
Learning signals defined on graphs with optimal transport and Gaussian process regression

Raphaël Carpintero Perez^{1,2} Sébastien Da Veiga³ Josselin Garnier² Brian Staber¹

¹ Safran Tech, Digital Sciences & Technologies, 78114 Magny-Les-Hameaux, France

² Centre de Mathématiques Appliquées, Ecole polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

³ Univ Rennes, Ensai, CNRS, CREST - UMR 9194, F-35000 Rennes, France

Abstract

In computational physics, machine learning has now emerged as a powerful complementary tool to explore efficiently candidate designs in engineering studies. Outputs in such supervised problems are signals defined on meshes, and a natural question is the extension of general scalar output regression models to such complex outputs. Changes between input geometries in terms of both size and adjacency structure in particular make this transition non-trivial. In this work, we propose an innovative strategy for Gaussian process regression where inputs are large and sparse graphs with continuous node attributes and outputs are signals defined on the nodes of the associated inputs. The methodology relies on the combination of regularized optimal transport, dimension reduction techniques, and the use of Gaussian processes indexed by graphs. In addition to enabling signal prediction, the main point of our proposal is to come with confidence intervals on node values, which is crucial for uncertainty quantification and active learning. Numerical experiments highlight the efficiency of the method to solve real problems in fluid dynamics and solid mechanics.

1 INTRODUCTION

Many problems in computational physics rely on solving partial differential equations on a domain with given geometry using the finite element method (FEM). Although much appreciated, FEM often involves meshes with highly refined discretization, which quickly becomes computationally intensive even with parallel computing. In particular, exploring how changes in the geometry impact some key quantities

of interest computed by FEM is an everyday task in engineering for design studies. Due to the associated computational cost, machine learning (ML) is a natural candidate to accelerate such design exploration: starting from an initial database of FEM simulations, a supervised model is trained to predict the FEM outputs from its inputs and is ultimately used as a proxy to evaluate new geometries with a negligible cost. But in this context, the supervised learning task actually involves inputs given as meshes, which can be modeled as graphs with continuous node attributes, different numbers of nodes and edges. In addition, the outputs can be scalar values but also physical quantities of interest defined on each node of the input graph, which we refer to as *signals defined on graphs* or *fields*.

These two specificities give rise to a supervised learning problem which is not in a standard form and thus calls for dedicated algorithms. A prominent example are graph neural networks (GNNs), which have shown remarkable success for this type of application by iteratively aggregating and transforming information from node neighbors. However they still have inherent limitations when the dataset is small and when it comes to predicting output values with associated uncertainties. ML models offering predictive uncertainties play a key role in a wide range of critical industrial applications as they can certify the quality of results (Jaber et al., 2024), can assist sequential design of experiments (Begoli et al., 2019; Vishwakarma et al., 2021) or can be plugged into Bayesian optimization workflows (Jones et al., 1998). Gaussian process (GP) regression is quite a popular approach in this small data setting with required uncertainties, but unfortunately suffers from shortcomings for FEM applications: even if the use of graph kernels can handle inputs defined as graphs, GPs do not generalize easily to complex outputs such as signals on graphs with varying and large number of nodes.

In this article, we introduce the Transported Output Signal GP regression (TOS-GP) that makes it pos-

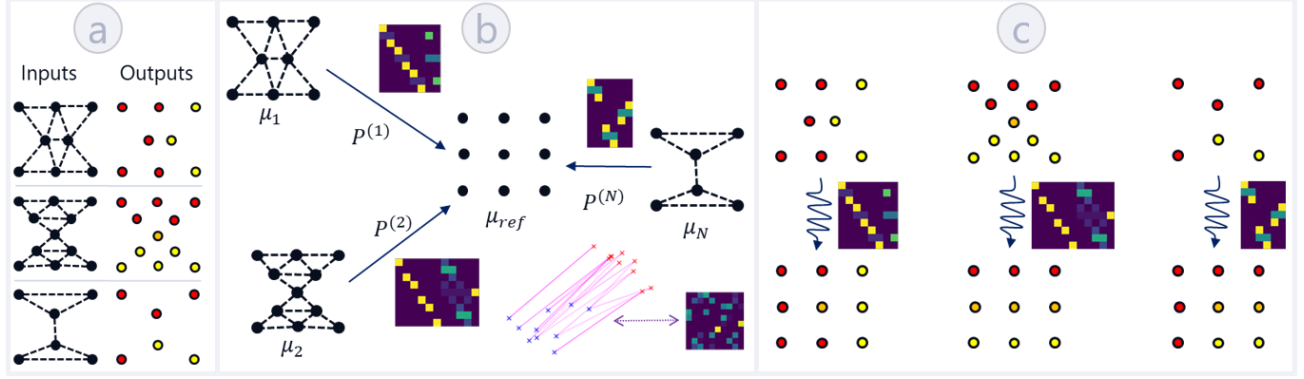


Figure 1: Summary of our approach. a) Inputs = Graphs. Outputs = Fields defined on the nodes. b) Step 1: obtaining transport plans to a reference measure. c) Step 2: Transferring signals to the reference measure.

sible to extend scalar Gaussian process regression to output signals thanks to a regressor-agnostic transformation of the outputs combined with dimension reduction. In a preliminary step, each input graph is pre-processed and encoded as an empirical measure supported on the nodes to which we assign their continuous attributes and those of their neighbors. We then select a reference measure and compute optimal transport plans between all input empirical measures and this reference: this step is crucial since it yields a transformation allowing to transport any information from any graph with any number of nodes to a common space, the reference measure support. In particular, each output signal is transported with the plan corresponding to its input graph: intuitively, this means that if nodes on two graphs match, then the signals defined on them should match too, see Figure 1 for an illustration. The supervised learning problem now writes in an amenable form, although possibly with a large number of new outputs equal to the size of the reference measure. The last step thus consists in applying a dimension reduction technique to finally learn a few independent scalar-valued GPs, taking graphs as inputs, for which we use the Sliced Wasserstein Weisfeiler-Lehman (SWWL) graph kernel Carpi ntero Perez et al. (2024).

The article is organized as follows. Section 2 first introduces the problem setting and recalls GP regression. Related work are discussed in Section 3. The TOS-GP methodology is described in Section 4. Numerical experiments are then performed in Section 5 and Section 6 with a focus on real datasets coming from FEM simulations.

2 PRELIMINARIES

We consider the task of learning a function $f : \mathcal{X} \rightarrow \mathcal{Y}$. \mathcal{X} denotes a space of undirected graphs with continu-

ous node attributes. Each element in \mathcal{X} can be written as $X = (V, E, w, \mathbf{F})$ where V is the set of $|V|$ nodes, E is a set of paired nodes, whose elements are called edges, and the function $w : E \rightarrow \mathbb{R}$ assigns the edge weights. The d -dimensional attributes of the nodes are gathered in the $|V| \times d$ matrix $\mathbf{F} = (\mathbf{F}_u)_{u \in V}$. The output space \mathcal{Y} is a set of signals defined on graphs of \mathcal{X} , that is to say:

$$\mathcal{Y} = \bigcup_{X=(V,E,w,\mathbf{F}) \in \mathcal{X}} \{Y : V \rightarrow \mathbb{R}\}.$$
 By abuse of notation, the signal can be denoted as the vector $\mathbf{Y} = (Y(1), \dots, Y(|V|))^T$. Remark that the outputs are permutation equivariant: for each permutation of the input nodes $(u_1, \dots, u_{|V|})$, the output signal needs to be permuted equivalently.

We assume that we are given a dataset \mathcal{D} consisting of N observations $\mathcal{D} = \{(X^{(i)}, \mathbf{Y}^{(i)})\}_{i=1}^N$, where the inputs may differ in terms of numbers of nodes and adjacency matrices, *i.e.*, we may have $|V^{(i)}| \neq |V^{(j)}|$ and $E^{(i)} \neq E^{(j)}$ for some $(i, j) \in \{1, \dots, N\}^2$.

For this problem, we focus on GP regression, which is a popular Bayesian approach for supervised learning in engineering (Williams and Rasmussen, 2006; Gramacy, 2020). In the particular case of a single output, we try to learn $g : \mathcal{X} \rightarrow \mathbb{R}$ from noisy training observations $y^{(i)} = g(X^{(i)}) + \epsilon^{(i)}$ for $i \in \{1, \dots, N\}$ of $g : \mathcal{X} \rightarrow \mathbb{R}$ at input locations $\mathbf{X} = (X^{(i)})_{i=1}^N$, where $\epsilon^{(i)} \sim \mathcal{N}(0, \eta^2)$ is an additive i.i.d. Gaussian noise. Let $\mathbf{g}_* := (g(X_*^{(i)}))_{i=1}^{N^*}$ be the values of g at new test locations $\mathbf{X}_* = (X_*^{(i)})_{i=1}^{N^*}$.

