



Multi-Fidelity Surrogate Models for High Dimensional Aerodynamic Data


PhD Student:
Javier Nieto-Centenero

Supervisor/s:
Rodrigo Castellanos (UC3M)
Esther Andres-Perez (INTA)

PhD Aerospace Engineering UC3M
Doctoral Meetings 2026
2nd and 3rd June 2026


What is Multi-Fidelity?

Potential Flow




What is Multi-Fidelity?

RANS / LES




Potential Flow



What is Multi-Fidelity?


RANS / LES



Wind Tunnel / Flying Test




Potential Flow



What is Multi-Fidelity?


RANS / LES




Wind Tunnel / Flying Test



Potential Flow



Multi-fidelity Model



Key idea:

Leverage the **abundance of inexpensive Low-Fidelity (LF)** data and the **accuracy of limited High-Fidelity (HF)** data to build effective surrogate models.


SPARSE HF

- Few HF points, distributed **sparsely**.
- **Examples:** experimental point measurements (sensors, strain gauges, ...)
- **Strategy:** reconstruct the full field.



HIGHER-RESOLUTION HF

- HF with **finer resolution**.
- **Examples:** finer resolution simulations (CFD with more nodes).
- **Strategy:** map or translate between fidelities.




SPARSE HF


- Few HF points, distributed **sparsely**.
- **Examples:** experimental point measurements (sensors, strain gauges, ...)
- **Strategy:** reconstruct the full field.




HIGHER-RESOLUTION HF

- HF with **finer resolution**.
- **Examples:** finer resolution simulations (CFD with more nodes).
- **Strategy:** map o translate between fidelities.






Snapshots of C_p
for different flight
conditions




X
Matrix of
snapshots


=



U
Spatial modes



S
Singular values



V^*
Temporal modes

$$X = U S V^*$$

$$\tilde{u} = \begin{bmatrix} 0 \\ \blacksquare \\ 0 \\ 0 \\ 0 \\ \blacksquare \\ 0 \\ 0 \\ \blacksquare \\ 0 \\ 0 \\ \vdots \\ \vdots \\ \vdots \end{bmatrix}$$

$\blacksquare$ Cp from wind tunnel tests located at their position in the full Cp vector.


$$\tilde{u} = P^T u \approx P^T \sum_{k=1}^d \tilde{a}_k \Psi_k = P^T U_d \tilde{a}$$

It is necessary to determine the coefficients $\tilde{a}$ that minimize the error in the approximation

$$\tilde{a} = \arg \min_a \|P^T U_d a - P^T u\|_2^2$$

Cp vector reconstructed:

$$u \approx U_d \tilde{a}$$



- **Gaussian Process (GP):** Joint Gaussian distribution over variables, specifying a prior over functions:

$$f(x) \sim \mathcal{GP}(m(x), k(x, x'))$$

- **GPR:** Non-parametric model ($y = f(x) + \varepsilon$) that outputs both the predictive **mean** ($\bar{f}_*$) and its associated **uncertainty** ($\text{var}(f_*)$).


Linear multi-fidelity model

$$f_t(\mathbf{x}) = \rho_t f_{t-1}(\mathbf{x}) + \delta_t(\mathbf{x})$$

We can express the model as a single Gaussian Process:


$$\begin{bmatrix} f_{t-1} \\ f_t \end{bmatrix} \sim GP \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} K_{t-1} & \rho K_{t-1} \\ \rho K_{t-1} & \rho^2 K_{t-1} + K_{\delta_t} \end{bmatrix} \right)$$

The problem is decoupled into many GPR problems as fidelities have. This resolutions scheme is hierarchical.




SPARSE HF

- Few HF points, distributed **sparsely**.
- **Examples:** experimental point measurements (sensors, strain gauges, ...)
- **Strategy:** reconstruct the full field.



HIGHER-RESOLUTION HF


- HF with **finer resolution**.
- **Examples:** finer resolution simulations (CFD more nodes).
- **Strategy:** map or translate between fidelities.





Multi-Fidelity Autoencoder Framework: Fuses Low and High-Fidelity (LF/HF) aerodynamic data using **Transfer Learning** and **Conformal Prediction** (CP) for uncertainty quantification.


Test Case (XRF1 Wing):

- **Pressure Coefficient (C_p) Dimensionality:** HF mesh contains 49,574 nodes, while the LF input represents a coarser grid of 4,000 nodes.
- **Train-Test Split (435 total cases):** Limited HF data with 65 cases for training and 22 for testing. LF training utilizes the 413 non-test cases.









Conformal Prediction (CP) is a **distribution-free framework** that requires **no prior statistical assumptions** and empirically estimates uncertainty bands from errors observed on a calibration dataset.

