Generative strategies to empower physics-based wave propagation with deep learning. Applications to earthquake engineering.

Filippo GATTI
Laboratoire de Mécanique Paris-Saclay UMR 9026

NAFEMS Teratec - ENSAM Paris, June 10, 2026

Acknowledgements

Thank you to my collaborators and mentors

Niccolo Perrone *Politecnico di Milano, MOX and LMPS, CentraleSupélec, Université Paris-Saclay*

Fanny Lehmann *ETH AI Center*

Hugo Gabrielidis *LMPS, CentraleSupélec, Université Paris-Saclay*

Mouad Bakhkhakh *LMPS, CentraleSupélec, Université Paris-Saclay & EDF R&D*

Michail Korres *EDF R&D*

Stefania Fresca *University of Washington*

Stéphane Vialle *LISN, Université Paris-Saclay, CentraleSupélec*

Fernando Lopez-Caballero *LMPS, CentraleSupélec, Université Paris-Saclay*

Didier Clouteau *LMPS, CentraleSupélec, Université Paris-Saclay*

Grants

NVIDIA 40000 GPU hours

IDRIS 110000 GPU hours

CINES 6700000 CPU hours

General context

SOURCE-TO-STRUCTURE ANALYSIS

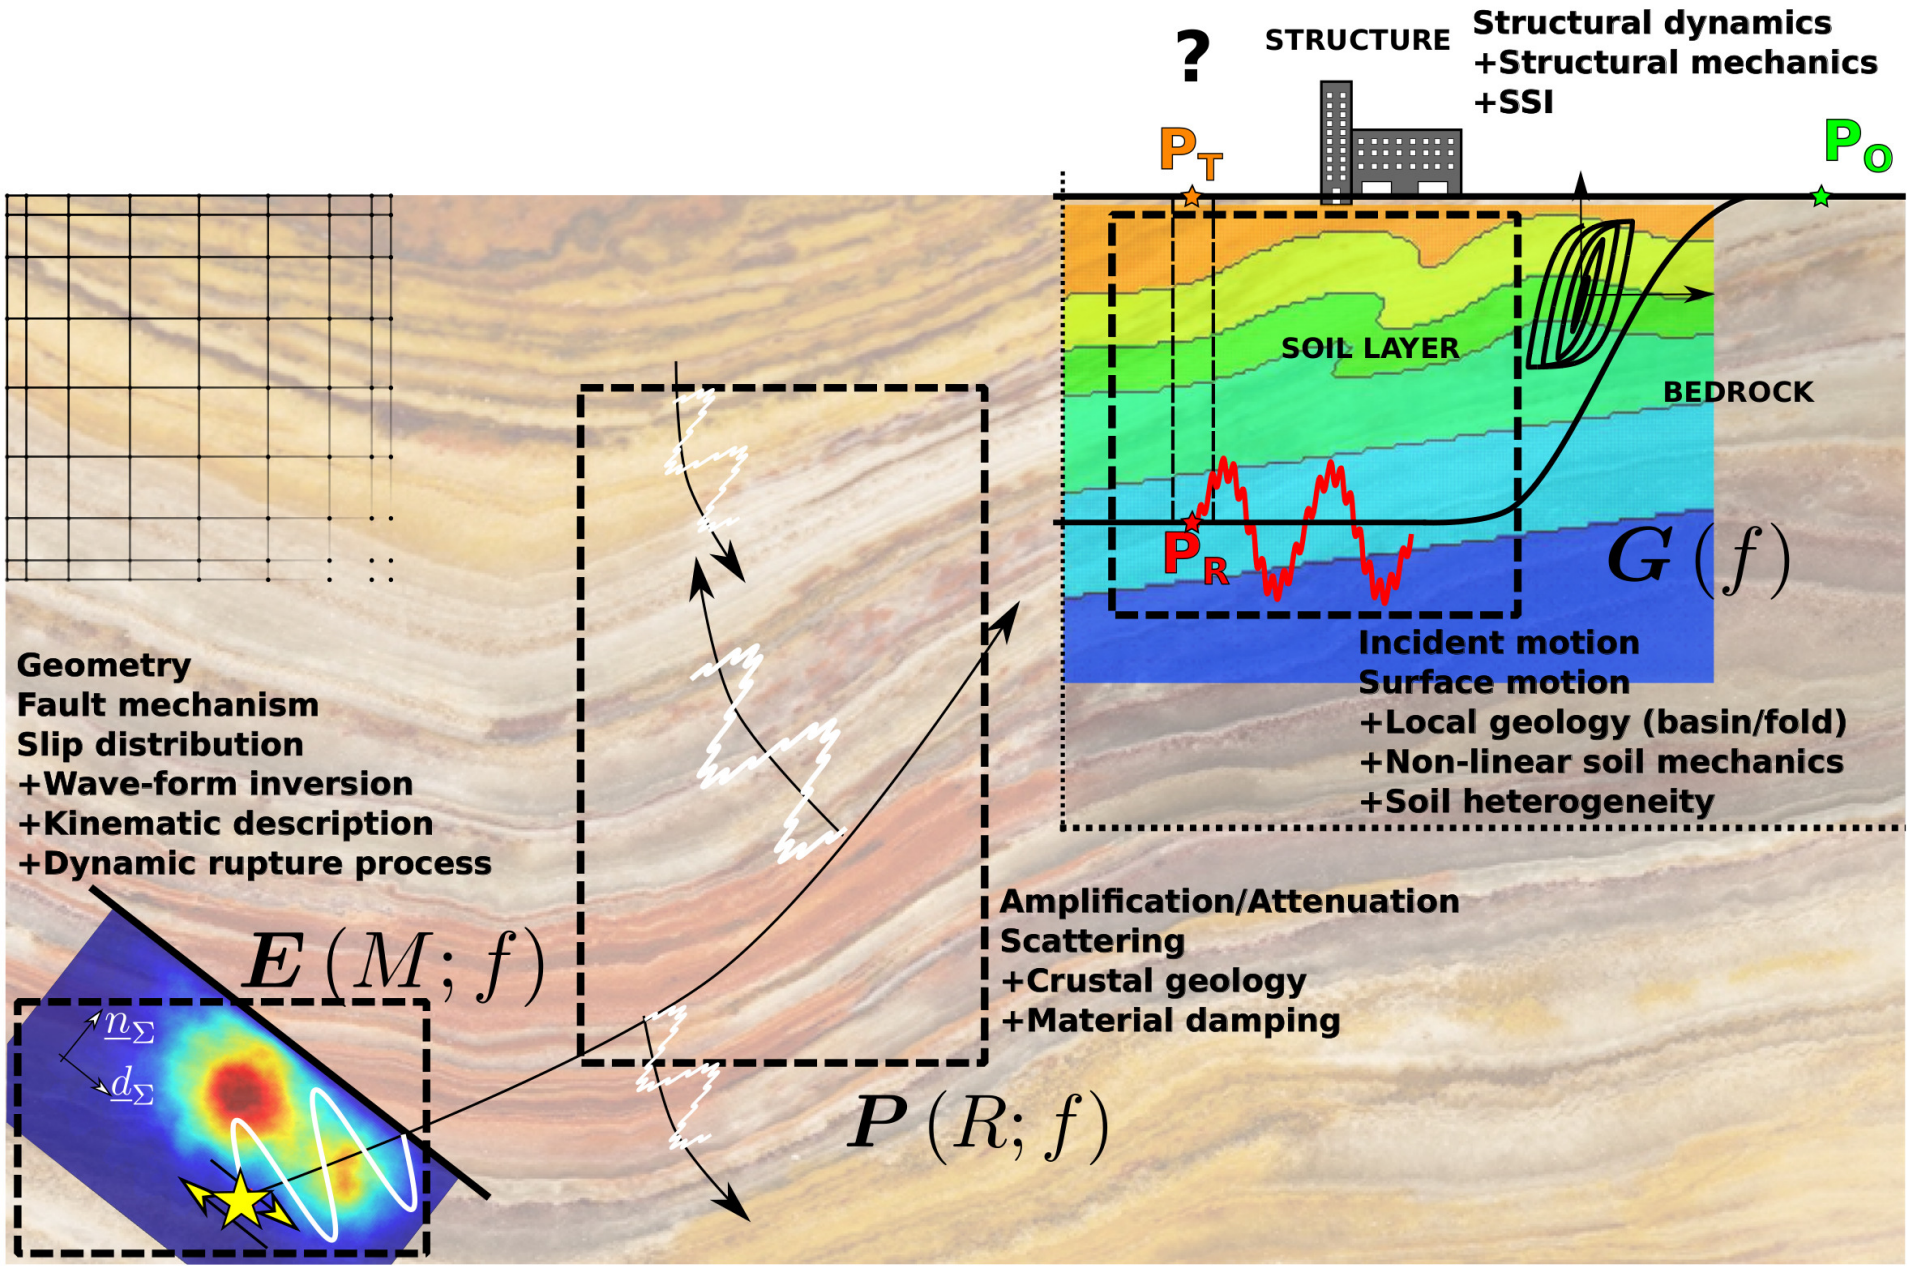

What we do

SEM3D [1, 6] : an open-source HPC code for earthquake simulation

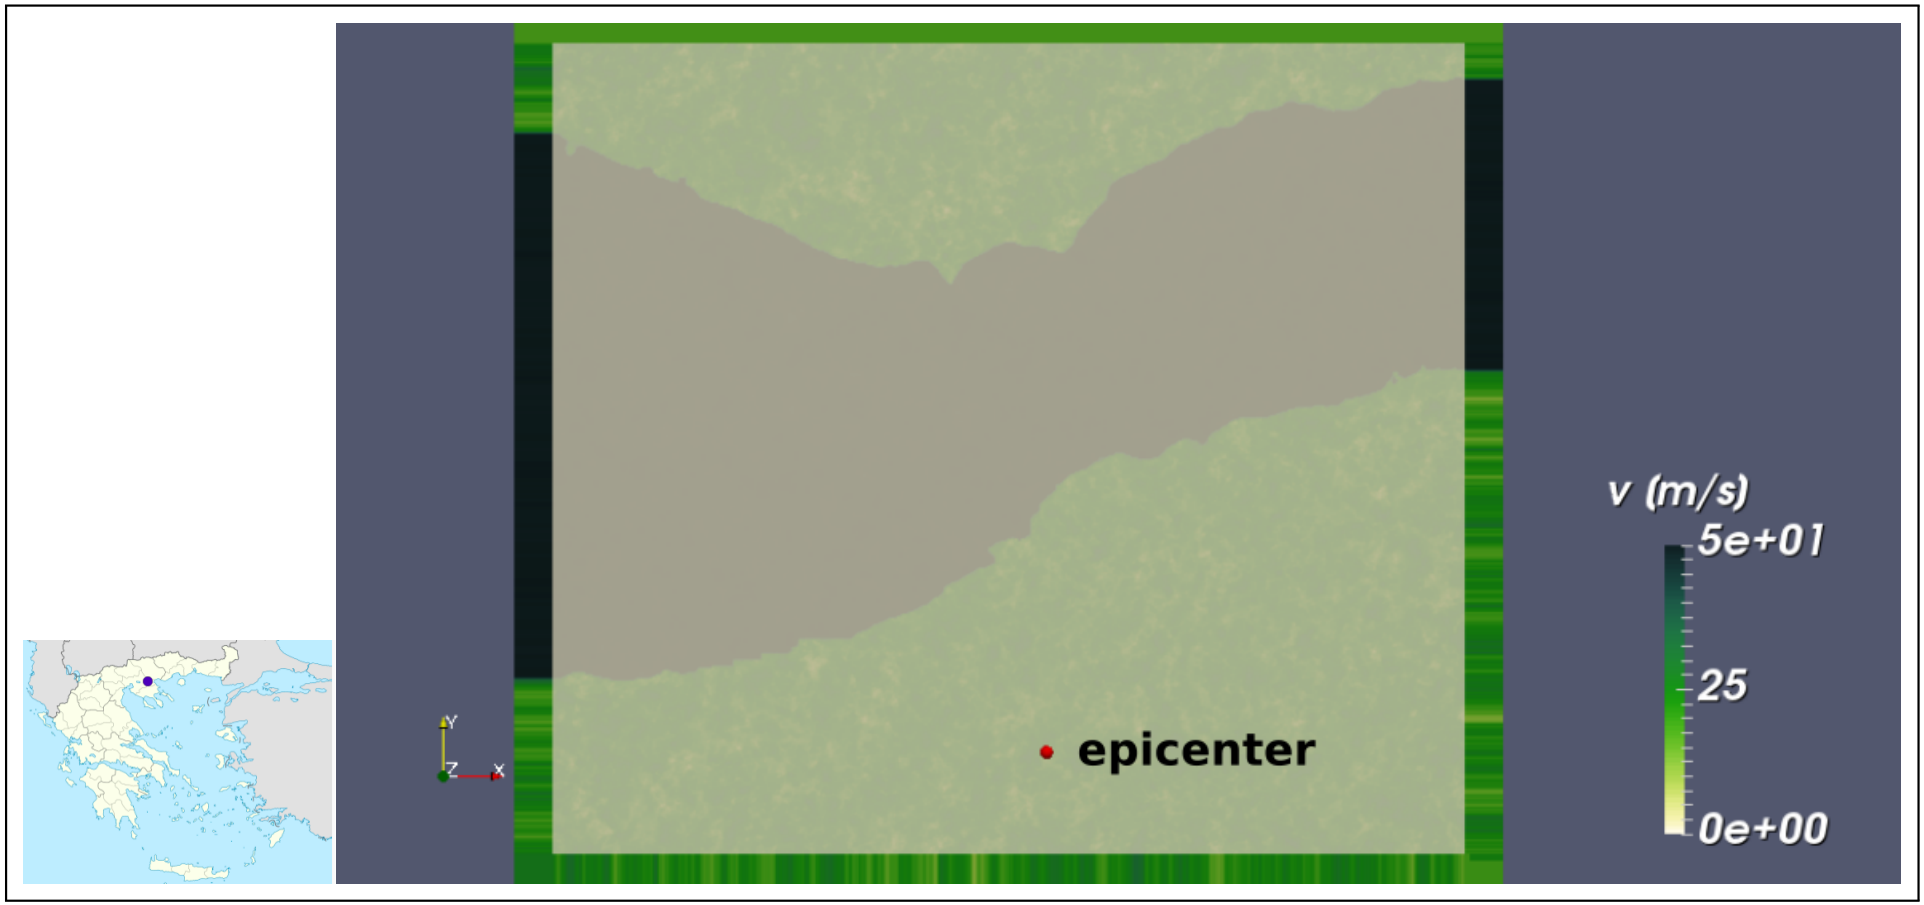

Figure – Earthquake simulation at the Volvi basin (Greece) [2]. 13 km × 11 km wide area. ≈7 h wall-time - 216 MPI cores @Ruche (≈62 days CPU time)

