5th Automotive CFD Prediction Workshop

Neural Surrogates for Automotive CFD:
Training Strategy, Subsampling, and Interpretability with RWR and Courant

Astrid Walle, Anuj Kumar, Josiah Bjorgaard, Nikolas Borrel Jensen, Niklas Heim, Matteo Salvador, Alexander Lavin

8–9 October 2026 | DAMA Technopole | Bologna, Italy

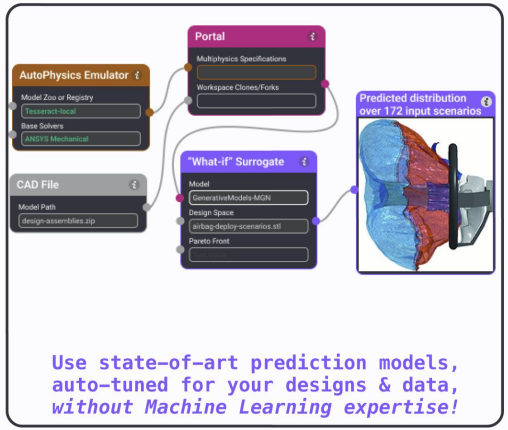
09/10/2026

Do not share or reproduce without express permission from Pasteur Labs or Institute for Simulation Intelligence authorized representatives.

What We Build at Pasteur Labs

Numerical Simulations and **Physics-AI** for industry-grade engineering workflows.

From **automated CAE data engines** to **Physics-AI models** and **fully differentiable optimization pipelines**.





Tesseract

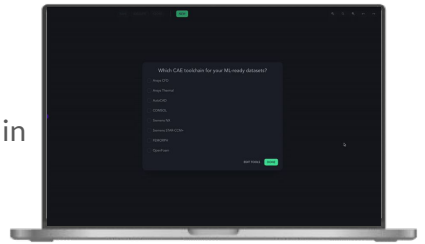
power

Testbeds: Automated CAE data engines

AutoPhysics: Pre-Conf. Components & Workflows Catalogue

runs in

Workspaces



PASTEUR LABS

DrivAerML

1. Our Flagship Models
2. Run Overview
3. Sequential vs. Combined vs. Separate Training
4. Data Scarcity
5. Downsampling Impact
6. Explainable Models
7. Just run everything in the SI Platform

Courant: Interpretable Neural Surrogates for CAE

- Neural surrogates speed up numerical simulations but behave as black boxes.
- Unlike FEM, engineers cannot easily inspect or validate what they are doing.
- Courant stays accurate while also being explainable *at scale*
    - By tying latent embeddings to spatial anchors and using an affine decoder
    - Every full-field prediction decomposes exactly into a local, adaptive, partition-of-unity-like basis.

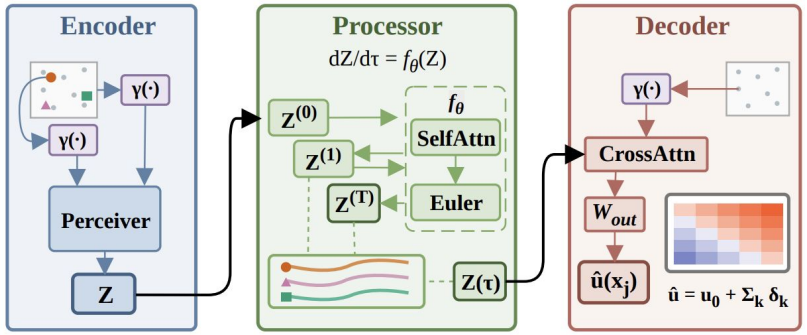
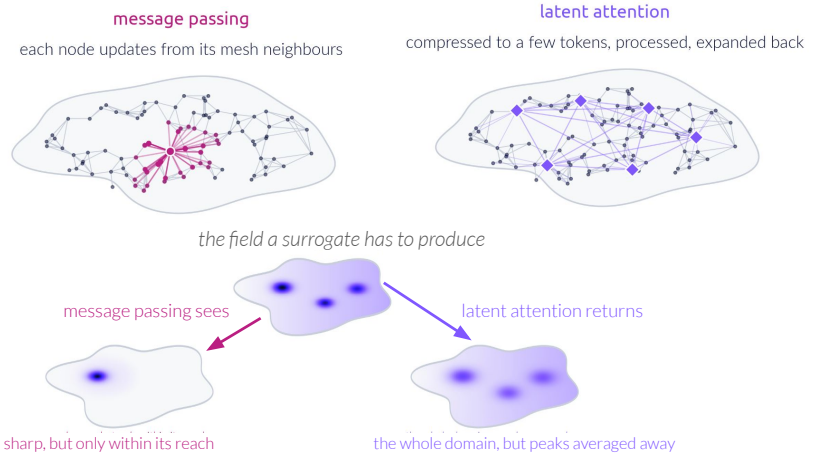
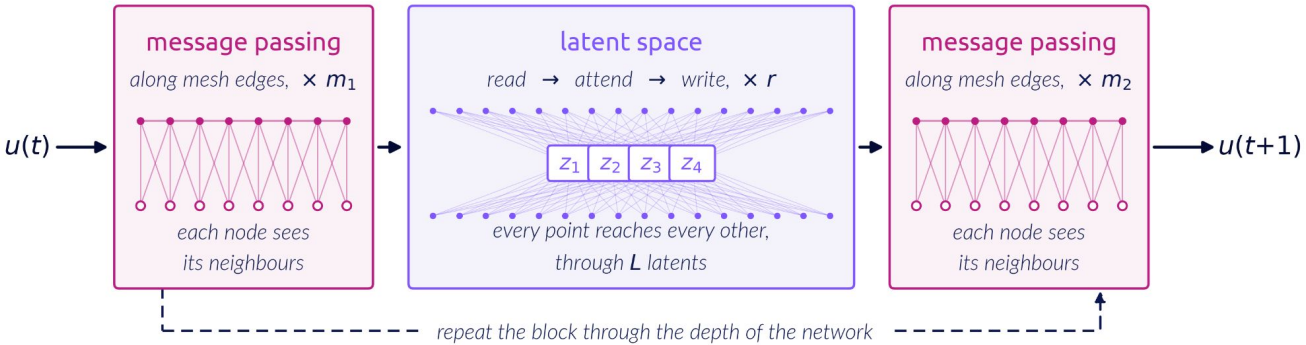


Figure 1 | Design of Courant across encoder, processor, and decoder. Selected point anchors are used to compress the input point cloud into latents using Perceiver layers. The processor evolves the latent anchors through a NeuralODE using self-attention for the right hand side operator. The decoder is a separate cross attention layer followed by a linear layer. At inference time, the latents can be decoded separately.

RWR: Neural PDE Surrogate with Local and Global Processing

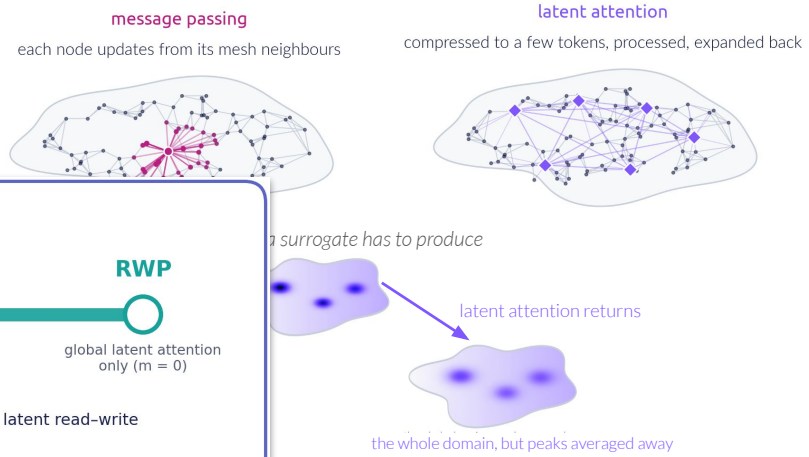
- Global surrogates (latent-token attention) act like a spatial low-pass filter and smooth over fine local detail.
- Local surrogates (mesh message-passing) capture fine details but cannot propagate information across large meshes.
- RWR interleaves both acting as the two halves of a multigrid cycle, one correcting low-frequency error, the other high-frequency error.





