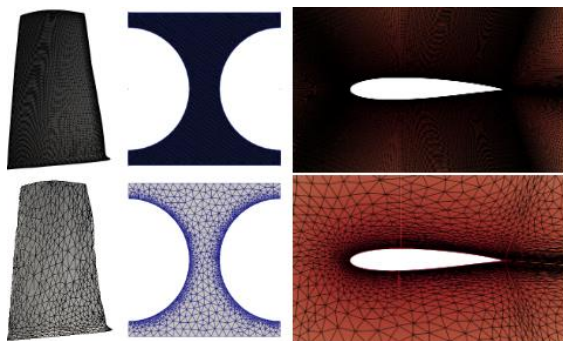
$$E_{\phi(G)} = [u_1^{\theta_1}, \dots, u_Q^{\theta_1}, \dots, u_1^{\theta_P}, \dots, u_Q^{\theta_P}]$$

- Complexity for the Gram matrix (sparse graphs):

$$O(\underbrace{NHn}_{\text{WL iterations}} + \underbrace{NP n \log n}_{\text{Quantiles}} + \underbrace{N^2 PQ}_{\text{Usual RBF kernel}})$$

SWWL: experiments on meshes

[Us]



RMSE (5 exp)

Kernel/Dataset	Rotor37 $\times 10^{-3}$	Rotor37-CM $\times 10^{-3}$	Tensile2d $\times 1$	Tensile2d-CM $\times 1$	AirFRANS $\times 10^{-4}$	AirFRANS-CM $\times 10^{-4}$
SWWL	1.44 ± 0.07	3.49 ± 0.15	0.89 ± 0.01	1.51 ± 0.01	7.56 ± 0.36	9.63 ± 0.54
WWL	-	3.51 ± 0.00	-	6.46 ± 0.00	-	14.4 ± 0.80
PK	-	4.18 ± 0.39	-	6.03 ± 4.58	-	8.94 ± 2.31

Time to build the Gram matrix

Kernel/Dataset	Rotor37	Rotor37-CM	Tensile2d	Tensile2d-CM	AirFRANS	AirFRANS-CM
SWWL	1min + 11s	4s + 11s	11s + 4s	2s + 4s	5min + 7s	15s + 7s
WWL	-	13min (*)	-	6min (*)	-	8h (*)
PK	-	1min	-	2min	-	15min

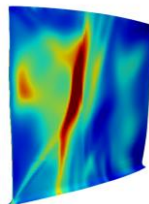
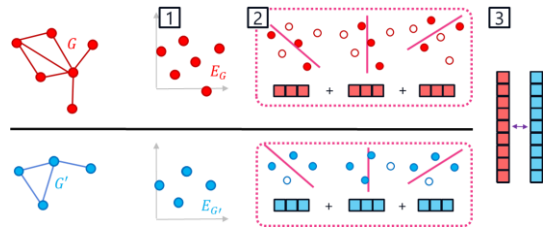
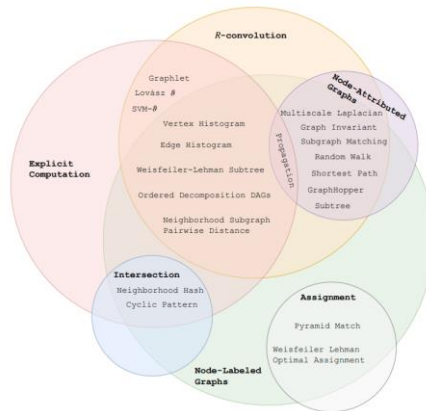
(*) in parallel, using 100 jobs

Conclusion and future work



Conclusion

- Limits of existing graph kernels
 - Many do not handle continuous attributes
 - Many do not scale well to large graphs
 - Many do not guarantee positive definiteness
 - Many are too dependent on the graph structure
- We propose the Sliced Wasserstein Weisfeiler Lehman (SWWL) kernel
 - Positive definite
 - Tractable for large graphs
 - Competitive results for mesh-based Gaussian process regression
- Future work
 - Extension to multiple outputs (e.g. vector fields)



References

▪ Graph kernels, Gaussian processes

- Nikolentzos, G., Siglidis, G., & Vazirgiannis, M. (2021). Graph kernels: A survey. *Journal of Artificial Intelligence Research*, 72, 943-1027.
- Kriege, N. M., Johansson, F. D., & Morris, C. (2020). A survey on graph kernels. *Applied Network Science*, 5(1), 1-42.
- Feragen, A., Kasenburg, N., Petersen, J., de Bruijne, M., & Borgwardt, K. (2013). Scalable kernels for graphs with continuous attributes. *Advances in neural information processing systems*, 26.
- Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning* (Vol. 2, No. 3, p. 4). Cambridge, MA: MIT press.