The earthquake engine : SEM3D

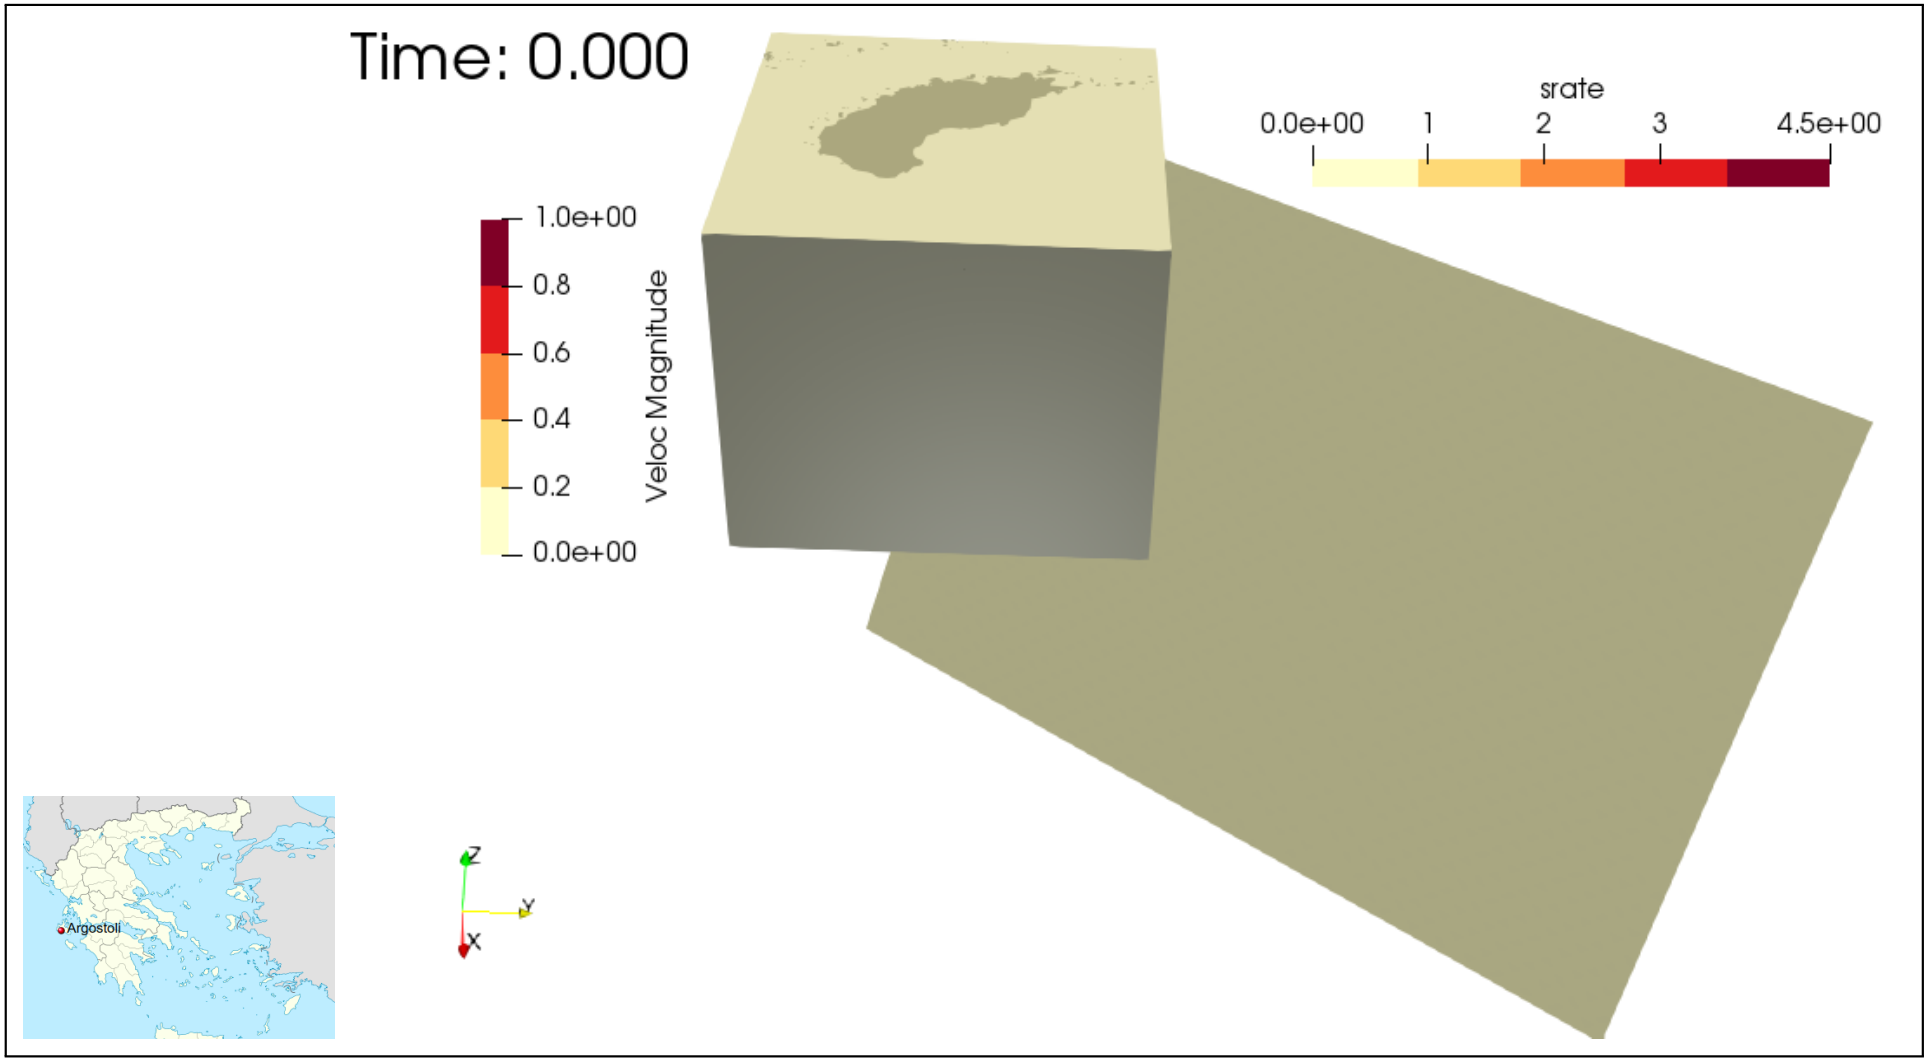

Figure – 1.3 $\cdot 10^{10}$ dofs 0-10 Hz Argostoli earthquake [6]. ≈20 h wall-time - 4000 MPI cores @Occigen (**3333 days CPU time**)

Earthquake simulation at several nuclear facilities

Dynamic SSI simulation of Unit 7 reactor building at Kashiwazaki-Kariwa NPP (Japan). Weak-coupling SEM3D/code_aster [3].