RWR: Neural PDE Surrogate with Local and Global Processing

- Global surrogates (latent-token attention) act like a spatial low-pass filter and smooth over fine local detail.
- Local surrogates (mesh message-passing) capture fine details but cannot propagate
meshes.
- RWR interleaves both acting
multigrid cycle, one correcti
other high-frequency error.



RWR: between MGN and RWP

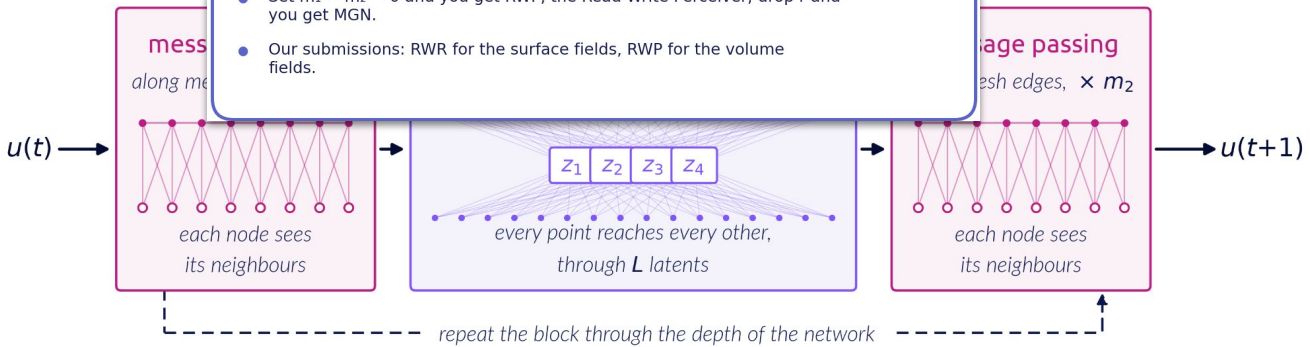
MGN — RWR — RWP

local message passing only (r = 0)

(m₁, r, m₂): both

global latent attention only (m = 0)

- RWR combines message-passing sweeps m (local) with latent read–write iterations r (global).
- Set $m_1 = m_2 = 0$ and you get RWP, the Read-Write Perceiver; drop r and you get MGN.
- Our submissions: RWR for the surface fields, RWP for the volume fields.