A zero-mean GP prior is placed on the function g (adding a constant mean is immediate, and has no impact in practice for our use cases). It follows that the joint distribution of the observed target values and the function values at the test locations writes (see, *e.g.*,

Williams and Rasmussen (2006))

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{g}_* \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K} + \eta^2 \mathbf{I} & \mathbf{K}_*^T \\ \mathbf{K}_* & \mathbf{K}_{**} \end{bmatrix}\right), \quad (1)$$

where \mathbf{K} , \mathbf{K}_{**} , \mathbf{K}_* are the train, test and test/train Gram matrices, respectively. The posterior distribution of \mathbf{g}_* , obtained by conditioning the joint distribution on the observed data, is also Gaussian: $\mathbf{g}_* | \mathbf{X}, \mathbf{y}, \mathbf{X}_* \sim \mathcal{N}(\bar{\mathbf{m}}, \bar{\Sigma})$ with mean and covariance given by

$$\bar{\mathbf{m}} = \mathbf{K}_* (\mathbf{K} + \eta^2 \mathbf{I})^{-1} \mathbf{y}, \quad (2)$$

$$\bar{\Sigma} = \mathbf{K}_{**} - \mathbf{K}_* (\mathbf{K} + \eta^2 \mathbf{I})^{-1} \mathbf{K}_*^T. \quad (3)$$

The mean of the posterior distribution is used as a predictor, and predictive uncertainties can be obtained through the covariance matrix. GP regression requires a positive definite kernel function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. When \mathcal{X} is a space of graphs, we rely on graph kernels.

When the output space \mathcal{Y} consists of signals on graphs, the learning problem possesses several characteristics that require special attention:

- Inputs $X^{(1)}, \dots, X^{(N)}$ can have different sizes, so the outputs $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(N)}$ do not have a fixed size common to all samples
- There is no natural ordering of the output dimensions, leading to the permutation equivariance of the signals
- The output dimension can be very large

3 RELATED WORK

GNNs and morphing. GNNs are indisputably the reference to build predictions on all the nodes of input graphs by relying on the message passing framework introduced by Gilmer et al. (2017) and extended by Battaglia et al. (2018), with recent advances targeting specifically solutions to physical systems, either by incorporating physical knowledge (Li et al., 2022) or by designing efficient GNNs architectures (Pfaff et al., 2020). GNNs are still limited by the need of a large sample size, and require a specific treatment to get access to uncertainties as elaborated in (Gawlikowski et al., 2023). Alternatively, the Mesh Morphing GP approach proposed by Casenave et al. (2024) consists in morphing the meshes to a reference shape, on which the output fields are then interpolated, followed by dimension reduction. This is close in spirit to our proposal, but this method primarily relies on morphing, which is limited to mesh data only, and meshes with the same topology. On the contrary our method is completely generic since we leverage optimal transport instead of morphing.

From single to multi-output prediction. When outputs are vectors in \mathbb{R}^d , it is possible to consider multi-output GPs (Alvarez et al., 2012) based on vector-valued RKHS (Carmeli et al., 2006), with the famous Linear Model of Coregionalization (LMC) of Journel and Huijbregts (1976), Intrinsic Coregionalization Model (ICM) of Goovaerts (1997) and Convolved Gaussian Processes of Alvarez and Lawrence (2011). However these approaches are not really suitable when the output dimension is large. When the output actually is the discretization of a continuous phenomenon, corresponding to infinite dimensional outputs, operator-valued kernels (Kadri et al., 2016; Owhadi, 2021) are a natural extension of vector-valued kernels when the output space writes as a space of functions. Instead, one can also use dimension reduction in the output space to decompose the problem into independent sub-problems while managing the uncertainty quantification (Kontolati et al., 2022). Recent applications show the use of Singular Value Decomposition (Nanty et al., 2017), wavelet transforms (Perrin et al., 2021) and autoencoders (Donnelly et al., 2024). Unfortunately, all methods mentioned so far assume that the outputs are defined on the same domain across the samples. An exception exists, when it is possible to measure a similarity between structured outputs via a positive definite kernel: the output kernel regression framework (Weston et al., 2002, 2005; Kadri et al., 2013; Brouard et al., 2016) performs prediction in the associated output RKHS instead. But ultimately it requires to solve a non-trivial pre-image problem at test time, requiring a minimization over the output space which is not tractable in the applications we consider.

Graph signal processing. We can also think of dimension reduction techniques based on the graph Laplacian like graph signal processing (Shuman et al., 2013), with many recent applications in ML (Dong et al., 2020; Ricaud et al., 2019). Such approaches usually learn functions defined on the nodes of a common graph (an input is a node, and an output is a scalar) either by regularizing with the Laplacian matrix (Venkitaraman et al., 2020) or by using spectral techniques (Zhi et al., 2023). The use of Laplacian eigenvectors however suffers from indeterminacies due to the choices of signs and basis, as described by Lim et al. (2022). In particular, comparing representations of signals defined on graphs of different sizes by projecting them onto their respective eigenvectors is irrelevant since comparing two eigenvectors of different sizes makes little sense.

Optimal transport. Using optimal transport to handle different domain distributions (Peyré et al., 2019; Montesuma et al., 2023) is not at all new, for ex-

ample color or texture transfer that involves transport plans between histograms (Ferradans et al., 2014). Optimal transport has also seen some applications in the domain of computational physics to interpolate between fields thanks to optimal transport between Gaussian models (Iollo and Taddei, 2022) or to register point clouds (Shen et al., 2021). In the context of domain adaptation (Courty et al., 2016), the recent approach of Courty et al. (2017) can handle changes not only in the input train and test distributions, but also in the output train and test domains by learning simultaneously a predictive function and the optimal transport plan between joint distributions. Even if they share similar ideas, these approaches cannot be used directly in our case since the domains are different for each input, and the output signals to be predicted cannot be used to obtain transport plans. To our knowledge, we propose the first approach to reuse a transport plan defined in feature space in order to change the support of output fields.

4 OPTIMALLY TRANSPORTED SIGNALS

Our TOS-GP methodology relies on the entropy-regularized Wasserstein distance to transfer the signals to a reference measure. Given this reference measure, we find the regularized optimal transport plans between input measures and the reference measure. We then transfer the signals to be expressed on the support of the reference measure using the latter transport plans. The transferred signals now have a fixed order and a fixed size, on which dimension reduction techniques can now be applied. We then learn the low-dimensional embeddings using independent GPs. Algorithm 1 and algorithm 2 respectively summarize the training and test stages.

The preliminary step consists in encoding the inputs of our original problem so that we can use optimal transport. Recall that $X^{(1)}, \dots, X^{(N)}$ are N graphs with n_1, \dots, n_N nodes and with respective feature matrices $\mathbf{F}^{(1)}, \dots, \mathbf{F}^{(N)}$. In order to leverage the adjacency structure of the graphs, we incorporate the continuous Weisfeiler-Lehman (WL) embeddings (Togninalli et al., 2019; Carpintero Perez et al., 2024) to the feature vector $\mathbf{F}_u^{(i)}$ of each node u , whose dimension becomes $d \times (H + 1)$ where H is the number of continuous WL iterations. Each input graph $X^{(i)}$ is finally encoded as an input empirical measure $\mu^{(i)} = \frac{1}{n_i} \sum_{u=1}^{n_i} \delta_{\mathbf{F}_u^{(i)}}$ supported by its features.

Regularized Wasserstein distance. We first recall the definition of the Wasserstein distance for arbitrary measures on $\mathbb{R}^s, s \geq 1$ (in our setting, $s =$

$d \times (H + 1)$).

Definition 1 (Wasserstein distance). *Let $s \geq 1$ be an integer, $r \geq 1$ be a real number, and μ, ν be two probability measures on \mathbb{R}^s having finite moments of order r . The r -Wasserstein distance is defined as*

$$W_r(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^s \times \mathbb{R}^s} \|\mathbf{x} - \mathbf{y}\|^r d\pi(\mathbf{x}, \mathbf{y}) \right)^{\frac{1}{r}}, \quad (4)$$

where $\Pi(\mu, \nu)$ is the set of all probability measures on $\mathbb{R}^s \times \mathbb{R}^s$ whose marginals w.r.t. the first and second variables are respectively μ and ν and $\|\cdot\|$ stands for the Euclidean norm on \mathbb{R}^s .