▪ Optimal transport

- Peyré, G., & Cuturi, M. (2019). Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6), 355-607.
- Togninalli, M., Ghisu, E., Llinares-López, F., Rieck, B., & Borgwardt, K. (2019). Wasserstein weisfeiler-lehman graph kernels. *Advances in neural information processing systems*, 32.
- Meunier, D., Pontil, M., & Ciliberto, C. (2022, June). Distribution Regression with Sliced Wasserstein Kernels. In *International Conference on Machine Learning* (pp. 15501-15523). PMLR.

Acknowledgments

- This work was supported by the French National Research Agency (ANR) through the SAMOURAI project under grant ANR20-CE46-0013.



Other approaches using Optimal Transport



Other approaches

- Many approaches with **GCNNs** and **message passing layers**

→ Continuous WL of torch_geometric

- Other node embedding: $a^{(i+1)}(v) = \sum_{u \in \mathcal{N}(v) \cup \{v\}} \frac{w(v,u)}{\sqrt{\deg(u) \deg(v)}} a^{(i)}(u)$

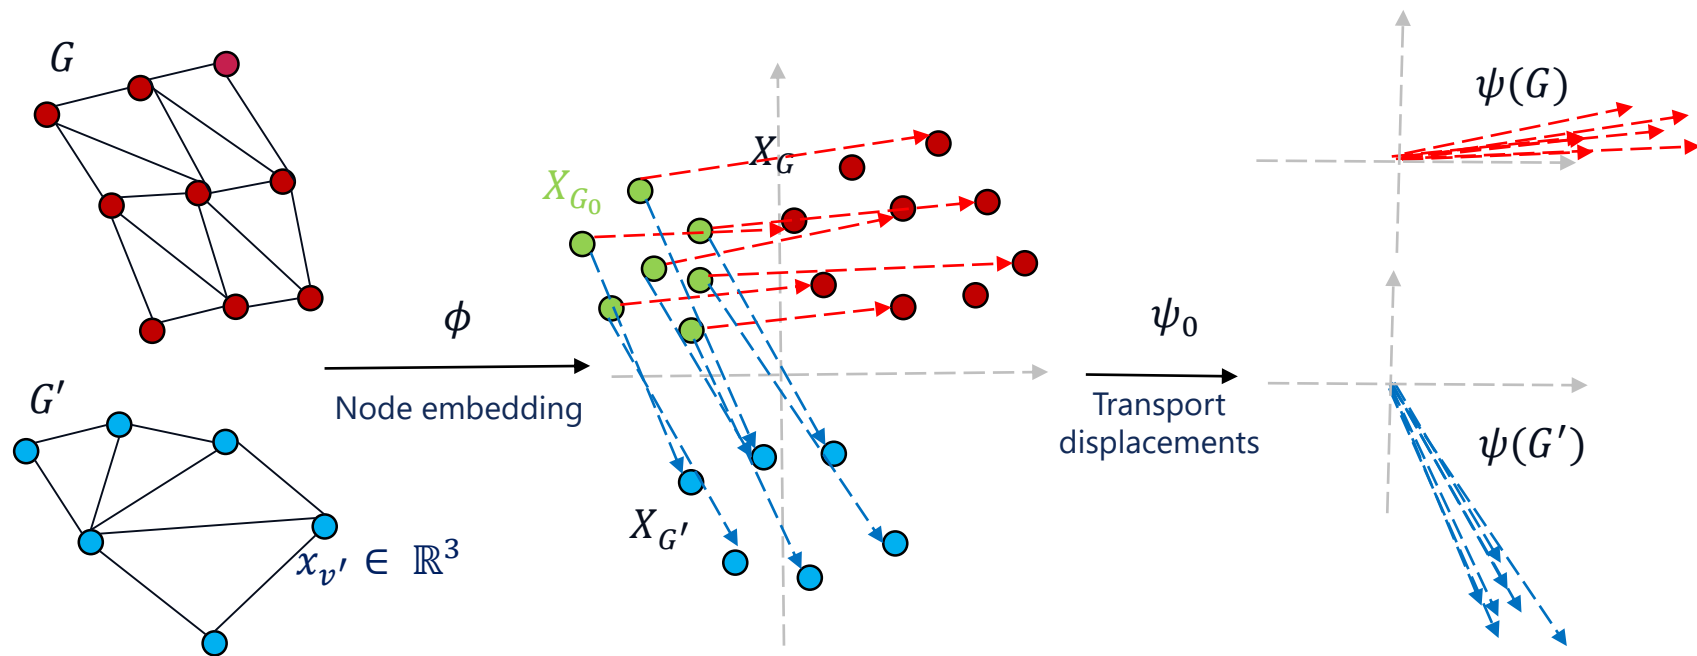
- Wasserstein embeddings with Linear Optimal transport [Kolouri et al., 2020]