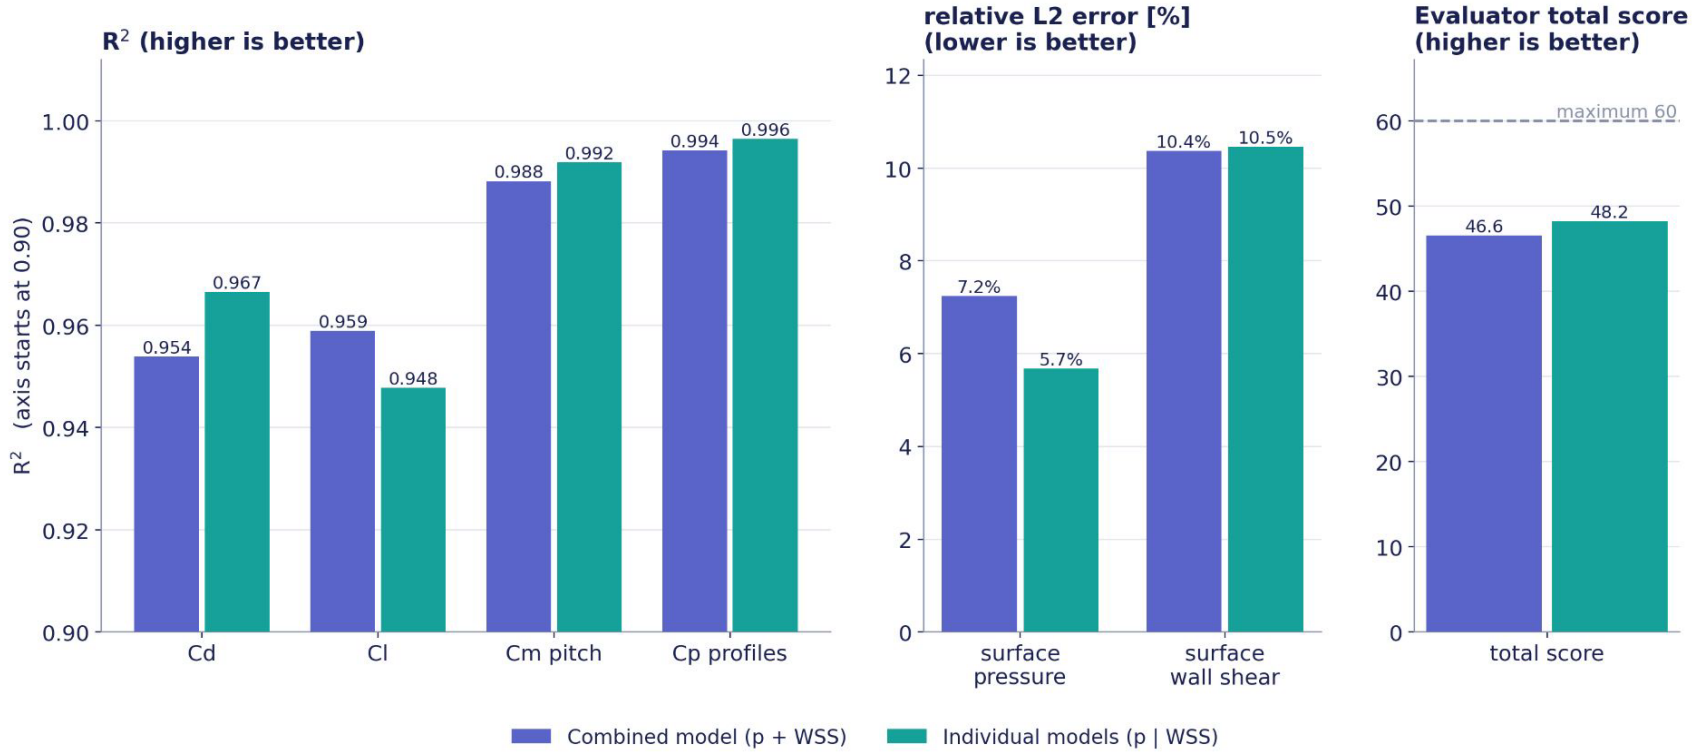


Run Overview

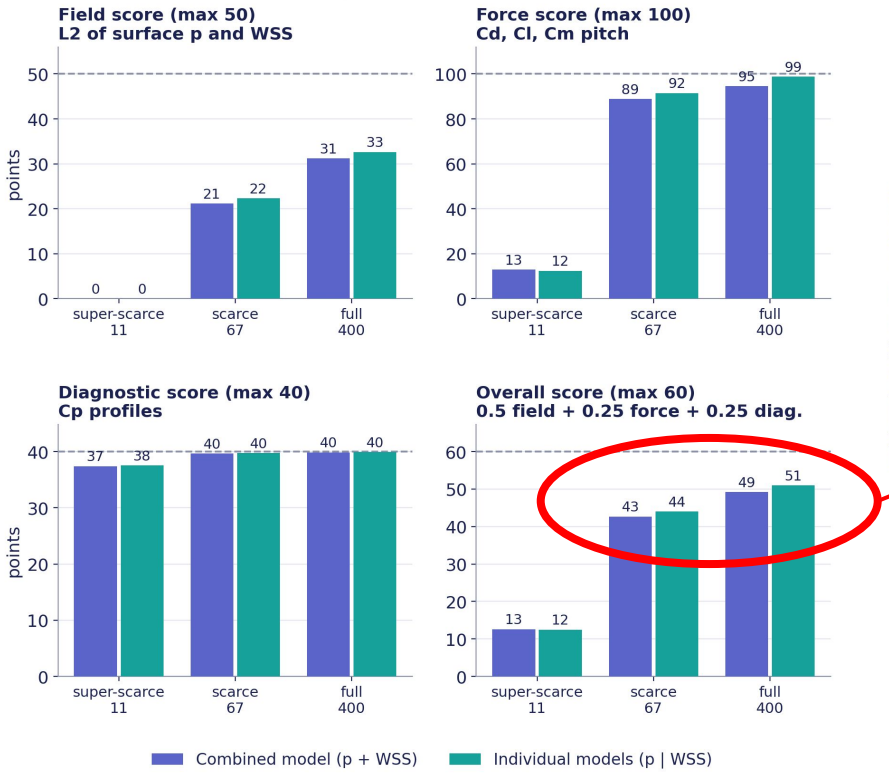
	Submission	Output	Split (train cases)	Strategy	Volume model	training sequence	scarcity	downsampling
training sequence	**combined_p_wss**	surface	full (400)	combined	–	●	●	
	individual_p_wss	surface	full (400)	separate	–	●	●	
	combined_u_p	surface + volume	full (400)	combined	RWP	●	●	
	individual_u_p	surface + volume	full (400)	separate	RWP	●	●	
	pin_u_p	surface + volume	full (400)	pin (p as input)	RWP	●	●	
	courant_combined_u_p	surface + volume	full (400)	combined	Courant	●	●	
	courant_individual_u_p	surface + volume	full (400)	separate	Courant	●	●	
	courant_pin_u_p	surface + volume	full (400)	pin (p as input)	Courant	●	●	
sample scarcity	**scarce_combined_p_wss**	surface	scarce (67)	combined	–		●	
	scarce_individual_p_wss	surface	scarce (67)	separate	–		●	
	scarce_combined_u_p	surface + volume	scarce (67)	combined	RWP		●	
	scarce_individual_u_p	surface + volume	scarce (67)	separate	RWP		●	
	super_scarce_combined_p_wss	surface	super-scarce (11)	combined	–		●	
	super_scarce_individual_p_wss	surface	super-scarce (11)	separate	–		●	
	super_scarce_combined_u_p	surface + volume	super-scarce (11)	combined	RWP		●	
	super_scarce_individual_u_p	surface + volume	super-scarce (11)	separate	RWP		●	
downsampling	**ds50**	surface, p only	full (400)	VTK 50 %	–			●
	ds50match	surface, p only	full (400)	full mesh	–			●
	ds90	surface, p only	full (400)	VTK 90 %	–			●
	ds90match	surface, p only	full (400)	full mesh	–			●