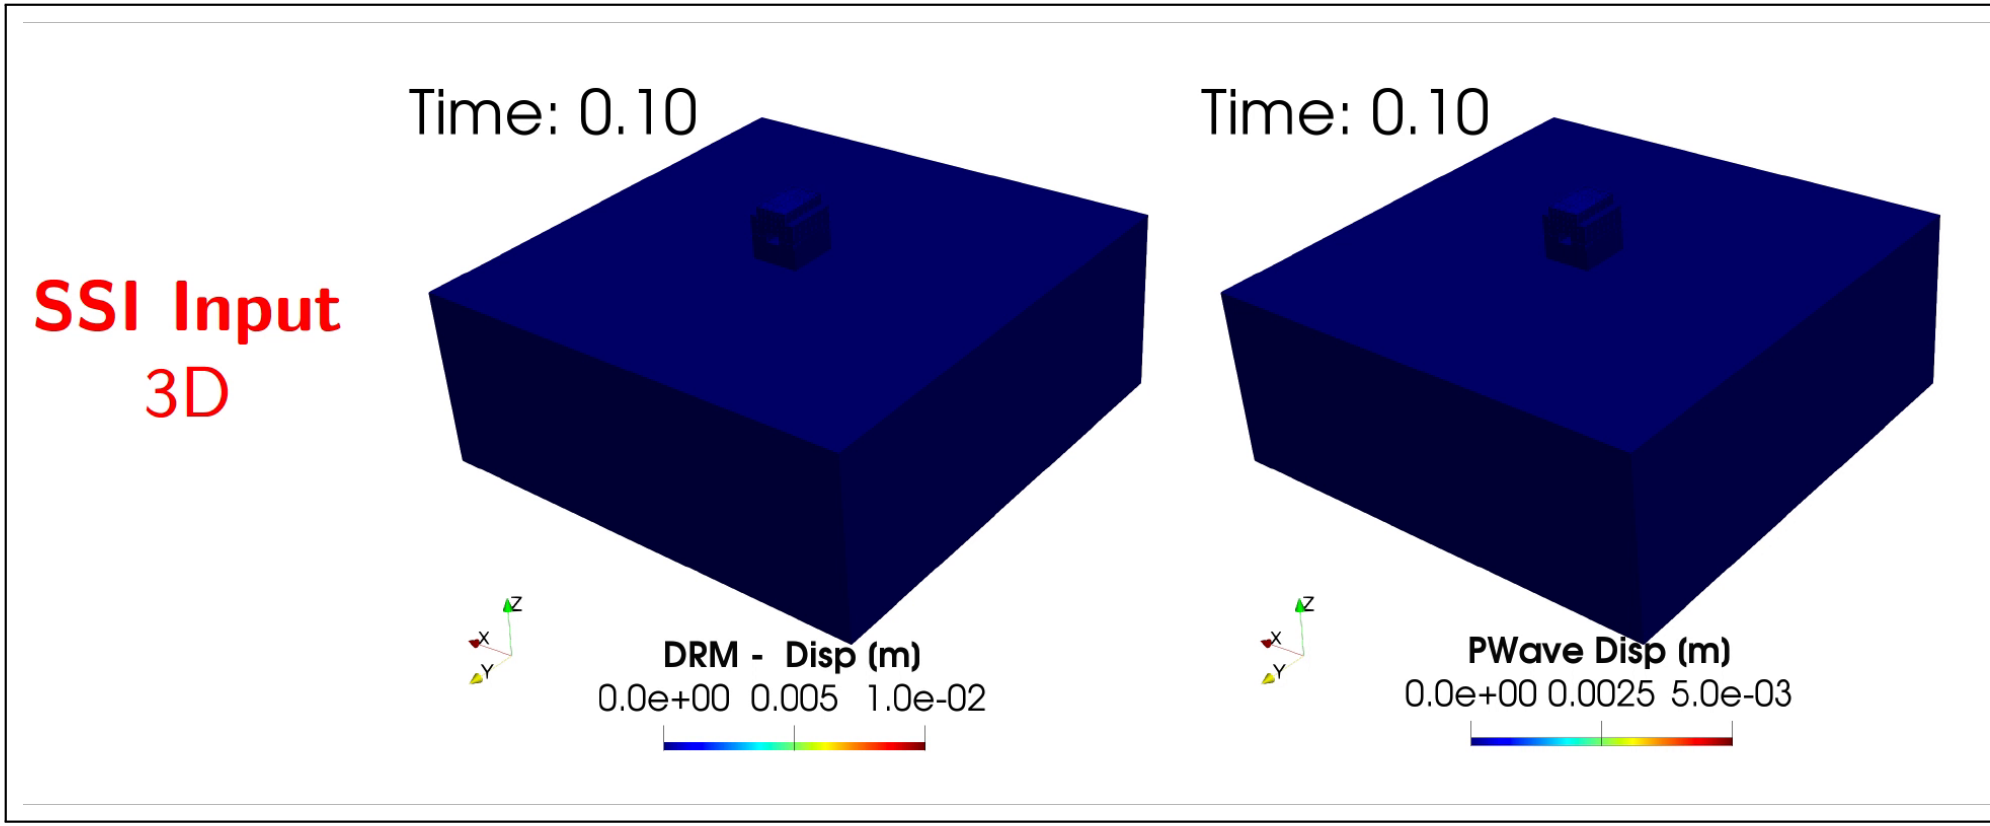

Figure – Courtesy of M. Korres EDF R&D.

Test numerical accuracy

All time-histories filtered at 7 Hz

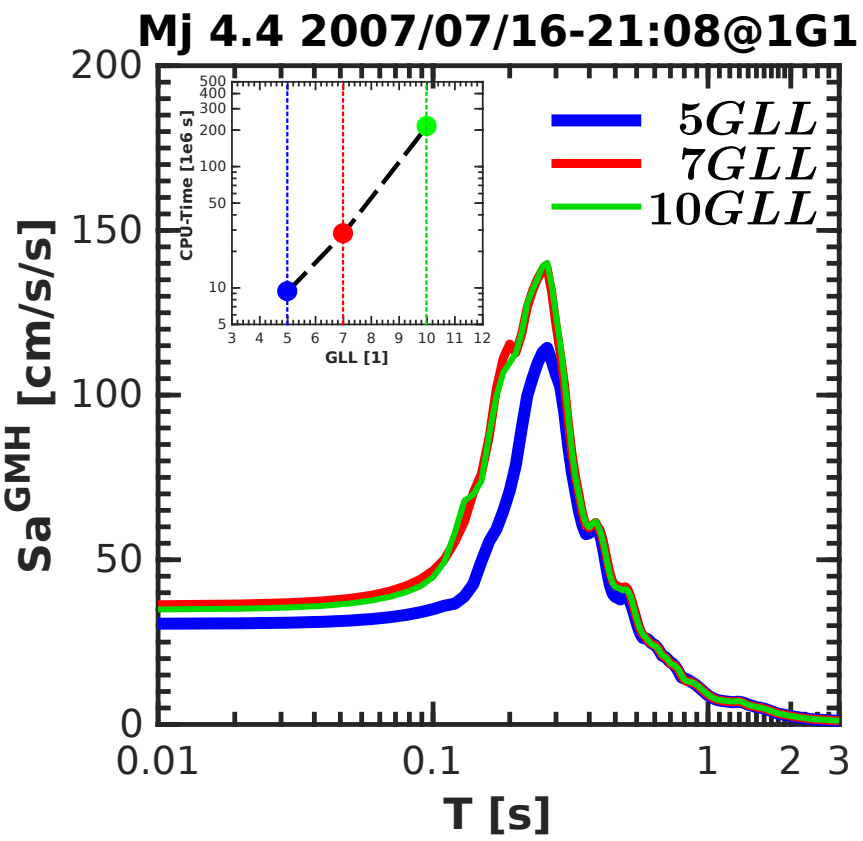

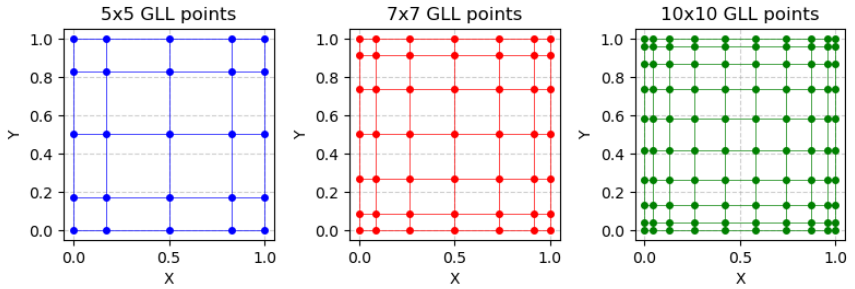

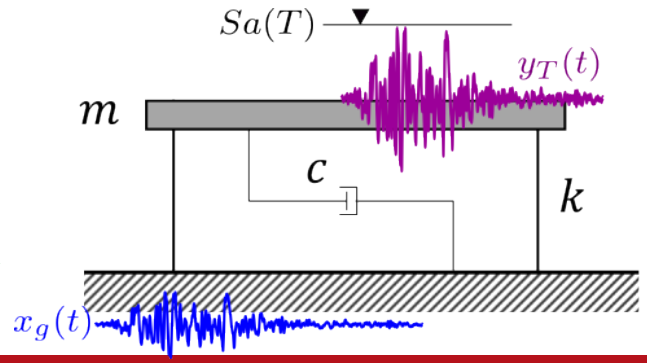

$f_{max}^{5GLL} \approx 3.5Hz$; $f_{max}^{7GLL} \approx 5.0Hz$; $f_{max}^{10GLL} \approx 7.0Hz$

Spinning-off : EASYRISK in a nutshell

SATT Paris-Saclay POC'UP 2023 funding

EASYRISK : an user-friendly platform for physics-based earthquake simulation and risk assessment at regional scale

- Pre-/Post-processing platform for SEM3D