In the following, we consider only discrete measures $\mu = \frac{1}{n} \sum_{u=1}^n \delta_{\mathbf{x}_u}$ and $\nu = \frac{1}{m} \sum_{v=1}^m \delta_{\mathbf{x}'_v}$ supported respectively on the points $\mathbf{x}_1, \dots, \mathbf{x}_n$ and $\mathbf{x}'_1, \dots, \mathbf{x}'_m \in \mathbb{R}^s$ (by abuse of language, we sometimes refer to the size of a measure as the size of its support). We denote by $U(n, m) = \{\mathbf{P} \in \mathbb{R}_+^{n \times m} : \mathbf{P} \mathbf{1}_m = \frac{1}{n} \mathbf{1}_n \text{ and } \mathbf{P}^T \mathbf{1}_n = \frac{1}{m} \mathbf{1}_m\}$ the set of admissible coupling matrices with mass preservation, where $\mathbf{1}_n$ is the column vector composed of ones. We denote by $\mathbf{C}^{\mu, \nu} := (c(\mathbf{x}_u, \mathbf{x}'_v))_{u,v}$ the $n \times m$ cost matrix, where $c : \mathbb{R}^s \times \mathbb{R}^s \rightarrow \mathbb{R}$ is the cost function $(c(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|^2$ for the r -Wasserstein distance with $r = 2$).

The Kantorovich formulation of optimal transport writes as follows:

$$L(\mu, \nu) := \min_{\mathbf{P} \in U(n, m)} \langle \mathbf{C}^{\mu, \nu}, \mathbf{P} \rangle. \quad (5)$$

When $n = m$, then there exists an optimal solution for (5) which is a scaled permutation matrix. To obtain approximate solutions to the original transport problem (5), one can add an entropic regularization term:

$$\begin{aligned} \mathcal{L}_\lambda(\mu, \nu, \mathbf{P}) &:= \langle \mathbf{C}^{\mu, \nu}, \mathbf{P} \rangle - \lambda H(\mathbf{P}), \\ L_\lambda(\mu, \nu) &:= \min_{\mathbf{P} \in U(n, m)} \mathcal{L}_\lambda(\mu, \nu, \mathbf{P}), \end{aligned} \quad (6)$$

where the discrete entropy of a coupling matrix is defined as $H(\mathbf{P}) := -\sum_{u,v} P_{u,v} (\log(P_{u,v}) - 1)$, and $\lambda \geq 0$ is the regularization parameter. When $\lambda > 0$, the objective is a λ -strongly convex function that has a unique minimizer $P_\lambda := \operatorname{argmin}_{\mathbf{P} \in U(n, m)} \mathcal{L}_\lambda(\mu, \nu, \mathbf{P})$. Solving the regularized problem has a two-fold benefit in our case. First, the $O(n^3 \log(n))$ complexity of the evaluation of the Wasserstein distance can be reduced to $O(n^2 \log(n))$ with entropic regularization, and the computation can be accelerated using GPU devices thanks to Sinkhorn iterations (the higher the coefficient, the quicker it will run). On the other hand, obtaining a more diffuse transport plan can be beneficial for our signal transfer application as it can induce a smoothing depending on the regularization parameter λ .

Signal transfer with optimal transport plans.

Given a reference measure μ_{ref} supported on n_{ref} points, and a common regularization parameter λ , we compute the entropy-regularized transport plan $\mathbf{P}_\lambda^{(i)}$ between each input measure $\mu^{(i)}$ and μ_{ref} :

$$\mathbf{P}^{(i)} = \underset{\mathbf{P} \in U(n_i, n_{\text{ref}})}{\operatorname{argmin}} \mathcal{L}_\lambda(\mu^{(i)}, \mu_{\text{ref}}, \mathbf{P}). \quad (7)$$

Such transport plans provide matches between several points of the involved measures, and the proportion of the masses that need to be split between them. A key ingredient of our method is to use the same transport plans in the output space, in order to obtain a representation of the signals with a fixed support given by the reference measure, called the *OT-transferred fields* given by:

$$\mathbf{T}^{(i)} = (n_{\text{ref}} \mathbf{P}^{(i)})^T \mathbf{Y}^{(i)} \in \mathbb{R}^{n_{\text{ref}}}. \quad (8)$$

Crucially, we can transfer back signals defined on the reference measure to their original support by applying the transposed of the transfer operator:

$$\tilde{\mathbf{Y}}^{(i)} = (n_i \mathbf{P}^{(i)}) \mathbf{T}^{(i)} \in \mathbb{R}^{n_i}. \quad (9)$$

Remark that this reconstruction cannot be exact as soon as $n_i \neq n_{\text{ref}}$ or $\lambda > 0$.

Dimension reduction. Since the transferred signals are now expressed on the same reference measure, any dimension reduction technique \mathcal{M}_{DR} can be applied to $\{\mathbf{T}^{(i)}\}_{i=1}^N$ in order to obtain some low-dimensional embeddings $\{\mathbf{C}^{(i)}\}_{i=1}^N = \mathcal{M}_{DR}(\{\mathbf{T}^{(i)}\}_{i=1}^N)$, where $\mathbf{C}^{(i)} = (C_j^{(i)})_{j=1}^Q \in \mathbb{R}^Q$. For instance, using Principal Component Analysis (PCA) (Jackson, 2005), $\mathbf{C}^{(i)}$ are the first Q PCA coefficients of each transferred signal obtained by multiplication with the projection matrix \mathbf{E}^T containing the first Q PCA eigenvectors.

Single-output GP regression. Approximating $f : \mathcal{X} \rightarrow \mathcal{Y}$ can now be broken down to learning Q independent functions: $g^{(j)} : \mathcal{X} \rightarrow \mathbb{R}$ using the respective datasets $\mathcal{D}_j = \{(X^i, C_j^{(i)})\}_{i=1}^N$ for $j = 1, \dots, Q$. Any positive definite graph kernel (Kriege et al., 2020; Nikolentzos et al., 2021) can be used, but here we choose to use the SWWL kernel of Carpintero Perez et al. (2024), which has shown good performance for scalar prediction when inputs are graphs with continuous attributes.

Alternatively, multi-output Gaussian processes like the ICM (Goovaerts, 1997) could be use in place of Gaussian process regression to explicitly incorporate the correlation matrix between PCA coefficients, but it would be limited to very low values of Q .

Algorithm 1 TOS-GP, training phase

Input: Train dataset $\mathcal{D} = \{(X^{(i)}, \mathbf{Y}^{(i)})\}_{i=1}^N$, reference measure μ_{ref} , regularization parameter λ
Output: (Linear) dimension reduction model \mathcal{M}_{DR} , list of Q GP models $\{\mathcal{M}_{gp,j}\}_{j=1}^Q$

- 1: **for** $i = 1, \dots, N$ **do**
- 2: $\mu^{(i)} \leftarrow \frac{1}{n_i} \sum_{u=1}^{n_i} \delta_{\mathbf{F}_u^{(i)}}$
- 3: $\mathbf{P}^{(i)} \leftarrow \underset{\mathbf{P} \in U(n_i, n_{\text{ref}})}{\operatorname{argmin}} \mathcal{L}_\lambda(\mu^{(i)}, \mu_{\text{ref}}, \mathbf{P})$
- 4: $\mathbf{T}^{(i)} \leftarrow (n_{\text{ref}} \mathbf{P}^{(i)})^T \mathbf{Y}^{(i)}$
- 5: $\tilde{\mathbf{Y}}^{(i)} \leftarrow (n_i \mathbf{P}^{(i)}) \mathbf{T}^{(i)}$
- 6: **end for**
- 7: $\{\mathbf{C}^{(i)}\}_{i=1}^N \leftarrow \mathcal{M}_{DR}(\{\mathbf{T}^{(i)}\}_{i=1}^N)$ \triangleright Reduced dimension Q
- 8: **for** $j = 1, \dots, Q$ **do**
- 9: $\mathcal{M}_{gp,j} \leftarrow GP(\{(X^i, C_j^{(i)})\}_{i=1}^N)$
- 10: **end for**

Algorithm 2 TOS-GP, test phase

Input: Test input X_* , reference measure μ_{ref} , regularization parameter λ , (linear) dimension reduction model \mathcal{M}_{DR} , list of Q GP models $\{\mathcal{M}_{gp,j}\}_{j=1}^Q$
Output: Predicted signal $\hat{\mathbf{Y}}_*$, uncertainties $\{(\hat{\mathcal{S}}_*)_{ii}\}_{i=1}^{n_*}$