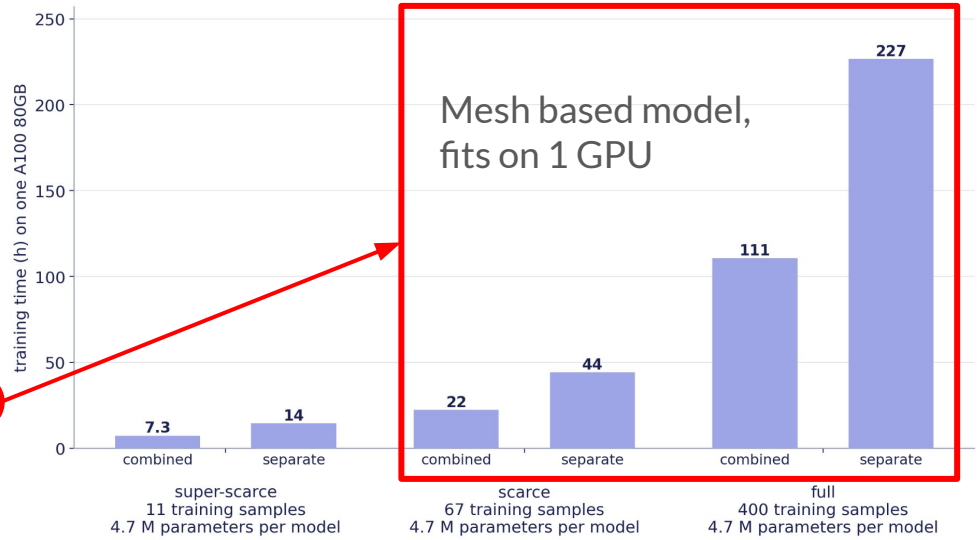
PASTEUR LABS

Surface - Full Split



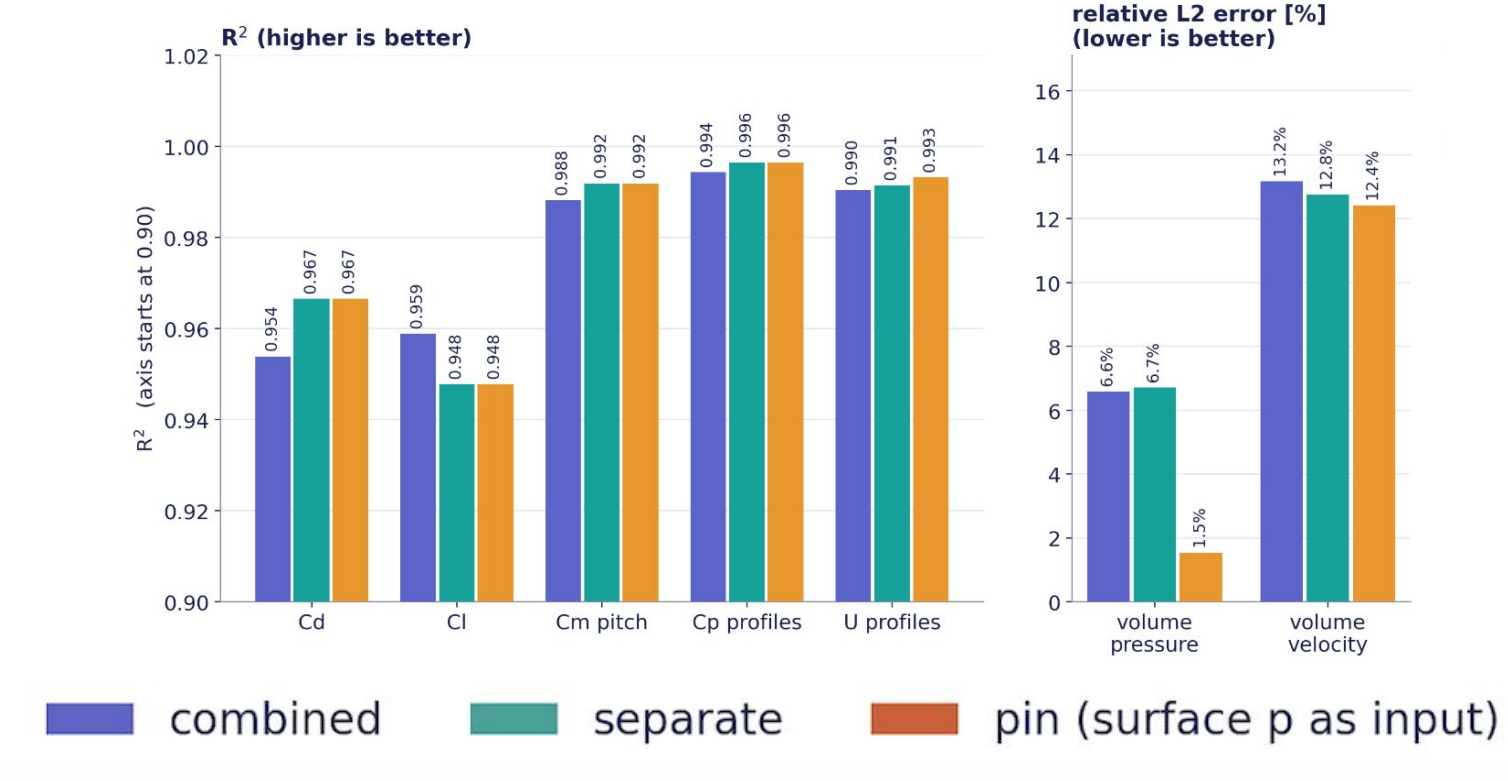
Surface - Scarce Data - Score Breakdown



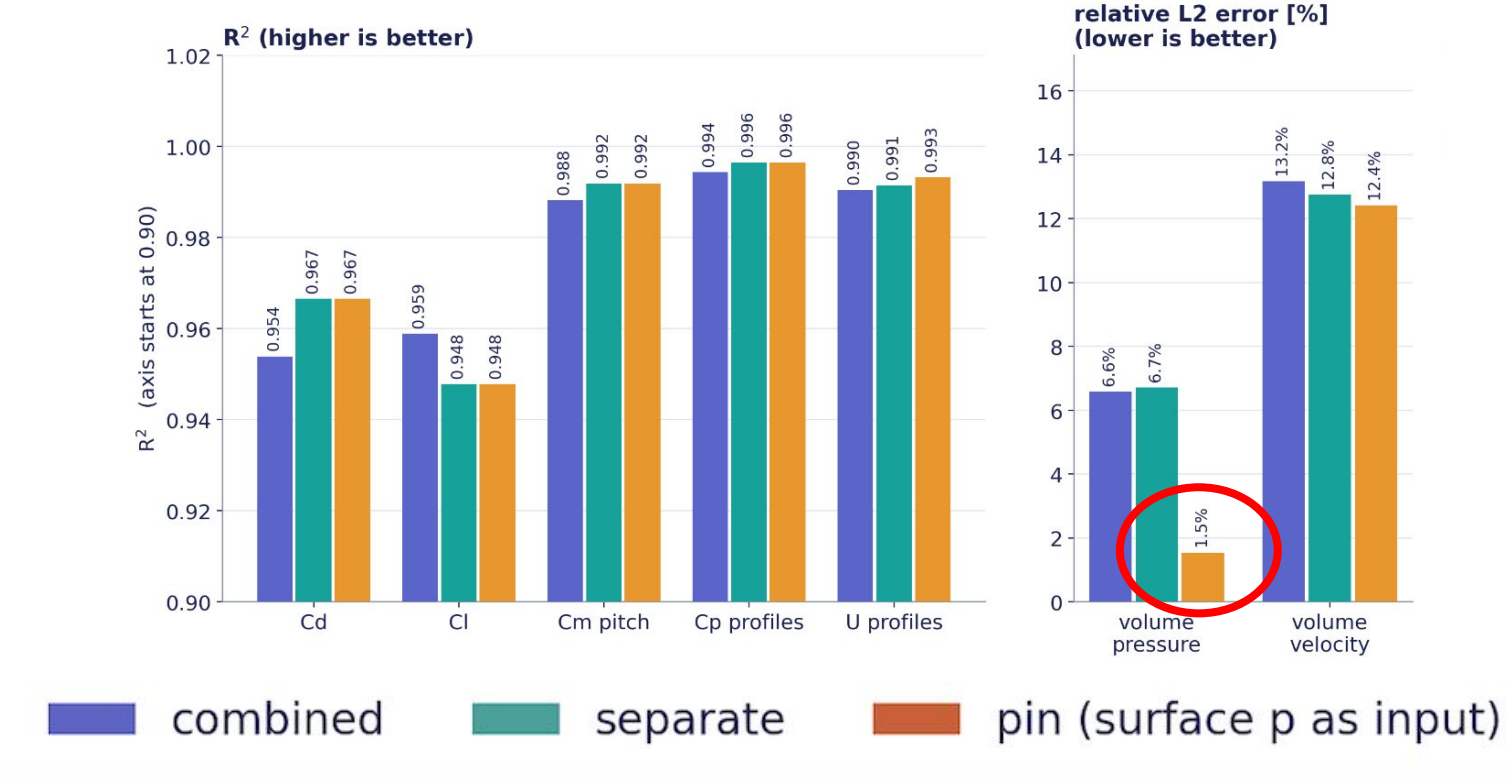


PASTEUR LABS

Submission Overview - Volume - Full Split



Submission Overview - Volume - Full Split



Training Strategy - Volume - Score Breakdown

Field: Volume pressure

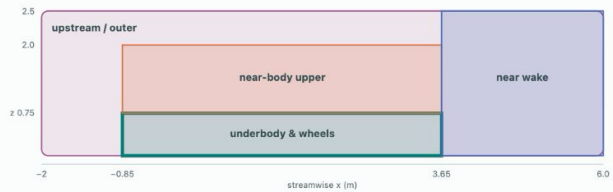
Regional error statistic: Typical case (median)

Zone details: Underbody & wheels envelope

3 compatible results shown on the exact volume pressure support. Regional values have zero official scoring weight.

Volume geometric zones · Volume pressure

Exact evaluator partition shown schematically



Underbody & wheels
−0.85 ≤ x < 3.65 m, |y| < 1.25 m, z < 0.75 m

Near-body upper
−0.85 ≤ x < 3.65 m, |y| < 1.25 m, 0.75 ≤ z < 2.0 m

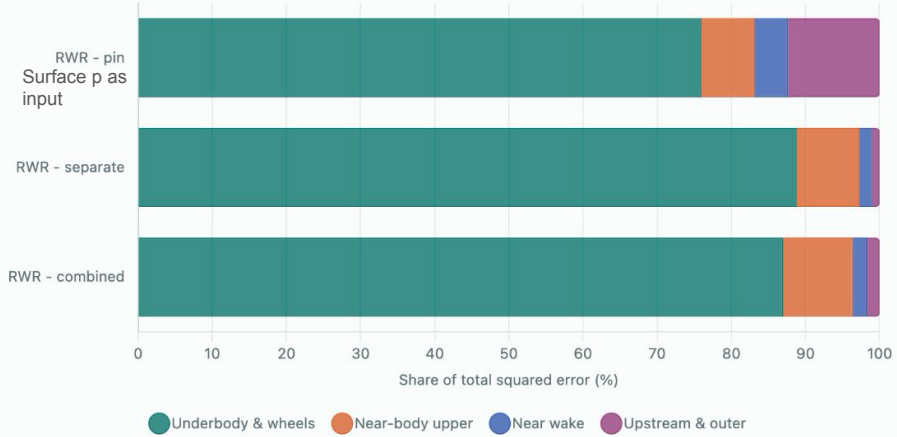
Near wake
3.65 ≤ x < 6.0 m, |y| < 2.0 m, z < 2.5 m

Upstream & outer
all remaining native volume cells

These are coarse Cartesian envelopes. The underbody-and-wheels zone is not an exclusive floor boundary-layer or wheel-patch mask.