- A **Non-conformity criterion** is used to quantify the deviation of the prediction within a calibration set. Use of specific metrics (e.g., absolute error, RMSE, ...).
- Uncertainty bandwidth is determined by a **target coverage level** ($1 - \delta$), enforced by calculating a **critical global quantile** from the empirical distribution of the calibration non-conformity scores.
- **Adaptive spatial modulation** can be achieved by scaling the uncertainty bounds according to the local distribution of errors across the domain.






Figure adapted from: Lei et al. (2018), *Distribution-Free Predictive Inference for Regression*, JASA 113(523).



Data Management

- Strict isolation between cross-fidelity data and test cases to prevent data leakage.




Data Management

- Strict isolation between cross-fidelity data and test cases to prevent data leakage.

LF Pre-training

- Leverages abundant LF data to extract core aerodynamic physics into a compact latent space.



Data Management


- Strict isolation between cross-fidelity data and test cases to prevent data leakage.

LF Pre-training

- Leverages abundant LF data to extract core aerodynamic physics into a compact latent space.

MSCP Calibration

- Clones the model into independent replicates to perform CP for B different calibration sets.
- Dynamically divides HF data into temporary training and calibration subsets per replicate.
- Records the optimal convergence epochs ($e^{(b)}$) and CP uncertainty margins ($R^{(b)}$) from each split.



Data Management

- Strict isolation between cross-fidelity data and test cases to prevent data leakage.

LF Pre-training


- Leverages abundant LF data to extract core aerodynamic physics into a compact latent space.

MSCP Calibration

- Clones the model into independent replicates to perform CP for B different calibration sets.
- Dynamically divides HF data into temporary training and calibration subsets per replicate.
- Records the optimal convergence epochs ($e^{(b)}$) and CP uncertainty margins ($R^{(b)}$) from each split.

Final Tuning & Aggregation

- Trains the final predictive model using 100% of the available high-fidelity data.
- Aggregates the B calculated CP margins using the median for robust uncertainty bands.




- **High-Dimensional Field Mapping:** Maps LF inputs ($D_{LF} = 4000$ nodes) to a HF mesh ($D_{HF} = 49574$ nodes) under severe training data constraint ($N = 65$ snapshots).
- **Transonic Shock Resolution:** Correctly maps localized pressure gradients and spatial positions of shock waves.

Test		
MAE	RMSE	R^2
0.046	0.084	0.965

- Low-Fidelity Correction:**
Model decouples from severely distorted LF inputs and accurately recovers HF pressure profiles.
- Statistical Band Optimization:**
The Normalized L_2 norm scoring function achieves a robust 95% pointwise coverage on test set while keeping narrow, engineering-actionable prediction interval widths (0.34).

	Normalized L_2 norm Coverage		
	Nominal	Pointwise	Band Width + Std
Test	0.05	0.95	[0.34, 0.17]



- **Low-Fidelity**

Model decouples from severely distorted LF inputs pressure profiles.

- **Statistical**

The Normalized L₂ norm scoring function achieves coverage on test set while keeping narrow, engineering-actionable prediction interval widths (0.34).

Javier Nieto-Centenero, Esther Andrés, Rodrigo Castellanos: “Multi-fidelity aerodynamic data fusion by autoencoder transfer learning”, 2025; [arXiv:2512.13069](https://arxiv.org/abs/2512.13069).

Paper under review in the **AIAA Journal**




WCCM–ECCOMAS 2026, 17th World Congress on Computational Mechanics & 10th European Congress on Computational Methods in Applied Sciences and Engineering, Munich, Germany, 19–24 July 2026.

$Mach = 0.60, \alpha = 0.0^\circ$

$Mach = 0.86, \alpha = 6.0^\circ$

$Mach = 0.90, \alpha = 3.3^\circ$



Test



Questions?


Multi-Fidelity Surrogate Models for High Dimensional Aerodynamic Data

PhD Student: Javier Nieto-Centenero
Supervisor/s: Rodrigo Castellanos
(UC3M), Esther Andres-Perez (INTA)


BACKUP SLIDES

Table 1 Model architecture and hyperparameters for Airfoils and Wing databases.