- 1: **for** $j = 1, \dots, Q$ **do**
- 2: $\hat{C}_{*,j}, \hat{\sigma}_{*,j} \leftarrow \mathcal{M}_{gp,j}(\{(X^i)_{i=1}^N\})$
- 3: **end for**
- 4: $\hat{\mathbf{C}}_* \leftarrow (\hat{C}_{*,j})_{j=1}^Q$
- 5: $\hat{\mathbf{T}}_* \leftarrow \mathcal{M}_{DR}^{-1}(\hat{\mathbf{C}}_*)$ \triangleright Transferred signal prediction
- 6: $\hat{\mathcal{S}}_{*,T} \leftarrow \mathcal{M}_{DR}^{-1}(\{\hat{\sigma}_j\}_{j=1}^Q)$ \triangleright Transferred signal UQ
- 7: $\mu_* \leftarrow \frac{1}{n_*} \sum_{u=1}^{n_*} \delta_{\mathbf{F}_{*,u}}$
- 8: $\mathbf{P}_* \leftarrow \underset{\mathbf{P} \in U(n_*, n_{\text{ref}})}{\operatorname{argmin}} \mathcal{L}_{\lambda_0}(\mu_*, \mu_{\text{ref}}, \mathbf{P})$
- 9: $\hat{\mathbf{Y}}_* \leftarrow (n_* \mathbf{P}_*) \hat{\mathbf{T}}_*$ \triangleright Signal prediction
- 10: $\hat{\mathcal{S}}_* \leftarrow (n_* \mathbf{P}_*) \hat{\mathcal{S}}_{*,T} (n_* \mathbf{P}_*)^T$ \triangleright Signal UQ

Prediction. For a new test input X_* with n_* nodes, we first encode it as an empirical measure and compute the transport plan to the reference measure \mathbf{P}_* . We then compute the single-output GP predictions to form the predicted coefficients vector $\hat{\mathbf{C}}_* = (\hat{C}_{*,j})_{j=1}^Q$. We apply the inverse of the dimension reduction method to get the predicted transferred signal $\hat{\mathbf{T}}_* = \mathcal{M}_{DR}^{-1}(\hat{\mathbf{C}}_*)$ on the reference measure, which is transported back to estimate the predicted signal $\hat{\mathbf{Y}}_* = (n_* \mathbf{P}_*) \hat{\mathbf{T}}_*$. Prediction uncertainties can be obtained through sampling in the reduced space according to the GP posterior distribution, and then going backward through the exact same reconstruction stages. Note that in the particular case of a linear dimension reduction technique such as PCA, having access to the projection matrix \mathbf{E}^T yields an analytical formula for prediction uncertainties. Denoting

$\sigma_{*,1}, \dots, \sigma_{*,Q}$ the posterior standard deviation for all Q PCA coefficients, the posterior distribution of the transferred signal is Gaussian with covariance given by $\hat{\mathcal{S}}_{*,T} = \mathbf{E} \text{Diag}(\sigma_{*,1}^2, \dots, \sigma_{*,Q}^2) \mathbf{E}^T$ and similarly for the reconstructed signal $\hat{\mathcal{S}}_* = (n_* \mathbf{P}_*) \hat{\mathcal{S}}_{*,T} (n_* \mathbf{P}_*)^T$.

Reference measure and hyperparameters.

There are several possible strategies for constructing the reference measure. One can take a barycenter of measures, or a uniform measure over a reference shape, typically the convex hull of the union of the supports of all train measures. The most natural and simplest option is to sub-sample from an original measure coming from the training set. A low discrepancy sequence is built from this empirical measure (using the Maximum Mean Discrepancy as described in the supplementary material) in order to reduce its support to a few representative points. The TOS-GP method also depends on several hyperparameters that need to be selected. When using PCA, the low-embedding dimension Q can be chosen to achieve 95% cumulative explained variance. Note, however, that limiting the number of PCA coefficients to smaller values is already sufficient for equivalent quality when reconstruction errors are already large. The regularization parameter λ as well as the number of continuous WL iterations can be chosen either in a supervised manner or independently of training, so as to minimize the reconstruction error on the train samples only. Let us emphasize, however, that the number of continuous WL iterations and the choice of the reference measure are far less influential than regularization, as illustrated in our experiments.

5 EXPERIMENTS

For our numerical experiments ¹, we focus on regression tasks involving large graphs from mesh-based simulations in computational fluid dynamics and mechanics ². TOS-GP is compared to three competing approaches: Mesh Morphing Gaussian Process (MMGP) (Casenave et al., 2024) and two state-of-the-art GNN architectures: GCNN (Kipf and Welling, 2016) and MGN (Pfaff et al., 2020). In particular, we perform an in-depth study of the impact of various hyperparameters: the choice of the reference measure, the number of WL iterations and the regularization parameter.

¹Code: https://gitlab.com/drti/tos_gp/.

²Datasets: https://plaid-lib.readthedocs.io/en/latest/source/data_challenges.html.

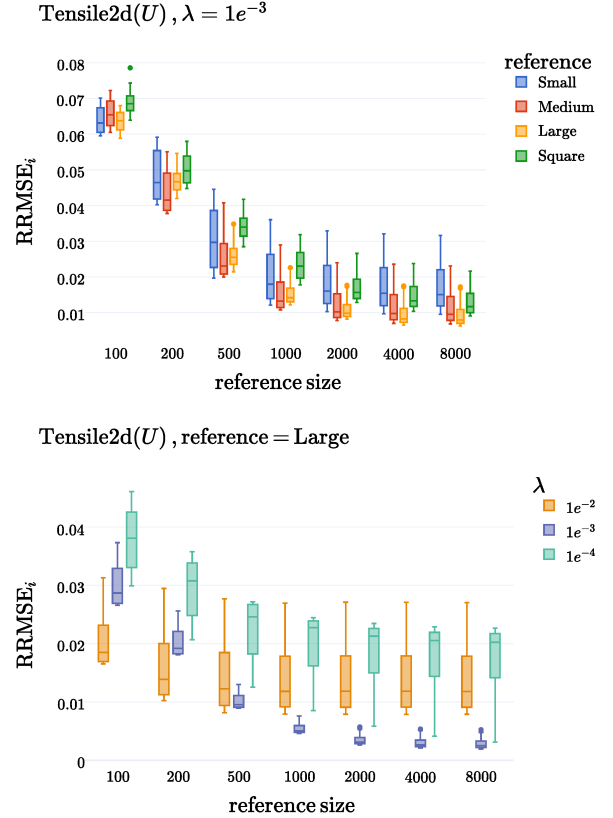


Figure 2: **Tensile2d** field **U**. Top: RRMSE vs reference size and reference measure with fixed $\lambda = 1e - 3$. Bottom: RRMSE vs reference size and λ for a fixed reference measure.

For all considered problems, we use PCA as a dimension reduction technique, and choose the number of coefficients so as to obtain 95% explained variance. The quality of trained models is assessed by computing the Relative Root Mean Square Error (RRMSE) following Casenave et al. (2024).

The RRMSE is defined for N_* ground truth test signals $\{\mathbf{Y}^{(i)}\}_{i=1}^{N_*}$, and the predictions $\{\hat{\mathbf{Y}}^{(i)}\}_{i=1}^{N_*}$ (where the signal i is of size n_{*i}) as:

$$\text{RRMSE} \left(\{\mathbf{Y}^{(i)}\}_{i=1}^{N_*}, \{\hat{\mathbf{Y}}^{(i)}\}_{i=1}^{N_*} \right)^2 = \frac{1}{N_*} \sum_{i=1}^{N_*} \text{RRMSE}_i^2, \quad (10)$$

where $\text{RRMSE}_i^2 = \frac{\|\mathbf{Y}^{(i)} - \hat{\mathbf{Y}}^{(i)}\|_2^2}{n_{*i} \|\mathbf{Y}^{(i)}\|_\infty^2}$ is the contribution of the i -th signals $\mathbf{Y}^{(i)}$ and $\hat{\mathbf{Y}}^{(i)}$ to RRMSE^2 . For each dataset, the methods are repeated 10 times to test their variability with respect to their own random initializations, and we report the empirical mean and standard deviation of these 10 RRMSE. Additional details on the GP model are given in the supplementary material.

Tensile2d dataset. We consider a two-dimensional problem in solid mechanics introduced by Casenave et al. (2023). The input geometries consist of 2D squares with two half circles that have been cut off in a symmetrical manner. These structures are subject to a uniform pressure field over the upper boundary and the material is modeled by a nonlinear elasto-viscoplastic law. The inputs of the problem are given by the mesh of the geometry (graphs with ~ 10000 nodes), and six scalar parameters that correspond to the material parameters and the input pressure applied to the upper boundary. The signal outputs of the problem are the horizontal component of the displacement field U and the tangential force per unit area acting in the horizontal direction on a surface normal to the second axis (shear stress σ_{12}) obtained with the **Zset** mechanical solver Garaud et al. (2019). There are 500 train and 200 test samples.