- Pooling by Sliced-Wasserstein (PSWE) [Naderializadeh., 2021]

- Template-based GNN with OT [Vincent-Cuaz et al., 2022]

Wasserstein embeddings

[Kolouri et al., 2020]

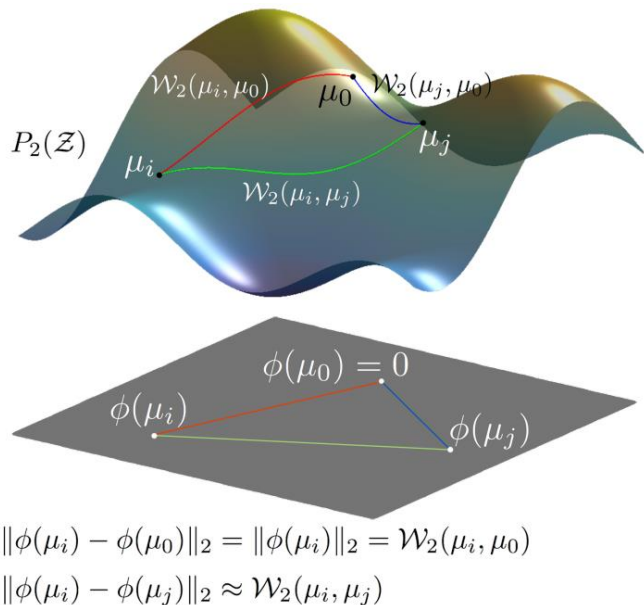


- Linear Wasserstein embedding (Linear Optimal transport LOT Framework)
- Transport displacements from a reference distribution to node embeddings

Wasserstein embeddings

[Kolouri et al., 2020]

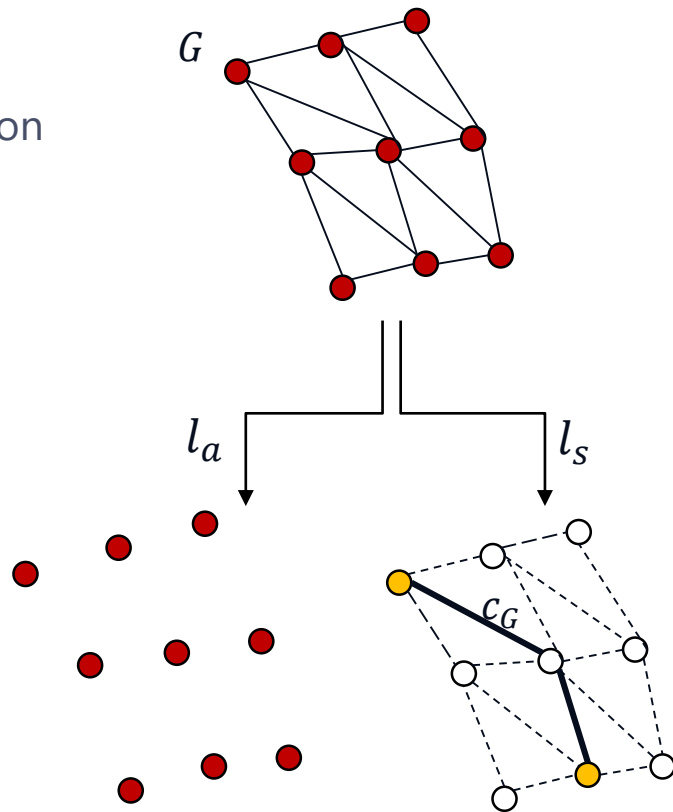
- Given a first node embedding $\phi: G \mapsto X_G \in \mathbb{R}^{|V_G| \times s}$
- $X_0 \in \mathbb{R}^{n_0 \times s}$ reference node embedding
- Linear Wasserstein embedding:
- $\psi_0(X_G) := (u_{G,0} - Id)\sqrt{n_0}$
- where $u_{G,0}$ is the Monge map that pushes X_0 to X_G
- New graph embedding: $\psi(G) := \psi_0(\phi(G)) \in \mathbb{R}^{n_0 \times s}$ of fixed size
- **Only N Monge map calculations needed**
- Choice of the reference embedding? (Not clear)