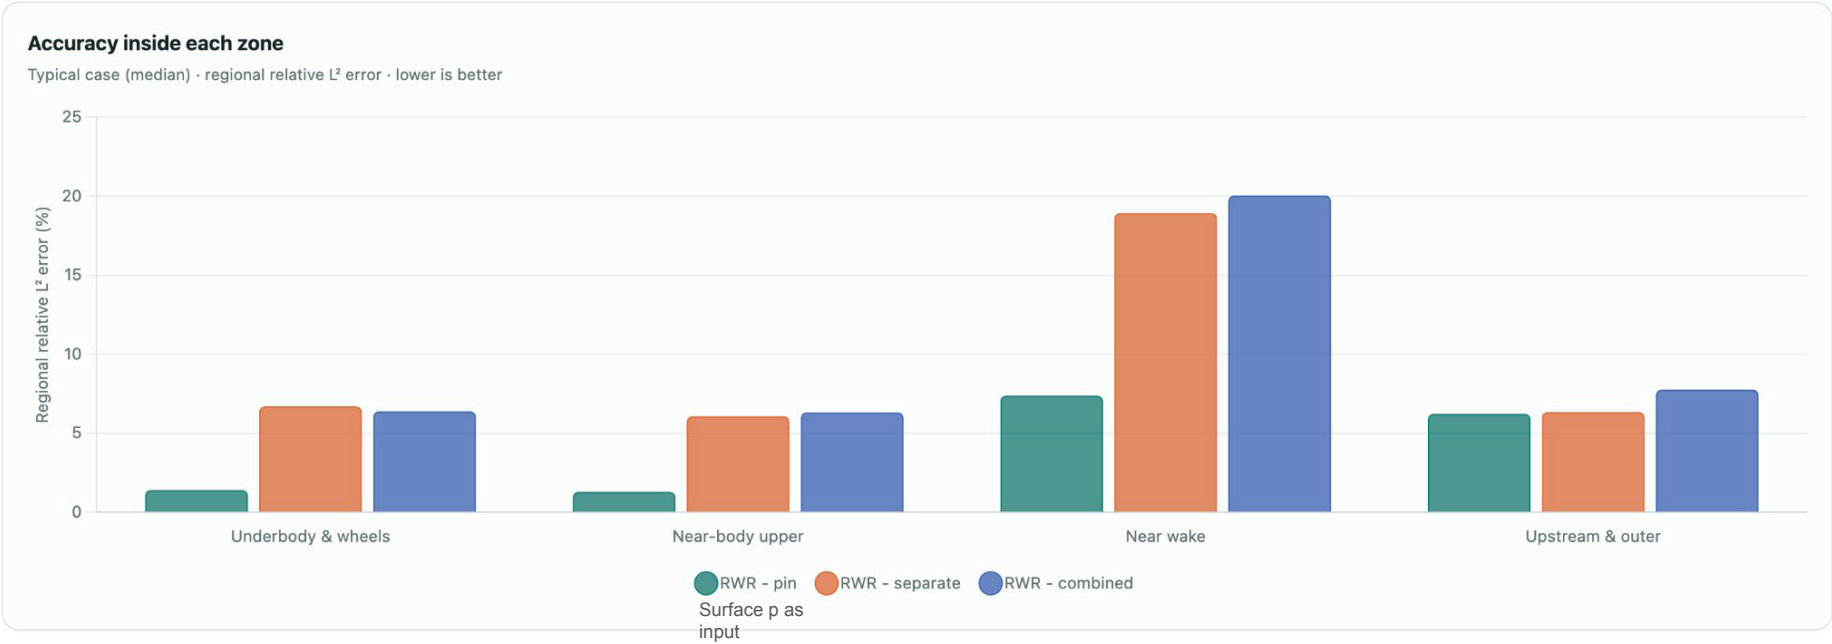
Where total error comes from

Share of the field's pooled squared error · each method totals 100%



PASTEUR LABS

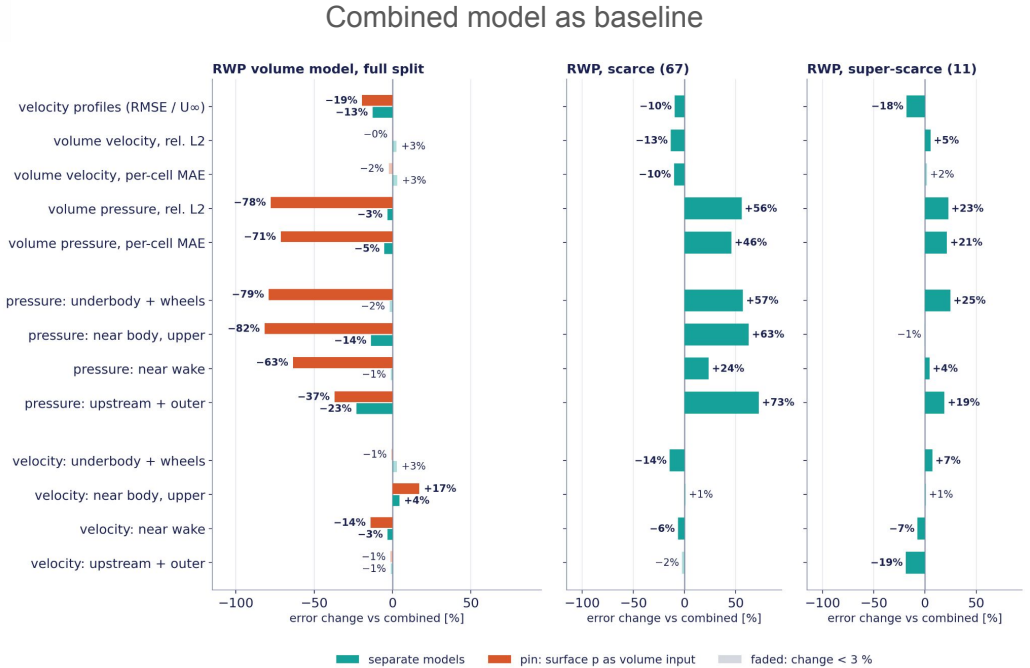
Training Strategy - Volume - Score Breakdown



Training Strategy

	full	scarce	super-scarce
surface models (RWR)			
Drag Cd	separate −56%	separate −33%	no clear winner
Lift Cl	separate −46%	combined −27%	separate −32%
Front lift Clf	no clear winner	separate −22%	no clear winner
Rear lift Clr	separate −46%	combined −46%	separate −34%
Pitch moment	no clear winner	combined −29%	no clear winner
Cp profiles	separate −17%	separate −11%	no clear winner
Surface pressure field	separate −14%	separate −6%	no clear winner
Surface wall shear field	no clear winner	no clear winner	separate −4%
p, low z, horiz. faces	separate −15%	separate −9%	separate −3%
p, low z, other faces	separate −12%	separate −3%	no clear winner
p, high z, horiz. faces	separate −27%	separate −12%	no clear winner
p, high z, other faces	separate −22%	separate −11%	separate −9%
volume models (RWP; pin only on the full split)			
Velocity profiles (RWP)	pin −8% vs separate	separate −10% vs combined	separate −18% vs combined
Volume velocity field (RWP)	no clear winner	separate −13% vs combined	combined −5% vs separate
Volume pressure field (RWP)	pin −77% vs separate	combined −36% vs separate	combined −18% vs separate
p underbody + wheels (RWP)	pin −79% vs separate	combined −36% vs separate	combined −20% vs separate
p near body, upper (RWP)	pin −79% vs separate	combined −39% vs separate	no clear winner
p near wake (RWP)	pin −63% vs separate	combined −19% vs separate	combined −4% vs separate
p upstream + outer (RWP)	pin −18% vs separate	combined −42% vs separate	combined −16% vs separate
U underbody + wheels (RWP)	no clear winner	separate −14% vs combined	combined −7% vs separate
U near body, upper (RWP)	combined −4% vs separate	no clear winner	no clear winner
U near wake (RWP)	pin −11% vs separate	separate −6% vs combined	separate −7% vs combined
U upstream + outer (RWP)	no clear winner	no clear winner	separate −19% vs combined