We consider four types of references: **Small**, **Medium**, **Large** and **Square**. They are described in the supplementary material. For each type, we consider several support sizes ranging from 100 to 8000. We first study the influence of these hyperparameters in Figure 2 for the **Tensile2d** field U . We notice that the RRMSE decreases as a function of the size of the reference measure, and seems to remain close to a constant beyond 1000 points. This suggests to choose the size of the reference measure as a tradeoff between computational time for OT and prediction errors. We also observe that the best regularization parameter is $\lambda = 1e-3$ which highlights the benefits of a well-chosen regularization. Interestingly, the **Large** reference measure with the largest surface gives slightly better scores with less variance than the **Medium** reference measure corresponding to the more central distribution. The **Square** reference gives equivalent performance when the size is high, but it seems to have more degraded predictions when the reference size is reduced compared to other references.

We also report in Table 1 the results of TOS-GP with an exhaustive grid search to identify the best regularization parameter and number of continuous WL iterations to minimize the reconstruction error on the training set, the experimental results of the concurrent methods being taken from Casenave et al. (2024).

Remark that we achieve competitive scores compared to the state-of-the-art methods on the **Tensile2d** dataset.




Figure 3: **Tensile2d**, σ_{12} : two test meshes (left and right) transported to a common reference. From top to bottom: output signals, predicted transferred signals, posterior standard deviation of the predicted transferred fields, posterior standard deviation of the predicted field.

Rotor37 dataset. The NASA rotor 37 case (Roy-nard et al., 2023) serves as a prominent example of a transonic axial-flow compressor rotor widely employed in computational fluid dynamics research. This dataset is made of 3D compressible steady-state Reynold-Averaged Navier-Stokes (RANS) simulations that model external flows (Ameri, 2009). The inputs

Table 1: Empirical means and standard deviations of the RRMSEs for all considered datasets and methods.

Method\Dataset	Rotor37(P)	Rotor37(T)	Tensile2d(U)	Tensile2d(σ_{12})
TOS-GP	3.4e-2 (6e-4)	9.6e-3 (2e-5)	2.2e-3 (8e-6)	5.6e-3 (3e-6)
GCNN	1.7e-2 (8e-4)	3.9e-3 (1e-4)	4.5e-2 (1e-2)	4.5e-2 (4e-3)
MGN	1.7e-2 (2e-3)	1.4e-2 (2e-3)	1.5e-2 (1e-3)	7.5e-3 (4e-4)
MMGP	7.2e-3 (5e-4)	8.2e-4 (1e-5)	3.4e-3 (4e-5)	2.4e-3 (2e-5)




Figure 4: **Rotor37** test field T for one test input mesh. From left to right: true field, predicted field (posterior mean), absolute error, posterior standard deviation.

are the FEM meshes of a 3D compressor blade (graphs with ~ 30000 nodes), as well as two additional scalar physical parameters corresponding to the rotational speed and the input pressure. The field outputs of the problem are the temperature T and the pressure P at each node. There are 1000 train and 200 test samples.

Similarly to the previous test case, we consider both a reference measure coming from the training set and the circumscribed sphere around the training samples, see Table 1 for the results obtained with the best set of hyperparameters. Here the RRMSE scores are comparable to those of GNN approaches but slightly lower than MMGP: this is due to the very strong discontinuity of the fields on the blade edge, which leads to a higher reconstruction error. Note that while the number of continuous WL iterations has little influence on **Tensile2d** outputs, it significantly helps with **Rotor37** output P as detailed in the supplementary material.

Prediction uncertainty. A key benefit of TOS-GP is its ability to assess prediction uncertainties. In Figure 3, we use the common **Medium** reference of size 1000 to shed light on the uncertainty propagation described in Section 4. After applying the inverse of the dimension reduction method, both predictions and uncertainties are expressed on the same reference measure (second and third lines). Predictive uncertainties are then propagated to the original space (fourth line). We observe more spread-out uncertainties on the right-hand input mesh, which is actually more difficult to predict due to more localized constraints (its RRMSE_i is equal to $2.9e-2$ versus $1.1e-2$ for the left-hand input mesh). Figure 4 highlights predictions for **Rotor37**’s field T along with their uncertainties and prediction errors. Interestingly, uncertainties are higher in “critical” zones where the signal is likely to have the most variations. These zones notably correspond to the areas where the error is the biggest. More figures detailing uncertainties for other data sets are given in the supplementary.

Implementation and computing infrastructure. We leverage a Python implementation of GP regression with (GPy, 2012). The optimal transport plans

are computed using the ott-jax library (Cuturi et al., 2022). All our analyses were performed on a hybrid (CPU/GPU) computational node using 1 Nvidia A100 GPU and 32 CPU cores (AMD EPYC Milan 7763) with a total of 128GB of RAM (4GB per core).

Computation times. The most computationally intensive step of TOS-GP is the identification of the transport plans during the preprocessing step. Transfer times depend on both the size of the measures involved and the regularization parameter. For **Tensile2d** using a reference measure of size 8000, computing one transport plan takes 6 seconds for $\lambda = 1e-3$. This corresponds to a total sequential preprocessing time of 2h38min for all training samples and all λ . Similarly, for **Rotor37** (temperature), computing one transport plan takes 45 seconds for a reference measure of size 15000 with $\lambda = 1e-7$, and a total sequential time of 28h20min for all λ and all training samples. But this preprocessing step is obviously embarrassingly parallel, reducing for example to 1min35s and 17min using 100 parallel jobs for **Tensile2d** and **Rotor37**, respectively. Computation times for concurrent methods are given in the supplementary material.

6 INPUT GEOMETRIES WITH VARYING TOPOLOGIES

In this section, we address a challenging regression problem involving topological changes in the input geometries. These topology variations can be handled by our methodology, which is not the case with MMGP. In addition, this dataset allows us to demonstrate the flexibility of the TOS-GP framework by considering different dimension reductions and kernels.

We introduce a two-dimensional multiscale problem in computational mechanics, where the goal is to solve a nonlinear boundary value problem to determine the macroscopic mechanical properties of a highly heterogeneous material at the microscopic scale as in Pasparakis et al. (2024). Four topologies are considered by generating microstructures with 5 to 8 pores. Mesh sizes range from 4119 to 7128 nodes. Examples of in-

put meshes which represent microstructures of porous hyperelastic materials are shown in the supplementary material. The mechanical problem is also parameterized by three input scalars (components of the right Cauchy deformation tensor) which define the boundary conditions. Train and test datasets of respective sizes 762 and 376 are generated by a constrained design of experiments. We consider the horizontal component of the displacement field U as the output signal.

The TOS-GP methodology is applied with a reference uniform grid of size 32×32 on the unit square. A regularization parameter $\lambda = 1e - 3$ is used. For the Gaussian process regression, we consider a kernel that computes the MMD distance between the centers of the pores and plugs them in a squared exponential kernel. For the dimension reduction part, we observe that the problem is highly non-linear and therefore cannot be easily reduced to a low dimensional space with PCA. We instead rely on convolutional autoencoders (Goodfellow et al., 2016) built on the transferred fields defined on the 32×32 unit grid. More specifically, the auto-encoder combines convolutional and 2D discrete Fourier layers based on the Fourier neural operator approach of Li et al. (2020), providing encodings in dimension $2 \times 8 \times 8 = 128$.

Figure 5 shows predictions for two sample meshes with 5 and 8 pores respectively. Even though their topologies change, regularized optimal transport allows to express output signals on the reference grid, thus enabling the predictions in this transferred space (see the rightmost figures). Here, we demonstrate that our method can seamlessly handle input meshes with varying topologies without requiring any modifications. It is worth emphasizing that to our best knowledge, such problems have so far only been addressed by relying on deep neural networks. A comparison with MGN is included in the supplementary.

7 CONCLUSION

We introduce TOS-GP, a first GP-based supervised model that can predict outputs defined as fields discretized on large graphs with continuous node attributes. It demonstrates similar predictive capabilities in comparison with state-of-the-art GNN architectures for several problems in computational physics, while ensuring uncertainty quantification on the nodes and involving few hyperparameters to tune. We focused here on datasets coming from physical simulations. But a key feature is that, since we make no assumption on the mesh topology nor on the graph adjacency, our methodology can also be readily extended to predict signals defined on point clouds or on meshes with topology changes. In addition, since the OT and




Figure 5: Two samples with different topologies from the multiscale dataset. From top to bottom: true field, predicted field, absolute error, predicted transferred signal.

dimension reduction steps are agnostic to the choice of the supervised model, one can consider replacing GP with any other single-output ML model. To go further, we observe that the regularized OT between measures of different sizes leads to an approximation error that may be high, which in practice can be compensated by an appropriate choice of the regularization parameter. But signals with higher variability, whose regularity varies greatly between inputs, needs specific care. As a perspective, in order to better take adjacency into account, we plan to investigate transport plans using the Fused Gromov Wasserstein distance (Vayer et al., 2019; Vincent-Cuaz et al., 2022), but this will certainly come with a significant computation overhead. Lastly, while our approach decouples output transfer and learning, one could also design variants of the work of Bachoc et al. (2020) and Bachoc et al. (2023) that reuse the same transport plans in both input and output spaces.