Fused Gromov-Wasserstein distance

[Vayer et al., 2019]

- $G = (V_G, E_G, l_a, l_s)$ with $l_a: V_G \rightarrow \mathbb{R}^3$ the coordinate function
- $l_s: V_G \rightarrow \Omega_G$ with (Ω_G, c_G) a metric space dependant of G
- $c_G: \Omega_G \times \Omega_G \rightarrow \mathbb{R}_+$ 'similarity' of points **in** G
(structure-dependent)
e.g. : $c_G(l_s(v_1), l_s(v_2)) = d_{PCC}(v_1, v_2 | G)$
- $a_i = l_a(v_i), s_i = l_s(v_i)$: attributes/structure of point i
- $\mu_G = \sum_{i=1}^{n_G} \frac{1}{n_G} \delta_{(a_i, s_i)}$: measure of G
- $C_G = [c_G(s_i, s_j)]_{1 \leq i, j \leq n_G}, C_{G'} = [c_{G'}(s'_i, s'_j)]_{1 \leq i, j \leq n_{G'}}$



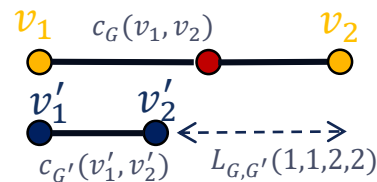
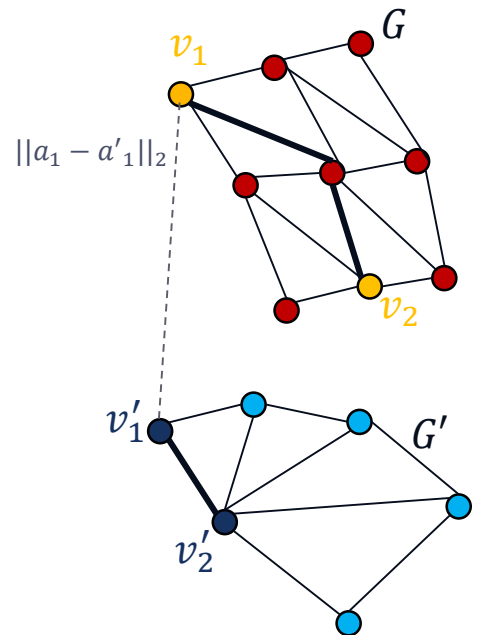
Fused Gromov-Wasserstein distance

[Vayer et al., 2019]

- $L_{G,G'} = |C_G[i,k] - C_{G'}[j,l]|_{i,j,k,l \in \mathbb{R}^{n_G \times n_{G'} \times n_G \times n_{G'}}$
- $M_{G,G'} = [\|a_i - a'_j\|_2]_{1 \leq i \leq n_G; 1 \leq j \leq n_{G'}} \in \mathbb{R}^{n_G \times n_{G'}}$
- $FGW_{q,\alpha}(\mu_G, \mu_{G'}) = \min_{\pi \in \Pi} \left(\alpha M_{G,G'}^q + (1 - \alpha) L_{G,G'}^q \otimes \pi, \pi \right)$