- Use Machine Learning and BigData tools
- End-to-end seismic risk assessment solutions
- Simulate blind earthquake scenarios

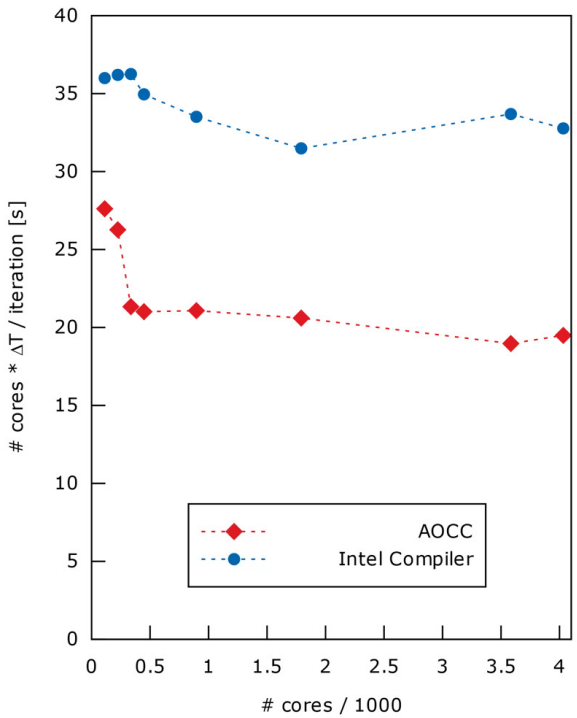

Make seismic RISK estimation EASY (on cloud)

Checkout blog article in collaboration with aneo

Earthquake Digital Twins:
The Power and Open Challenges of Physics-Based Simulation

WHAT PHYSICS-BASED SIMULATIONS (SEM3D) CAN DO

Source-to-Structure Analysis

SEM3D models the complete physics of an earthquake, including active fault rupture, wave scattering through crustal geology, and non-linear soil mechanics at the surface.