References

- Alvarez, M. A. and Lawrence, N. D. (2011). Computationally efficient convolved multiple output gaussian processes. *The Journal of Machine Learning Research*, 12:1459–1500.
- Alvarez, M. A., Rosasco, L., Lawrence, N. D., et al. (2012). Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3):195–266.
- Ameri, A. (2009). Nasa rotor 37 cfd code validation glenn-lt code. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, page 1060.
- Bachoc, F., Béthune, L., Gonzalez-Sanz, A., and Loubes, J.-M. (2023). Gaussian processes on distributions based on regularized optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pages 4986–5010. PMLR.
- Bachoc, F., Suvorikova, A., Ginsbourger, D., Loubes, J.-M., and Spokoiny, V. (2020). Gaussian processes with multidimensional distribution inputs via optimal transport and Hilbertian embedding. *Electronic Journal of Statistics*, 14(2):2742 – 2772.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Begoli, E., Bhattacharya, T., and Kusnezov, D. (2019). The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence*, 1(1):20–23.
- Brouard, C., Szafranski, M., and d’Alché Buc, F. (2016). Input output kernel regression: Supervised and semi-supervised structured output prediction with operator-valued kernels. *Journal of Machine Learning Research*, 17(176):1–48.
- Carmeli, C., De Vito, E., and Toigo, A. (2006). Vector valued reproducing kernel hilbert spaces of integrable functions and mercer theorem. *Analysis and Applications*, 4(04):377–408.
- Carpintero Perez, R., Da Veiga, S., Garnier, J., and Staber, B. (2024). Gaussian process regression with sliced Wasserstein Weisfeiler-Lehman graph kernels. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, pages 1297–1305.
- Casenave, F., Roynard, X., and Staber, B. (2023). Tensile2d: 2D quasistatic non-linear structural mechanics solutions, under geometrical variations.
- Casenave, F., Staber, B., and Roynard, X. (2024). Mmgp: a mesh morphing gaussian process-based machine learning method for regression of physical problems under nonparametrized geometrical variability. *Advances in Neural Information Processing Systems*, 36.
- Courty, N., Flamary, R., Habrard, A., and Rakotomamonjy, A. (2017). Joint distribution optimal transportation for domain adaptation. *Advances in neural information processing systems*, 30.
- Courty, N., Flamary, R., Tuia, D., and Rakotomamonjy, A. (2016). Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 39(9):1853–1865.
- Cuturi, M., Meng-Papaxanthos, L., Tian, Y., Bunne, C., Davis, G., and Teboul, O. (2022). Optimal transport tools (ott): A jax toolbox for all things wasserstein. *arXiv preprint arXiv:2201.12324*.
- Dong, X., Thanou, D., Toni, L., Bronstein, M., and Frossard, P. (2020). Graph signal processing for machine learning: A review and new perspectives. *IEEE Signal processing magazine*, 37(6):117–127.
- Donnelly, J., Daneshkhah, A., and Abolfathi, S. (2024). Forecasting global climate drivers using gaussian processes and convolutional autoencoders. *Engineering Applications of Artificial Intelligence*, 128:107536.
- Ferradans, S., Papadakis, N., Peyré, G., and Aujol, J.-F. (2014). Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3):1853–1882.
- Garaud, J.-D., Rannou, J., Bovet, C., Feld-Payet, S., Chiaruttini, V., Marchand, B., Lacourt, L., Yastrebov, V., Osipov, N., and Quilici, S. (2019). Z-set-suite logicielle pour la simulation des matériaux et structures. In *14ème Colloque National en Calcul des Structures*.
- Gawlikowski, J., Tassi, C. R. N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., et al. (2023). A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1):1513–1589.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Goovaerts, P. (1997). *Geostatistics for natural resources evaluation*. Oxford University Press, USA.

- GPy (since 2012). GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>.
- Gramacy, R. B. (2020). *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. CRC press.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773.
- Iollo, A. and Taddei, T. (2022). Mapping of coherent structures in parameterized flows by learning optimal transportation with gaussian models. *Journal of Computational Physics*, 471:111671.
- Jaber, E., Blot, V., Brunel, N., Chabridon, V., Remy, E., Iooss, B., Lucor, D., Mougeot, M., and Leite, A. (2024). Conformal approach to gaussian process surrogate evaluation with coverage guarantees. *arXiv preprint arXiv:2401.07733*.
- Jackson, J. E. (2005). *A user’s guide to principal components*. John Wiley & Sons.
- Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13:455–492.
- Journel, A. G. and Huijbregts, C. J. (1976). Mining geostatistics.
- Kadri, H., Duflos, E., Preux, P., Canu, S., Rakotomamonjy, A., and Audiffren, J. (2016). Operator-valued kernels for learning from functional response data. *Journal of Machine Learning Research*, 17(20):1–54.
- Kadri, H., Ghavamzadeh, M., and Preux, P. (2013). A generalized kernel approach to structured output learning. In *International Conference on Machine Learning*, pages 471–479. PMLR.
- Kipf, T. N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kontolati, K., Loukrezis, D., Giovanis, D. G., Vandanapu, L., and Shields, M. D. (2022). A survey of unsupervised learning methods for high-dimensional uncertainty quantification in black-box-type problems. *Journal of Computational Physics*, 464:111313.
- Kriege, N. M., Johansson, F. D., and Morris, C. (2020). A survey on graph kernels. *Applied Network Science*, 5(1):1–42.
- Li, Z., Huang, D. Z., Liu, B., and Anandkumar, A. (2022). Fourier neural operator with learned deformations for pdes on general geometries. *arXiv preprint arXiv:2207.05209*.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. (2020). Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.
- Lim, D., Robinson, J., Zhao, L., Smidt, T., Sra, S., Maron, H., and Jegelka, S. (2022). Sign and basis invariant networks for spectral graph representation learning. *arXiv preprint arXiv:2202.13013*.
- Montesuma, E. F., Mboula, F. N., and Souloumiac, A. (2023). Recent advances in optimal transport for machine learning. *arXiv preprint arXiv:2306.16156*.
- Nanty, S., Helbert, C., Marrel, A., Pérot, N., and Prieur, C. (2017). Uncertainty quantification for functional dependent random variables. *Computational Statistics*, 32:559–583.
- Nikolentzos, G., Siglidis, G., and Vazirgiannis, M. (2021). Graph kernels: A survey. *Journal of Artificial Intelligence Research*, 72:943–1027.
- Owhadi, H. (2021). Notes on operator valued kernels, feature maps and gaussian processes.
- Pasparakis, G. D., Graham-Brady, L., and Shields, M. D. (2024). Bayesian neural networks for predicting uncertainty in full-field material response. *arXiv preprint arXiv:2406.14838*.
- Perrin, T., Roustant, O., Rohmer, J., Alata, O., Naulin, J., Idier, D., Pedreros, R., Moncoulon, D., and Tinard, P. (2021). Functional principal component analysis for global sensitivity analysis of model with spatial output. *Reliability Engineering & System Safety*, 211:107522.
- Peyré, G., Cuturi, M., et al. (2019). Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P. (2020). Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*.
- Ricaud, B., Borgnat, P., Tremblay, N., Gonçalves, P., and Vanderghynst, P. (2019). Fourier could be a data scientist: From graph fourier transform to signal processing on graphs. *Comptes Rendus. Physique*, 20(5):474–488.
- Roynard, X., Casenave, F., and Staber, B. (2023). Rotor37: a 3D CFD RANS dataset, under geometrical variations of a compressor blade.
- Sejdinovic, D., Sriperumbudur, B., Gretton, A., and Fukumizu, K. (2013). Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *The Annals of Statistics*, pages 2263–2291.