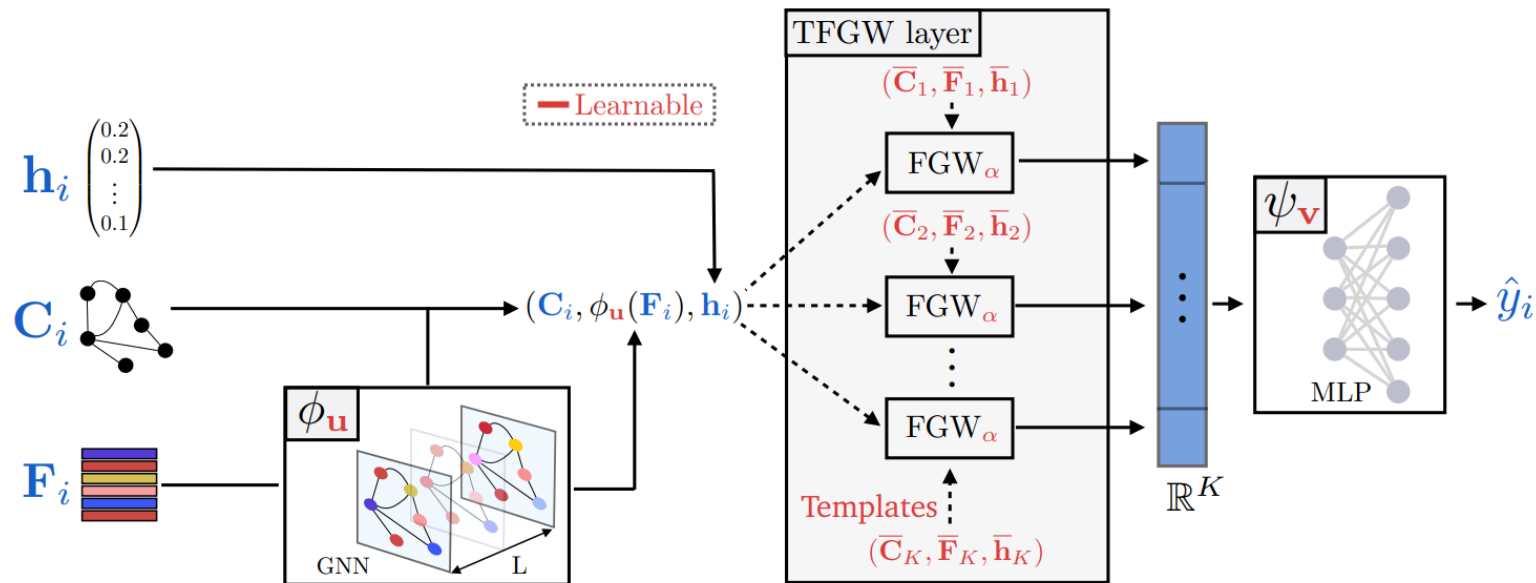
Wasserstein Gromov-Wasserstein

- **Issue:** $k(G, G') = e^{-\gamma FGW_{q,\alpha}(\mu_G, \mu_{G'})}$ is **not positive definite**



Template based GNN with OT

[Vincent-Cuaz et al., 2022]



Graph Convolutional Gaussian Processes

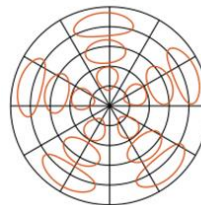
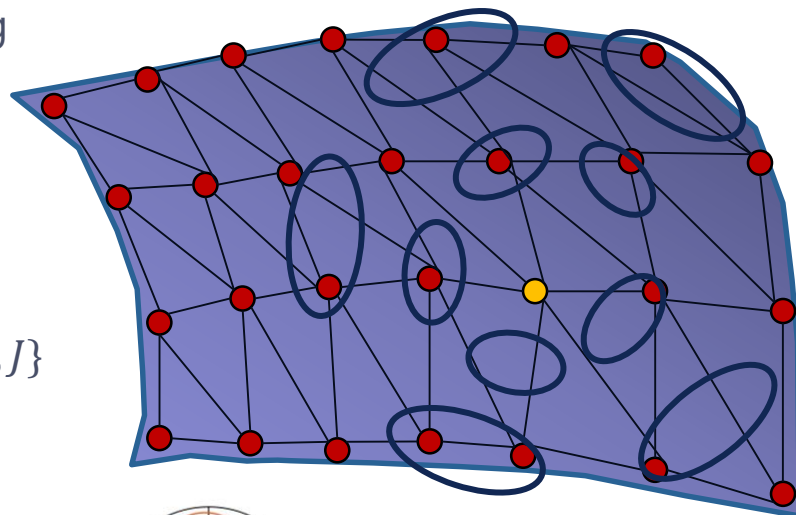
[Walker et al., 2019]