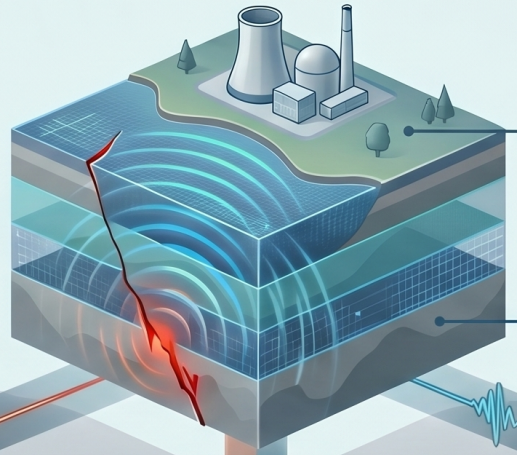

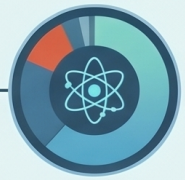

Billions of Degrees of Freedom

State-of-the-art simulations, such as the Argostoli earthquake model, involve up to 15 billion degrees of freedom to capture regional wave propagation.

Critical Infrastructure Protection

These simulations are used to assess the safety of nuclear facilities by modeling now reactors react to dynamic seismic loads in 3D.

THE REALITY OF RETRO-ENGINEERING PAST EARTHQUAKES

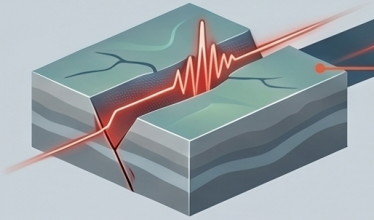

The 0–2 Hz Frequency Bottleneck

Current accurate retro-engineering is largely restricted to a limited frequency range (8-2 Hz or 8-18 Hz), which misses higher-frequency "vibrational fingerprints" critical for structural design.

Hard but Not Impossible

Reconstructing past earthouahee is achievable through high-fidelity digital hsms, but it is limited by significant rechnical constraints.

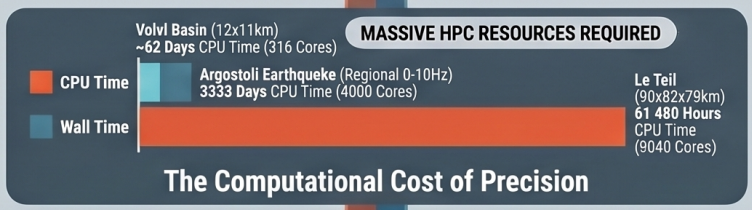

The Computational Cost of Precision

High Inherent Uncertainty

Models suffer from poor resolution compared to the intricacy of the Earth's crust, and accurate geological measurements (like shear wave velocity) are often unavailable.

THE OPEN CHALLENGES: LOOKING TO THE FUTURE

The Problem of Multiple Geological Models

Researchers must determine how to synthesize results and manage high variability when multiple conflicting geological interpretations of the same site exist.

Missing Observations

A major hurdle is how to build reliable physics-based models for regions where no historical earthquake recordings or borehole data are available.

Unlocking PSHA via Reduced Costs

To enable Probabilistic Seiemic Hazard Assessment (PSHA), we must reduce the "prohibitive" computational cost of running thousands of scenario simulations.

Real-Time Multi-Scenario Simulation

The ultimate goal is to move from months of supercomputing to real-time simulation of multiple scenarios, likely through AI surrogate models like Diffusion Transformers (DiT1D).

LMPS Laboratoire de Mécanique Paris-Saclay
CentraleSupélec
université PARIS-SACLAY

M_W4.9 Le Teil earthquake (2019) [4]