- Shen, Z., Feydy, J., Liu, P., Curiale, A. H., San Jose Estepar, R., San Jose Estepar, R., and Nethammer, M. (2021). Accurate point cloud registration with robust optimal transport. *Advances in Neural Information Processing Systems*, 34:5373–5389.
- Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. (2013). The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE signal processing magazine*, 30(3):83–98.
- Togninalli, M., Ghisu, E., Llinares-López, F., Rieck, B., and Borgwardt, K. (2019). Wasserstein Weisfeiler-Lehman graph kernels. *Advances in neural information processing systems*, 32.
- Vayer, T., Courty, N., Tavenard, R., and Flamary, R. (2019). Optimal transport for structured data with application on graphs. In *International Conference on Machine Learning*, pages 6275–6284. PMLR.
- Venkitaraman, A., Chatterjee, S., and Handel, P. (2020). Gaussian processes over graphs. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5640–5644. IEEE.
- Vincent-Cuaz, C., Flamary, R., Corneli, M., Vayer, T., and Courty, N. (2022). Template based graph neural network with optimal transport distances. *Advances in Neural Information Processing Systems*, 35:11800–11814.
- Vishwakarma, G., Sonpal, A., and Hachmann, J. (2021). Metrics for benchmarking and uncertainty quantification: Quality, applicability, and best practices for machine learning in chemistry. *Trends in Chemistry*, 3(2):146–156.
- Weston, J., Chapelle, O., Vapnik, V., Elisseeff, A., and Schölkopf, B. (2002). Kernel dependency estimation. *Advances in neural information processing systems*, 15.
- Weston, J., Schölkopf, B., and Bousquet, O. (2005). Joint kernel maps. In *International Workshop Conference on Artificial Neural Networks*, pages 176–191. Springer.
- Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. MIT press.
- Zhi, Y.-C., Ng, Y. C., and Dong, X. (2023). Gaussian processes on graphs via spectral kernel learning. *IEEE Transactions on Signal and Information Processing over Networks*.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Not Applicable]
 - (b) Complete proofs of all theoretical results. [Not Applicable]
 - (c) Clear explanations of any assumptions. [Not Applicable]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Yes]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]

5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Supplementary Material for Learning signals defined on graphs with optimal transport and Gaussian process regression

A More experimental details

This section provides more information about the experiments described in the main paper. We summarize the datasets used for the experiments. We give more details on the Gaussian process model, on the choice of the reference measures and on the hyperparameter selection. We also provide a more precise analysis of the preprocessing and training times.

A.1 Datasets

Details about the datasets are given in Table 2. The attributes column corresponds to the dimension of the continuous attributes associated to each node.

Table 2: Summary of the datasets. (*): fixed number of nodes and adjacency structure

Dataset	Train+Test samples	Mean number of nodes	Mean number of edges	Number of attributes	Number of input scalars	Number of output fields
Rotor37*	1000+200	29773	77984	3	2	2
Tensile2d	500+200	9425.6	27813.8	2	6	2
Multiscale	764+376	4591.6	16721.6	2	3	1

For **Tensile2d**, we consider three references in the train set respectively called **Small**, **Medium** and **Large** with original size 6143, 9733 and 11627, respectively. These **Small**, **Medium**, **Large** inputs actually correspond to the largest, medium-size and smallest radius of the half circles as shown in Figure 6. Such measures are then sub-sampled to obtain new supports of pre-defined sizes included in the original supports.




Figure 6: **Small**, **Medium** and **Large** input meshes from the **Tensile2d** dataset (from left to right). The meshes are coarsened here for visualization purpose only.

A.2 Gaussian Processes

The GP kernel function is chosen as the tensorized kernel

$$k(\mathfrak{X}, \mathfrak{X}') := \sigma^2 k_{\text{SWWL}}(X, X') \prod_{\ell=1}^m c_{5/2}(|s_\ell - s'_\ell|) \quad (11)$$

where $\mathfrak{X} = (X, s_1, \dots, s_m)$ and $\mathfrak{X}' = (X', s'_1, \dots, s'_m)$ denote the input graphs and the m input scalars s_1, \dots, s_m , $c_{5/2}$ is the Matérn-5/2 covariance function, σ^2 is a variance parameter and k_{SWWL} is a SWWL kernel with 50 projections and 500 quantiles. The lengthscale parameters of the SWWL and Matérn-5/2 kernels are optimized simultaneously by maximizing the marginal log-likelihood (Williams and Rasmussen, 2006).

A.3 Choice of the reference measure

The reference measure can be chosen in various ways: selecting an arbitrary measure within the training dataset or computing a barycenter, for instance. In our work, we build a reference measure by subsampling a selected measure in the training dataset using a procedure described in the sequel of this section. Let X be an arbitrary graph with node features $\mathbf{F}_1, \dots, \mathbf{F}_n$. Let then $\mu = \frac{1}{n} \sum_{u=1}^n \delta_{\mathbf{F}_u}$ be the associated empirical measure. We subsample the measure μ using a procedure that minimizes the Maximum Mean Discrepancy (MMD) Gretton et al. (2012).

Definition 2 (Maximum Mean Discrepancy). *Let x and y be random variables defined on a topological space \mathcal{Z} , with respective Borel probability measures p and q . Let $k : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ be a kernel function and let $\mathcal{H}(k)$ be the associated reproducing kernel Hilbert space. The maximum mean discrepancy between p and q is defined as*

$$\text{MMD}_k(p, q) = \sup_{\|f\|_{\mathcal{H}(k)} \leq 1} |\mathbb{E}_{x \sim p}[f(x)] - \mathbb{E}_{y \sim q}[f(y)]|.$$

The MMD admits the following closed-form expression:

$$\text{MMD}_k(p, q)^2 = \mathbb{E}_{x \sim p, x' \sim p}[k(x, x')] + \mathbb{E}_{y \sim q, y' \sim q}[k(y, y')] - 2\mathbb{E}_{x \sim p, y \sim q}[k(x, y)], \quad (12)$$

which can be estimated thanks to U- or V-statistics. Here k is chosen as the distance-induced kernel $k(x, y) = \|x\| + \|y\| - \|x - y\|$ which has been shown to be characteristic by Sejdinovic et al. (2013).

The MMD subsampling algorithm 3 selects $m < n$ particles by greedily minimizing the MMD. At the $(i + 1)$ -th iteration, a new particle F_j is selected by minimizing the MMD between μ and the empirical measure given by

$$\mu_{i+1}(j) = \frac{1}{i+1} \sum_{\ell \in \mathcal{P}_i} \delta_{\mathbf{F}_\ell} + \frac{1}{i+1} \delta_{\mathbf{F}_j}, \quad j = 1, \dots, n,$$

where $\mathcal{P}_i \subset \{1, \dots, n\}$ denotes the set of particles indices already selected by the previous i iterations.

Algorithm 3 MMD subsampling

Input: Empirical measure μ , kernel k , subsample size m

Output: Subsampled measure μ'

- 1: $\pi_1 \leftarrow \operatorname{argmin}_{j \in \{1, \dots, n\}} \text{MMD}_k^2(\mu, \delta_{\mathbf{F}_j})$
 - 2: $\mathcal{P}_1 = \{\pi_1\}$
 - 3: **for** $i = 1, \dots, m - 1$ **do**
 - 4: $\pi_{i+1} \leftarrow \operatorname{argmin}_{j \in \{1, \dots, n\} \setminus \mathcal{P}_i} \text{MMD}_k^2(\mu, \mu_{i+1}(j))$
 - 5: $\mathcal{P}_{i+1} \leftarrow \mathcal{P}_i \cup \{\pi_{i+1}\}$
 - 6: **end for**
 - 7: $\mu' \leftarrow \frac{1}{m} \sum_{j \in \mathcal{P}_m} \delta_{\mathbf{F}_j}$
-

When the input measures lie in a two-dimensional space, we can also choose a shape corresponding to the convex hull of the supports of all train measures and then choose a uniform distribution of suitable size on it. This is less obvious in dimension three when the measure is the discretization of a 2-manifold. For **Rotor37**, we can use

a circumscribed sphere to the union of all the supports of train measures. In practice, the scores obtained with such a reference are very similar to those obtained with a reference measure from the train. It is further possible to consider an optimal transport barycenter of the input measures as a reference. For **Tensile2d**, we observe that the barycenter is close to the input measure **Medium** already considered. For **Rotor37**, the barycenter no longer corresponds to the discretization of a 2-manifold, which is less relevant than choosing a train measure as a reference. Instead, it is possible to select a "representative" measure from the set of train measures. This can be done using MMD again, but this time in a graph space. The MMD subsampling algorithm 3 is applied to the set of train graphs using the SWWL kernel (with lengthscales equal to the median of pairwise SWWL pseudo-distances), with only one step to identify a central train input in the sense of SWWL pseudo-distances. In practice, we again observe few differences when taking the first input or the latter one as reference measures due to the proximity of the geometries.

A.4 Hyperparameter selection

In addition to the choice of the reference measure, the method involves various hyperparameters such as the number of continuous WL iterations, the regularization parameter and the number of PCA coefficients. For **Tensile2d**, we obtain subsampled versions of the **Small**, **Medium** and **Large** references with 100, 200, 500, 1000, 2000, 4000, 8000 points each with the MMD procedure described previously. A uniform distribution on the convex hull of all inputs (a square) with the same sizes is also considered. We vary the regularization parameter λ in $\{1e-4, 1e-3, 1e-2\}$. For **Rotor37**, we consider the first input measure in the dataset as a reference with 100, 1000, 5000, 10000 and 15000 points. Some preliminary tests with other references showed that it was not relevant to change it with another train input. The regularization parameter λ is selected in $\{1e-7, 1e-6, 1e-5, 1e-4\}$ for the temperature field and in $\{1e-8, 1e-7, 1e-6\}$ for the pressure. The number of continuous WL iterations is selected in $\{0, 1, 2, 3\}$ for all the problems.