- Graph Convolutional Gaussian Processes
- Local patches around vertices are defined using Spatial-domain charting

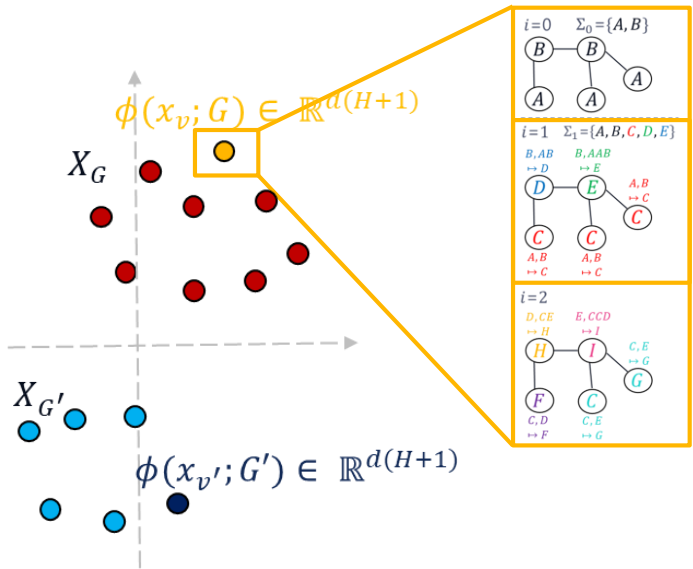
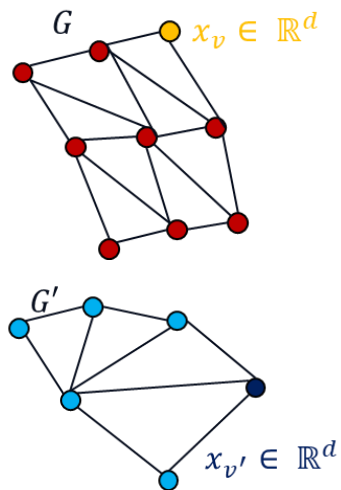
- J : number of bins
- Convolution operator on the graph signal

$\psi: V \rightarrow \mathbb{R}^3 :$

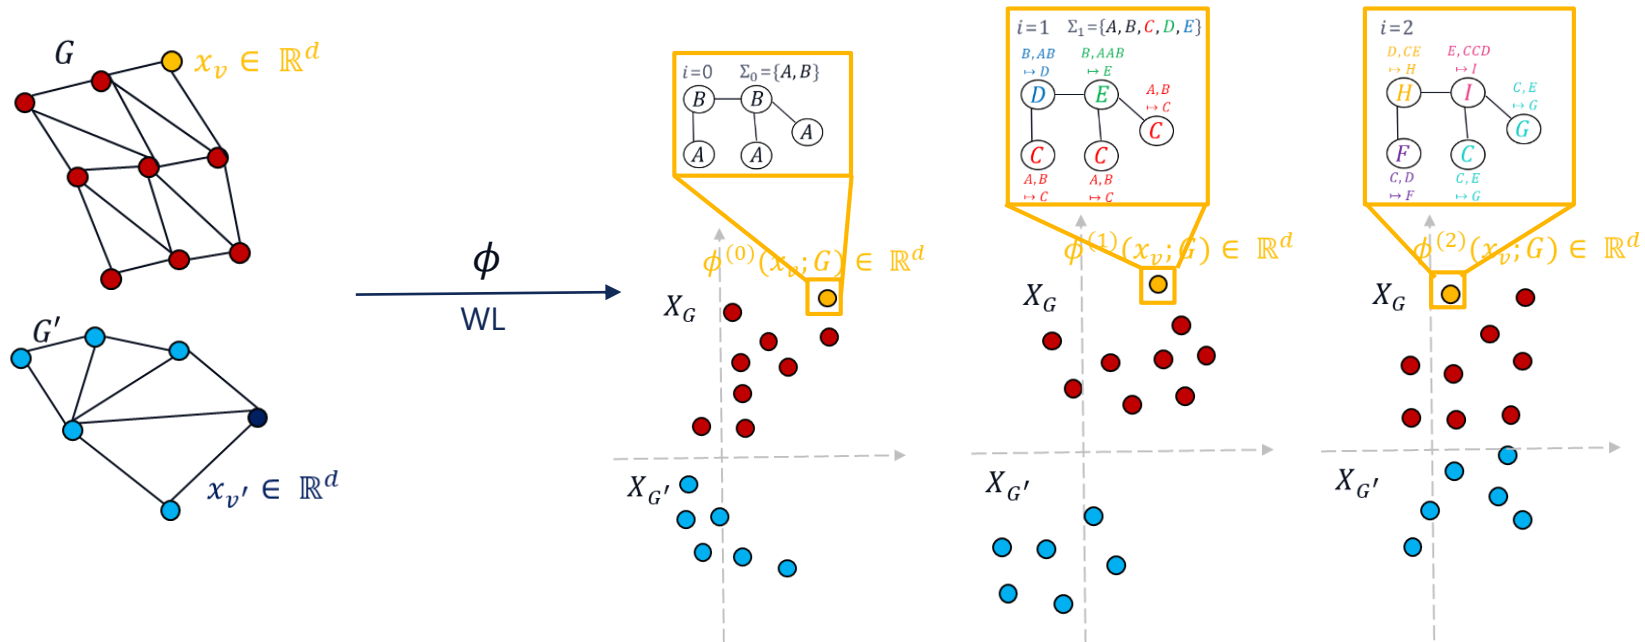
- $D_j(v) \psi = \sum_{u \in V} \psi(u) u_j(u, v) \quad \forall j \in \{1, \dots, J\}$
- u_j : geodesic polar weighting function e.g.