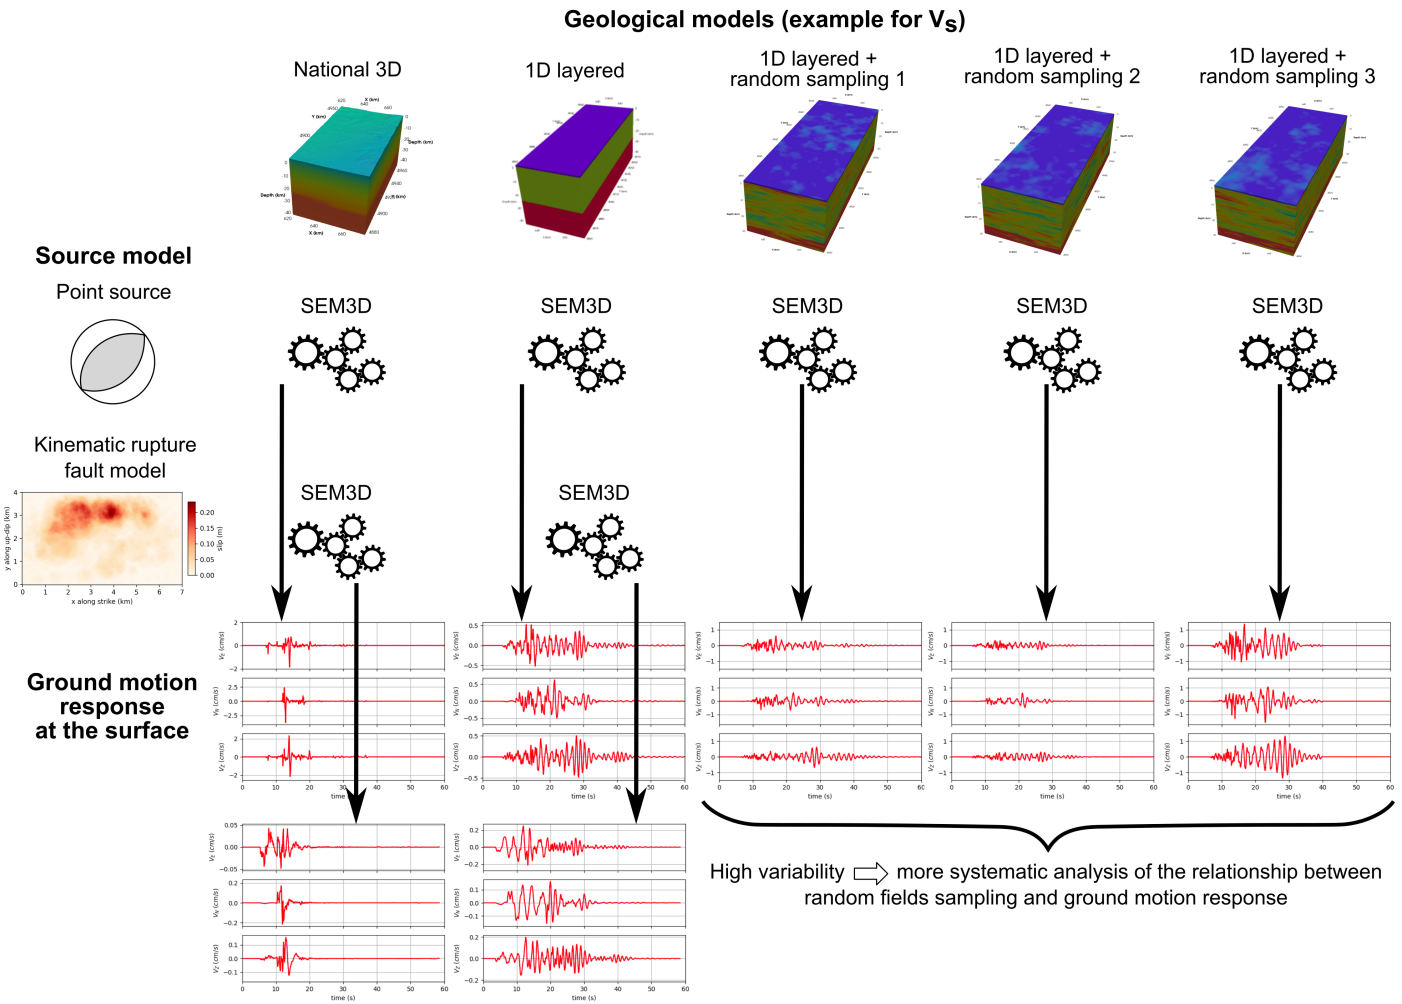

Figure – 80 km × 92 km × 79 km - 18.3 million hexahedra. 60 s of simulated signal 61440h CPU time - 2048 MPI cores AMD Milan @2.45GHz (AVX2) on TGCC.

M_W4.9 Le Teil earthquake (2019) [4]