The number of PCA coefficients is chosen so as to obtain a cumulative variance ratio greater than 95%. If this number is smaller than 4, we decide to use a minimum number of coefficients equal to 4. In practice, we thus need to learn 4 coefficients for both fields of **Tensile2d** and **Rotor37(T)**, and 51 coefficients for **Rotor37(P)**. Remark that taking only 10 coefficients for **Rotor37(P)** gives a similar prediction error between test and predicted transferred fields due to an already high approximation error coming from the optimal transport part. In the case of **Tensile2d**, we add a warm-start initialization of the lengthscales, variances and nuggets based on a common GP trained to learn all coefficients simultaneously in order to help find a better optimum of the likelihood function.

Results given in Table 1 correspond to the following parameters:

- **Tensile2d(U)**: **Large** reference measure with size 8000, $\lambda = 1e-3$, 1 WL iteration, 4 PCA coefficients
- **Tensile2d(σ_{12})**: **Large** reference measure with size 8000, $\lambda = 1e-3$, 2 WL iterations, 4 PCA coefficients
- **Rotor37(T)**: reference measure with size 15000, $\lambda = 1e-6$, 1 WL iteration, 10 PCA coefficients
- **Rotor37(P)**: reference measure with size 15000, $\lambda = 1e-8$, no WL iteration, 10 PCA coefficients

A.5 Detailed computation times

The most costly steps are the obtention of transport plans during the preprocessing step and the learning phase with Gaussian processes. Transfer times depend on both the size of the measures involved and the regularization parameter. For **Tensile2d** using the **Large** reference measure with size 8000, computing optimal transport plans takes 3, 6 and 10 seconds for $\lambda = 1e-2$, $\lambda = 1e-3$ and $\lambda = 1e-4$ respectively. This corresponds to a total sequential preprocessing time of 9500 seconds (2h38min) for all regularizations, which can be easily reduced by computing transport plans in parallel (in practice, only 1min35s for the entire train dataset with 100 parallel jobs). Considering all measure types (four choices), sizes (four choices with supports of respective sizes 100, 200, 500, 1000, 2000, 4000, and 8000) and regularizations (three choices) involved in Figure 2, this corresponds to a total sequential time of 142830 seconds (24 minutes using 100 parallel jobs). Surprisingly, changing the number of continuous WL iterations (and therefore the size of the points forming the input and output empirical measures) has a negligible impact on transport plan computation time. Remark that for the test phase, a preprocessing step to obtain the test transport plans is also necessary. It takes 1200 seconds in total for the 200 test outputs (again possible to parallelize to be broken down to a 12 seconds using 100 jobs). Gaussian processes take 8min40s to train (with 3 restarts).

Similarly, for **Rotor37**, transport plans take 1, 4, 16, 34 and 52 seconds when the reference is the first input being subsampled to have respective support size equal to 100, 1000, 5000, 10000 and 15000 with $\lambda = 1e - 8$. Fixing the support size to 15000, transport plans take respectively 52, 45, 40, 6 and 2 seconds for $\lambda = 1e - 8, 1e - 7, 1e - 6, 1e - 5, 1e - 4$. For the pressure field, considering all sizes (100, 1000, 5000, 10000 and 15000) and regularizations (three choices, $\lambda = 1e - 8, 1e - 7, 1e - 6$), this corresponds to a total sequential time of 316430 seconds (53 minutes using 100 parallel jobs). For the temperature field, considering all sizes (100, 1000, 5000, 10000 and 15000) and regularizations (four choices, $\lambda = 1e - 7, 1e - 6, 1e - 5, 1e - 4$), this corresponds to a total sequential time of 206127 seconds (34 minutes using 100 parallel jobs). Gaussian processes take 2h to train (with 3 restarts) for the field P and 11min for the field T. In comparison, GCNN takes respectively 24h and 1h for **Rotor37** and **Tensile2d** and MGN takes respectively 13h and 7h for **Rotor37** and **Tensile2d** for each hyperparameter set according to Casenave et al. (2024).

B More experimental results

In this section, we add experimental results. In particular, we present visuals of the predicted fields for all the datasets in Section 5 in addition to those of the main paper. We also give details about the errors, and the influence of the number of continuous WL iterations. Lastly, we present an additional comparison with MGN for the multiscale hyperelasticity problem of Section 6.

B.1 Predicted fields and uncertainties




Figure 7: From top to bottom: **Tensile2d** fields U and T , **Rotor37** field P for one test input mesh. From left to right: true field, predicted field, absolute error, standard deviation from the posterior law.

Figure 7 displays the predictions and uncertainties for the datasets of Section 5. Similarly to Figure 4, uncertainties are concentrated in areas where the signal varies the most and also seem to correspond to high error zones. For **Rotor37**(P), we observe a strange behavior with the prediction, which shows small areas on the left. This is due to a very large signal discontinuity on the slice of the mesh (with very high signal values), which is sent to the face via the regularized optimal transport plan. Unfortunately, this error is found as early as the transfer phase, and taking smaller regularizations does not improve this transfer error. Both Figure 4 and Figure 7 represent the first test input of the **Rotor37** dataset. Figure 7 and Figure 3 respectively use the test inputs with the largest and smallest radius of the half circles.

B.2 Prediction errors

The detail of the prediction error due to the reconstruction and due to the Gaussian process model is given in Table 3. The approximation error $\text{RRMSE}\left(\{\mathbf{Y}^{(i)}\}_{i=1}^{N*}, \{\hat{\mathbf{Y}}^{(i)}\}_{i=1}^{N*}\right)$ corresponds to the error between true and approximated signals (that is to say the signal obtained by transferring the true field, and then transferring it back without any prediction) while the transferred prediction error $\text{RRMSE}\left(\{\mathbf{T}^{(i)}\}_{i=1}^{N*}, \{\hat{\mathbf{T}}^{(i)}\}_{i=1}^{N*}\right)$ corresponds to the error between the true and predicted signals in the transferred space. Note that the total error $\text{RRMSE}\left(\{\mathbf{Y}^{(i)}\}_{i=1}^{N*}, \{\hat{\mathbf{Y}}^{(i)}\}_{i=1}^{N*}\right)$ is not the sum of these errors. For all the considered problems, the error is dominated by the approximation with optimal transport.

Table 3: RRMSE for the successive stages of TOS-GP: errors between test and approximated signals (Approximation), errors between test and predicted transferred signals (Transferred prediction), and errors between test and predicted signals (Total).

Stage\Dataset	Rotor37(P)	Rotor37(T)	Tensile2d(U)	Tensile2d(σ_{12})
Approximation	3.29e-2	9.51e-3	1.90e-3	4.55e-3
Transferred Prediction	2.59e-2	2.08e-3	1.35e-3	3.37e-3
Total	3.36e-2	9.63e-3	2.23e-3	5.57e-3

Table 4 shows RRMSE scores as a function of the number of continuous WL iterations selected. This parameter has little influence on the error.

Table 4: RRMSE scores depending on the number of continuous WL iterations.

WL iterations\Dataset	Rotor37(P)	Rotor37(T)	Tensile2d(U)	Tensile2d(σ_{12})
0	4.38e-2	9.63e-3	2.91e-3	9.60e-3
1	3.36e-2	9.82e-3	2.23e-3	6.41e-3
2	3.52e-2	1.01e-2	2.35e-3	5.57e-3
3	3.71e-2	1.04e-2	2.34e-3	5.59e-3

B.3 Multiscale hyperelasticity problem: comparison with MGN

Figure 8 shows four meshes corresponding to the four different topologies of the multiscale problem with 5 to 8 pores. Both the positions of the pores and their number can vary between samples.

In addition to the results for field prediction with TOS-GP, we carried out experiments with MGN in order to compare them. We highlight the predicted fields with both methods for two samples with 5 and 7 pores respectively in Figure 9. We note that predictions with MGN do not have the expected regularity of a solution to such a computational mechanics problem. Moreover, MGN predicts poorly parts of the field with high absolute values, with a kind of compression towards the mean values. Such shortcomings are absent when using TOS-GP. The previous observations are also confirmed by the global errors over the entire dataset, as the test RRMSE of MGN is equal to 0.053 compared to 0.044 for TOS-GP.




Figure 8: Examples of input meshes from the multiscale problem.




Figure 9: Two samples with different topologies from the multiscale dataset. From left to right: true field, predicted field with TOS-GP, predicted field with MGN.