Future work : Anisotropic SWWL?



Future work : Anisotropic SWWL?

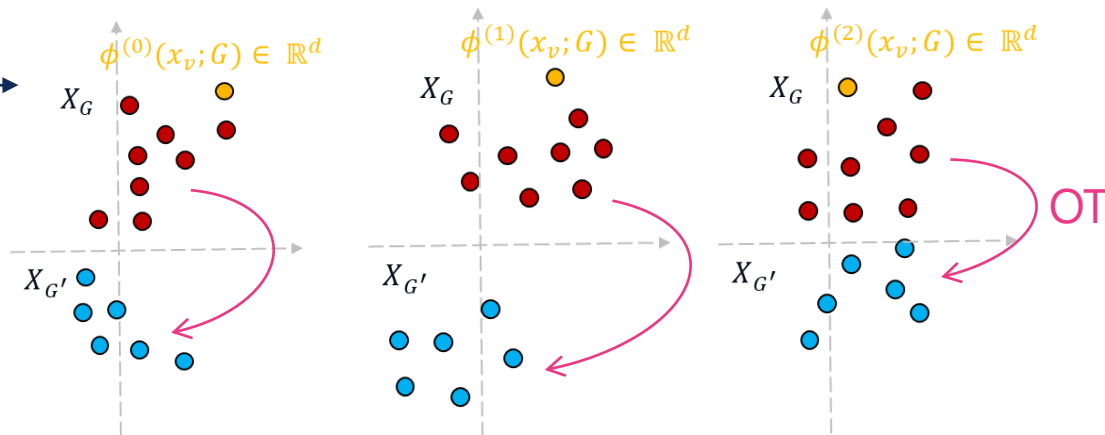
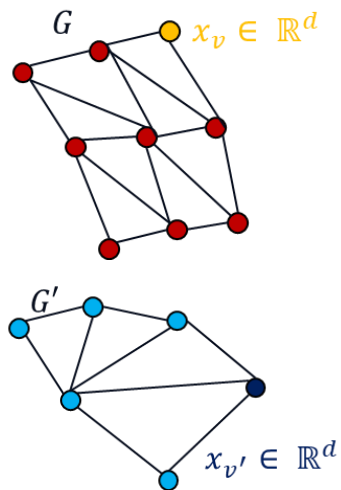


Future work : Anisotropic SWWL?

Anisotropic SWWL:

$$\phi^{(i)} : G \mapsto X_G^{(i)} \in \mathbb{R}^{|V_G| \times d} \quad (i\text{-th iteration of WL})$$

$$k_{ASWWL}(G, G') = e^{-\sum_{i=0}^H \lambda_i \widehat{\text{SW}}_2^2(\phi^{(i)}(G), \phi^{(i)}(G'))}$$



**POWERED
BY TRUST**