separate combined pin Surface p as input



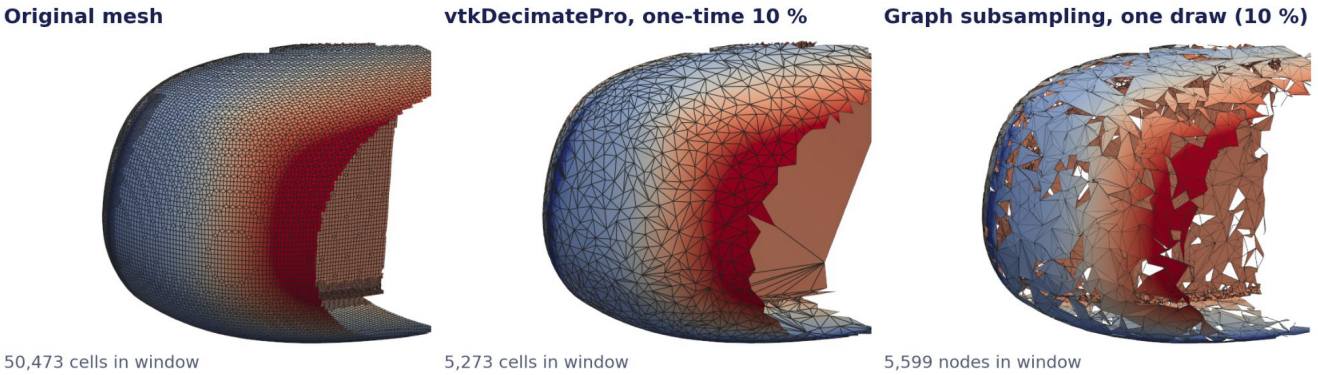
PASTEUR LABS

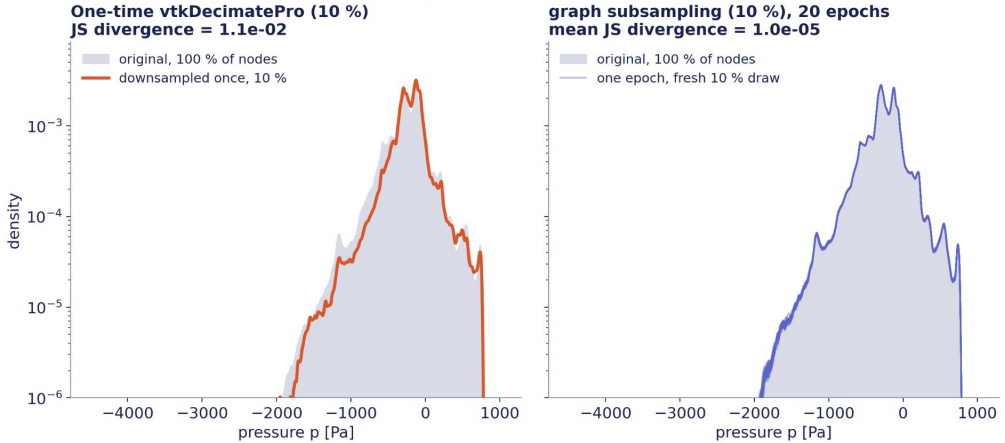
Downsampling - Why and What is the Impact?

Why: Fit the problem to the resources at hand

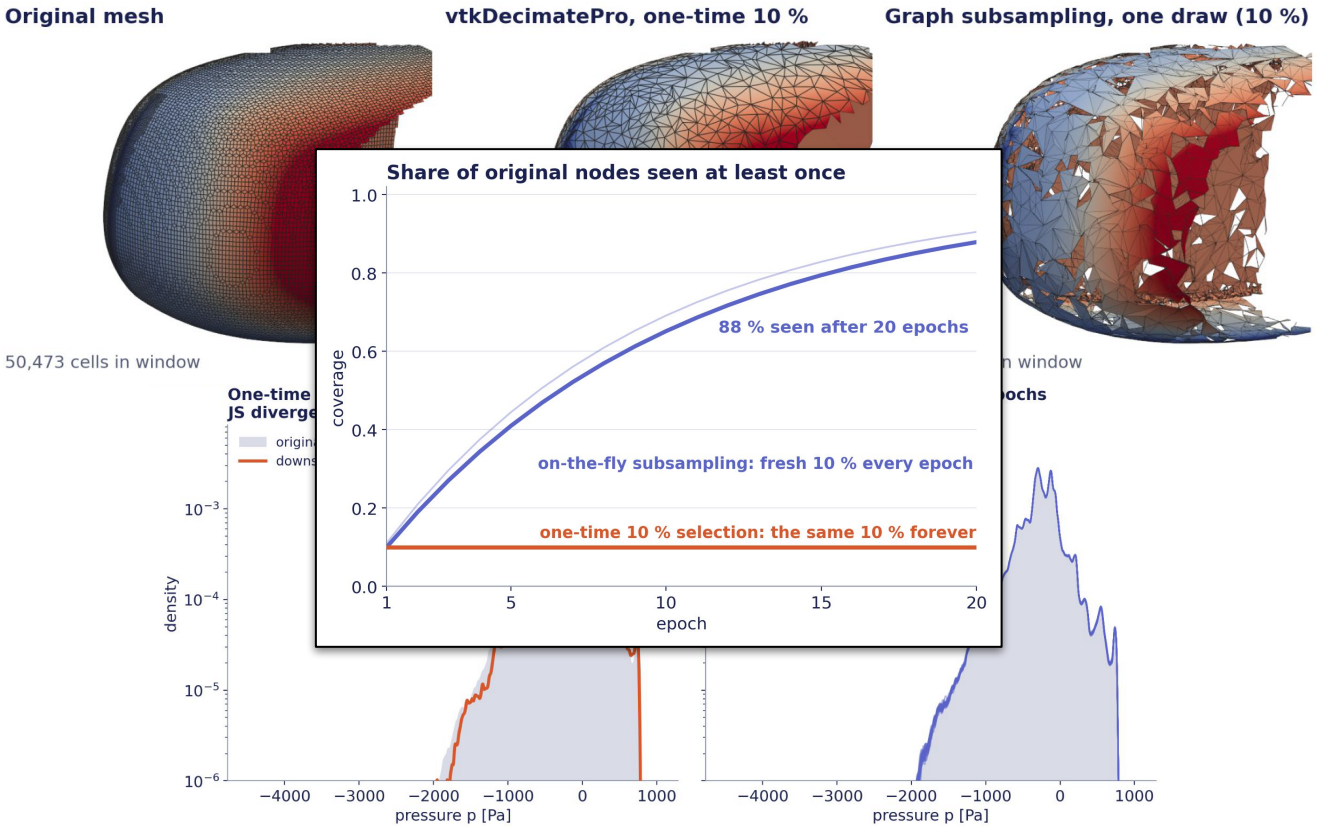
method	how it picks 10 % of the nodes	cost
ONE-TIME VTK REMESHING / DECIMATION		
vtkDecimatePro + sample ★	mesh decimation (topology kept), cell values averaged to points	109 s, once
vtkQuadricDecimation	quadric-error mesh decimation, point data interpolated	144 s, once
vtkQuadricClustering	uniform spatial bins, one representative per bin	39 s, once
ACVD remesh	isotropic remesh, fields re-interpolated onto it	69 s, once
ON-THE-FLY SUBSAMPLING, EVERY EPOCH		
random subsample	fresh random 10 % of the nodes each epoch (point cloud)	0.16 s / epoch
graph subsample ★	random node drop on the cell-adjacency graph, mesh kept	1.15 s / epoch