Database	Airfoils	Wing
Optimizer	Adam	
Layers Type	Fully connected	
Activation Type	ReLU	
Pre-training Learning Rate	1×10^{-4}	1×10^{-4}
Fine-tuning Learning Rate	1×10^{-4}	1×10^{-5}
Input size (D_{LF})	[40, 66, 104, 156, 222, 260]	4000
Encoder	64 - 32 - 16	1024 - 512 - 256 - 64
Latent Space	3	2
Decoder	16 - 32 - 16	64 - 256 - 512 - 1024
Up-Scaler	$1.5 \times D_{LF}$	8000
Output size (D_{HF})	260	49574
MSCP Splits (B)	30	30




(a) P • 0.5 ■ 0.6 ▲ 0.7 ◆ 0.8 XX • • •



(b) α • • •

Latent spaces: (a) Airfoils , (b) Wing



Predictive C_p for NACA 6513 (Airfoils database)

Model performance

LF/HF Dim (%)	Test			Complementary Test		
	MAE	RMSE	R^2	MAE	RMSE	R^2
15	0.034	0.110	0.995	0.039	0.100	0.996
25	0.027	0.080	0.997	0.032	0.079	0.997
40	0.025	0.070	0.998	0.028	0.070	0.998
60	0.018	0.055	0.999	0.020	0.054	0.999
85	0.019	0.064	0.998	0.021	0.063	0.998
100	0.025	0.090	0.996	0.022	0.064	0.998
100 (*)	0.018	0.053	0.999	0.021	0.058	0.999

MSCP Coverage

	LF/HF Dim (%)	Normalized L_2 norm Coverage			L_∞ norm Coverage		
		Nominal	Pointwise	Band width + Std	Nominal	Pointwise	Band width + Std
Test	15	0.75	0.96	[0.17, 0.14]	0.95	0.99	[0.50, 0.39]
	25	0.60	0.96	[0.16, 0.14]	0.95	0.99	[0.50, 0.39]
	40	0.80	0.97	[0.15, 0.11]	0.95	0.99	[0.50, 0.43]
	60	0.85	0.97	[0.14, 0.10]	0.95	0.99	[0.45, 0.36]
	85	0.90	0.98	[0.14, 0.10]	0.95	0.99	[0.46, 0.32]
	100	0.85	0.97	[0.18, 0.22]	0.95	0.99	[0.54, 0.57]
	100 (*)	0.75	0.97	[0.12, 0.09]	0.95	0.99	[0.25, 0.18]
Complementary Test	15	0.66	0.93	[0.17, 0.14]	0.91	0.98	[0.50, 0.39]
	25	0.59	0.93	[0.16, 0.14]	0.93	0.98	[0.50, 0.39]
	40	0.78	0.95	[0.15, 0.11]	0.93	0.99	[0.50, 0.43]
	60	0.80	0.96	[0.14, 0.10]	0.92	0.99	[0.45, 0.36]
	85	0.84	0.96	[0.14, 0.10]	0.93	0.99	[0.46, 0.32]
	100	0.78	0.98	[0.18, 0.22]	0.97	0.99	[0.54, 0.57]
	100 (*)	0.69	0.95	[0.12, 0.09]	0.91	0.98	[0.25, 0.18]




- **Strong reliability across the dataset:** The framework consistently achieves >95% pointwise coverage for most snapshots, showing that performance is stable and not driven by a few high-performing cases.
- **Localized failure modes:** Miscoverage is mainly concentrated in shock-dominated regions (up to ~19.3%), with overall well-calibrated uncertainty elsewhere across the wing surface.

Model performance

Test			Complementary test		
MAE	RMSE	R^2	MAE	RMSE	R^2
0.046	0.084	0.965	0.049	0.087	0.963

MSCP Coverage

	Normalized L_2 norm Coverage			L_∞ norm Coverage		
	Nominal	Pointwise	Band Width + Std	Nominal	Pointwise	Band Width + Std
Test	0.05	0.95	[0.34, 0.17]	0.41	0.99	[0.90, 0.46]
Complementary Test	0.003	0.95	[0.34, 0.17]	0.33	0.99	[0.90, 0.46]



- (a) MSCP performance for Wing database.
(b) Individual predictions outside uncertainty bands.

- **Objective:** Construct prediction bands with statistical guarantee: $\mathbb{P}(\mathbf{y} \in C_\delta(\mathbf{x})) \geq 1 - \delta$
- **Local normalization** (heteroscedastic scaling): $r_{i,j} = \frac{|\epsilon_{i,j}|}{s_j}$, $s_j = \text{Std}(\epsilon_{i,j})$, $\epsilon_{i,j} = y_{i,j} - M_j(x_i)$
- **Non-conformity score:** $S_i = \sqrt{\frac{1}{D} \sum_{j=1}^D r_{i,j}^2}$, $S_i = \max_j r_{i,j}$
(RMSR) (L_∞ norm)
- **Calibration** (quantile selection): $k_s = S_{(\lceil (n_{cal}+1)(1-\delta) \rceil)}$
- **Adaptative uncertainty band:** $C_\delta(\mathbf{x}_*) = [\mathbf{M}(\mathbf{x}_*) - \mathbf{R}, \mathbf{M}(\mathbf{x}_*) + \mathbf{R}]$, with $\mathbf{R} = k_s \mathbf{s}$

$$\tilde{u} = \begin{bmatrix} 0 \\ \blacksquare \\ 0 \\ 0 \\ 0 \\ \blacksquare \\ 0 \\ 0 \\ 0 \\ \blacksquare \\ 0 \\ 0 \\ \vdots \\ \vdots \\ \vdots \end{bmatrix}$$