Downsampling - Impact

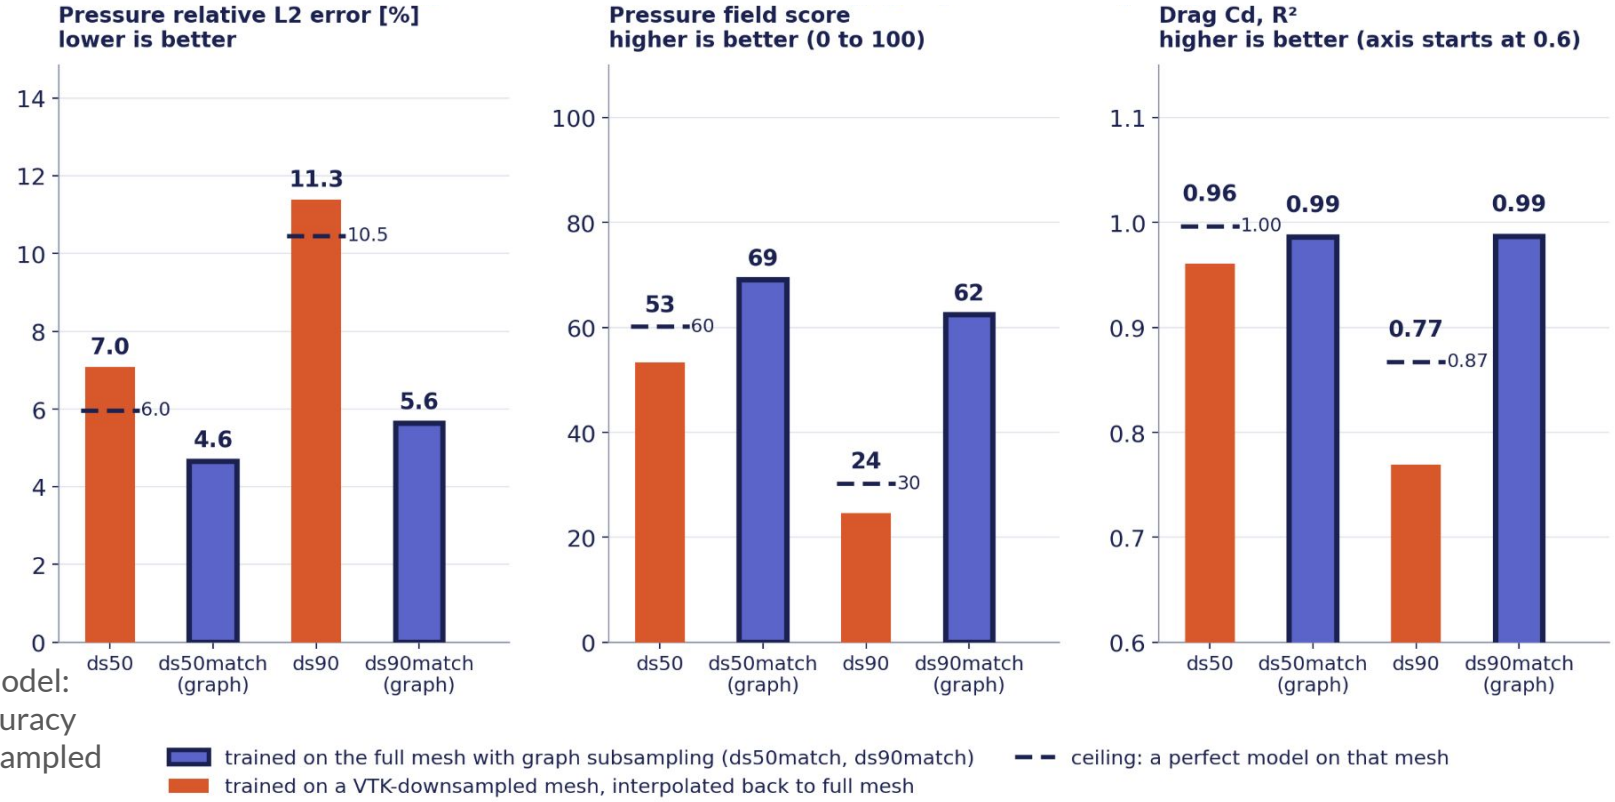




Downsampling - Impact



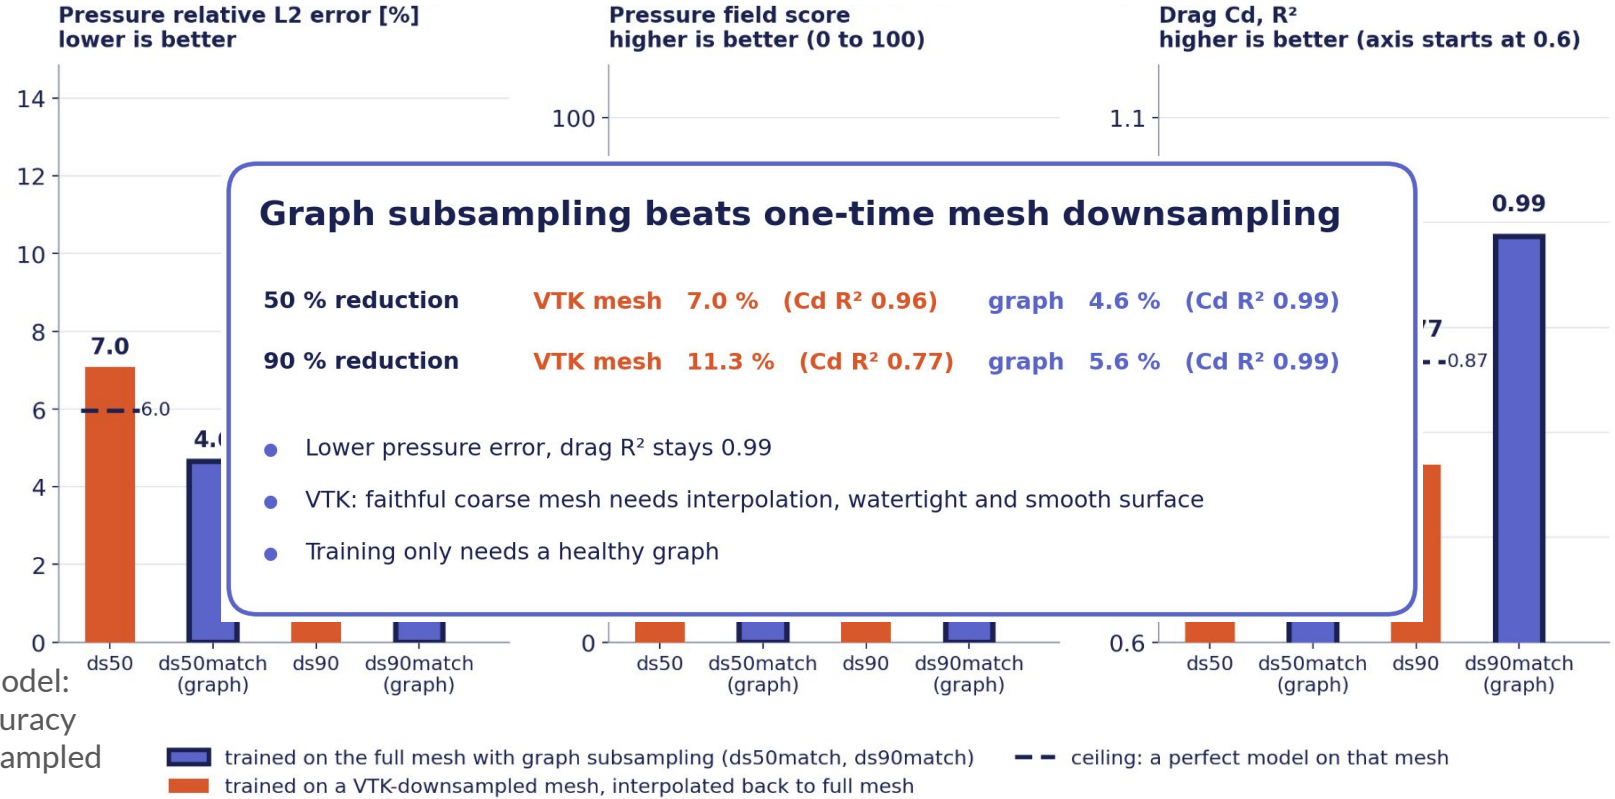
Downsampling - Comparison



Perfect model: 100% accuracy on downsampled mesh

PASTEUR LABS

Downsampling - Comparison

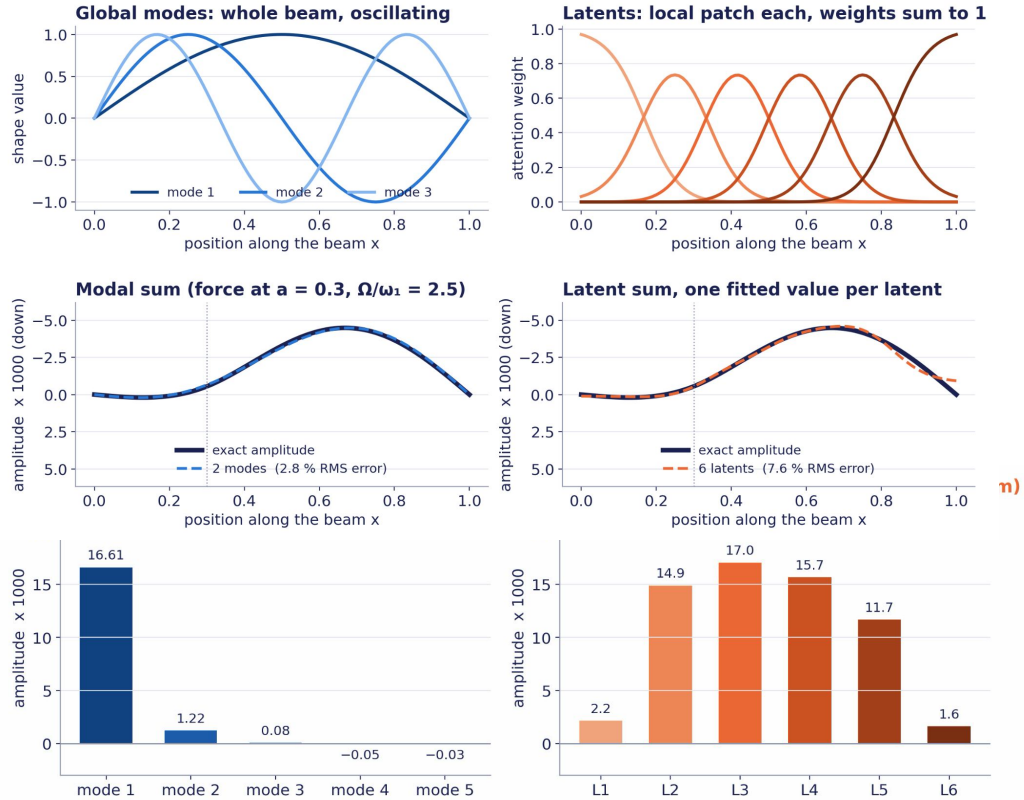


Perfect model: 100% accuracy on downsampled mesh

PASTEUR LABS

Explainable Models - Latent Decomposition - Analogy Harmonic Modes

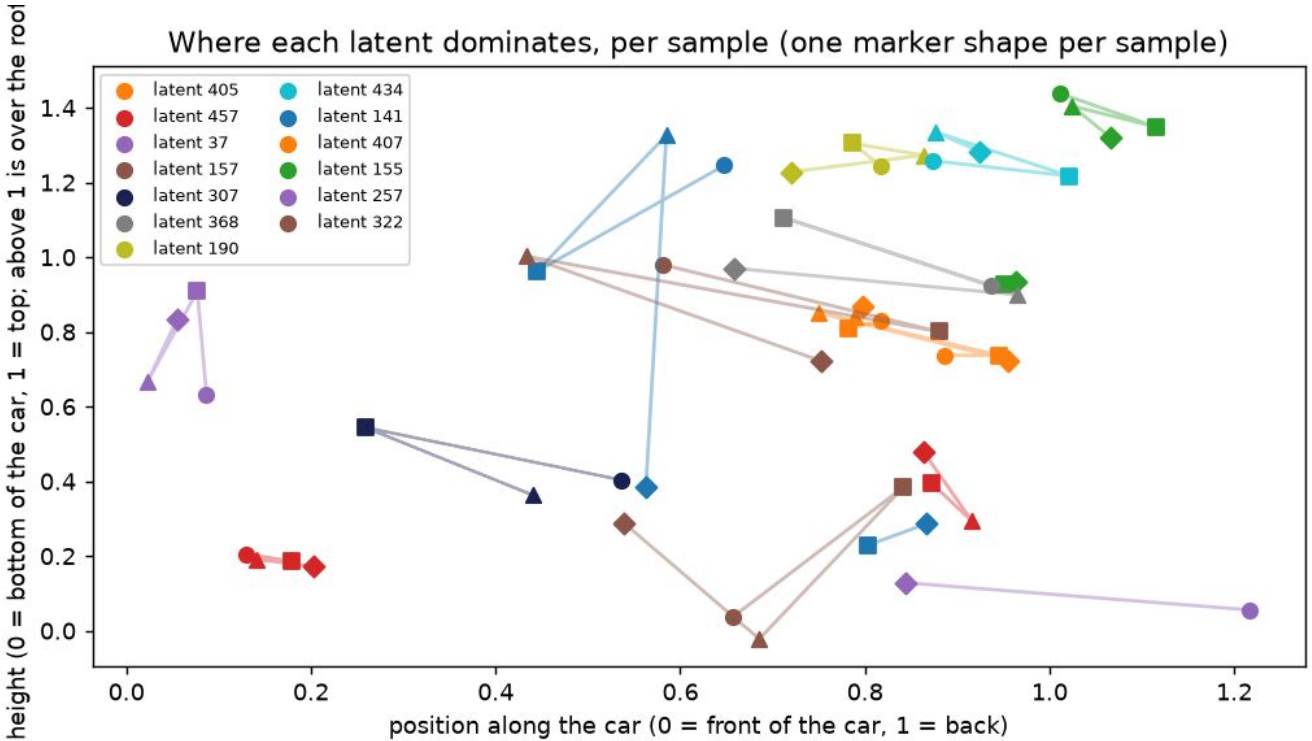
Harmonic Modes

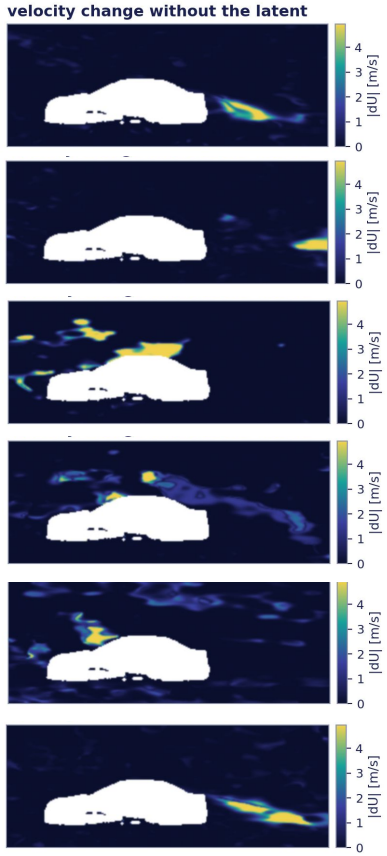


Latent Decomp.

PASTEUR LABS

Explainable Models

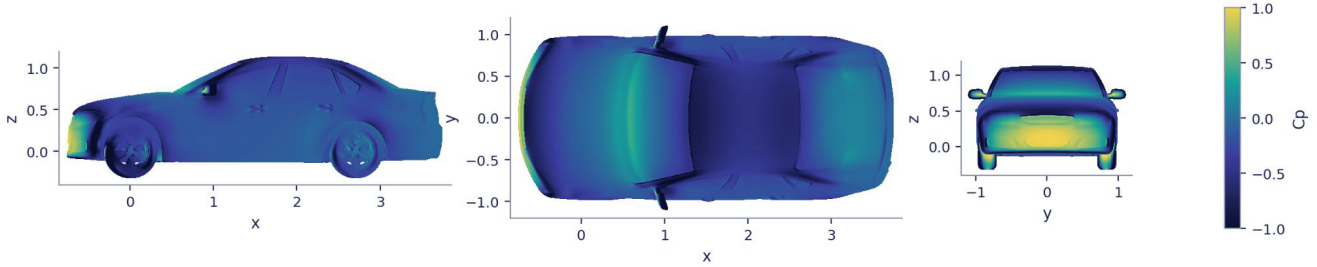




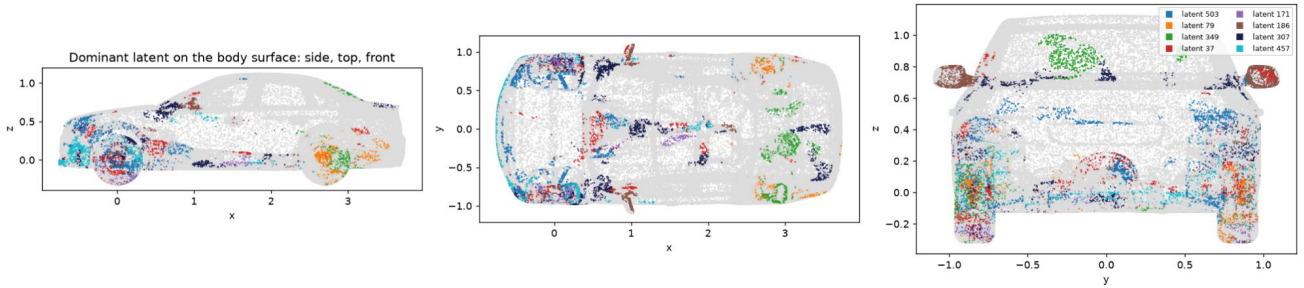
PASTEUR LABS

Explainable Models

Time-averaged Cp on the body surface



Dominant latent at each surface point (colours = latent ids)



Chained Inference Workflow in SI Platform