$\blacksquare$ Cp from wind tunnel tests located at their position in the full Cp vector.

$$\tilde{u} = P^T u \approx P^T \sum_{k=1}^d \tilde{a}_k \Psi_k = \underbrace{P^T U_d}_X \tilde{a}$$

It is necessary to determine the coefficients $\tilde{a}$ that minimize the error in the approximation


$$\tilde{a} = \arg \min_a \|P^T U_d a - P^T u\|_2^2$$

Cp vector reconstructed:

$$u \approx U_d \tilde{a}$$

If $X = P^T U_d \in \mathbb{R}^{r \times d}$ has a full column rank, the minimization has as its solution:


$$M = 0.86, \alpha = -4^\circ$$



Ordinary Least Squares Estimator

$$\tilde{a} = (X^T X)^{-1} X^T \tilde{u}$$

$$M = 0.86, \alpha = -4^\circ$$



Ridge Estimator

$$\tilde{a} = (X^T X + \lambda I)^{-1} X^T \tilde{u}$$

GPR is a non-parametric Bayesian approach for inference.

We first assume a Gaussian process prior: $f(x) \sim GP(m(x), k(x, x'))$

The collection of training points and test points are joint multivariate Gaussian distributed:

$$\begin{bmatrix} y \\ f_* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mu \\ \mu_* \end{bmatrix}, \begin{bmatrix} K(X, X) + \sigma_n^2 I & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix} \right)$$

$$\begin{aligned} f^* | X, y, X^* &\sim \mathcal{N}(\bar{f}^*, \Sigma^*) \\ \bar{f}^* &= \mu^* + K(X^*, X)[K(X, X) + \sigma_n^2 I]^{-1}(y - \mu) \\ \Sigma^* &= K(X^*, X^*) - K(X^*, X)[K(X, X) + \sigma_n^2 I]^{-1}K(X, X^*) \end{aligned}$$

Hyperparameters Optimization

$$\log p(\mathbf{y}|\mathbf{X}) = -\frac{1}{2}\mathbf{y}^\top [K(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I}]^{-1}\mathbf{y} - \frac{1}{2}\log |K(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I}| - \frac{n}{2}\log(2\pi)$$


GPR KERNELS

RBF Kernel

$$k(x, x') = \sigma_k^2 \cdot \exp\left(-\frac{(x - x')^2}{2l^2}\right)$$

Kernel Matérn 3/2

$$k(x, x') = \sigma_k^2 \cdot \left(1 + \frac{\sqrt{3}\|x - x'\|_2}{l}\right) \cdot \exp\left(-\frac{\sqrt{3}\|x - x'\|_2}{l}\right)$$




Linear multi-fidelity model (Kennedy and O'Hagan)

$$f_t(\mathbf{x}) = \rho_t f_{t-1}(\mathbf{x}) + \delta_t(\mathbf{x})$$

We can express the model as a Gaussian Process:

$$\begin{bmatrix} f_{t-1} \\ f_t \end{bmatrix} \sim GP \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} K_{t-1} & \rho K_{t-1} \\ \rho K_{t-1} & \rho^2 K_{t-1} + K_{\delta_t} \end{bmatrix} \right)$$



Le Gratiet & Garnier (2014) developed an efficient scheme where the problem is decoupled into many GPR problems as fidelities have. This resolutions scheme is hierarchical.

$$\bar{\mathbf{f}}_{*t} = \rho_t \bar{\mathbf{f}}_{*t-1} + \mu_{\delta_t} + K(\mathbf{X}_*, \mathbf{X}_t) [K(\mathbf{X}_t, \mathbf{X}_t) + \sigma_{nt}^2 \mathbf{I}]^{-1} [\mathbf{y}_t - \rho_t \bar{\mathbf{f}}_{*t-1} - \mu_{\delta_t}]$$
$$\text{cov}(\mathbf{f}_{*t}) = \rho_t^2 \text{cov}(\mathbf{f}_{*t-1}) + K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X}_t) [K(\mathbf{X}_t, \mathbf{X}_t) + \sigma_{nt}^2 \mathbf{I}]^{-1} K(\mathbf{X}_t, \mathbf{X}_*)$$