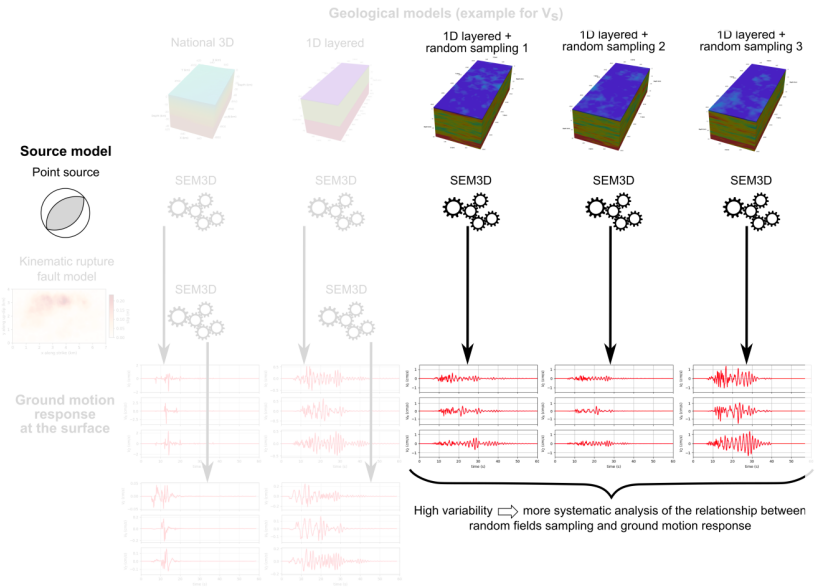

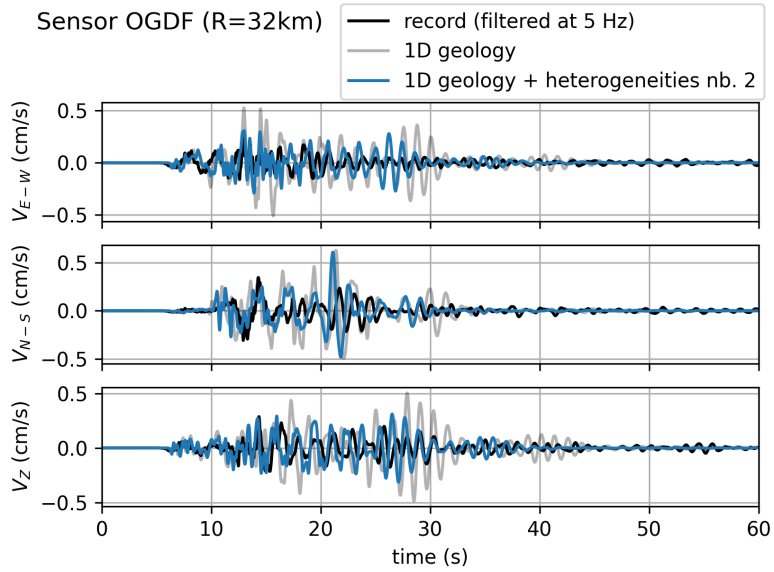

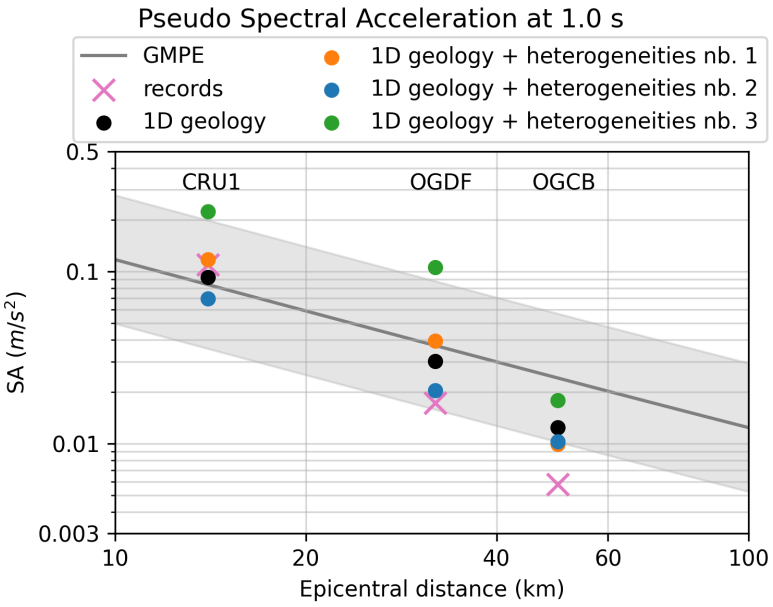

- Large geology uncertainty (small wavelengths)
- Impact on uncertain seismic response
- 40.000 heterogeneous geologies (9.6 km $\times$ 9.6 km $\times$ 9.6 km with 32.768 samples each)
  2621440000 parameters $\rightarrow$ Curse of dimensionality !
- What are the *eigen-geologies* ?
- Seismic signature of each *eigen-geology* ?

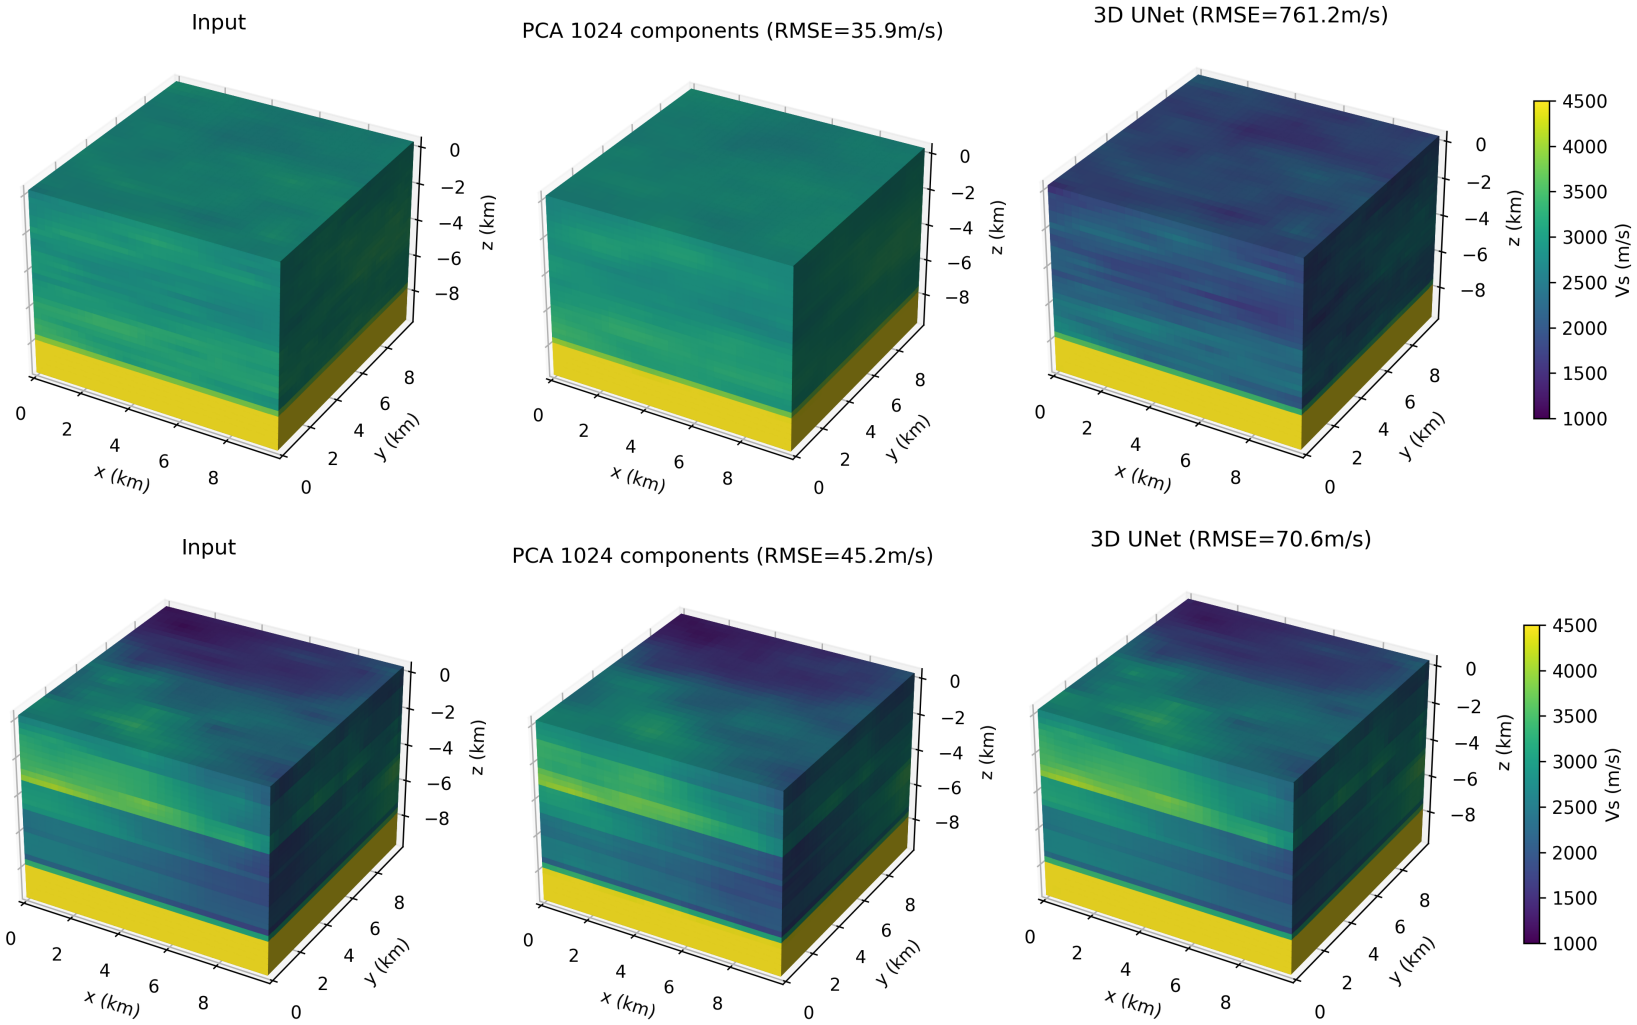
Input
PCA 1024 components (RMSE=35.9m/s)
3D UNet (RMSE=761.2m/s)
Input
PCA 1024 components (RMSE=45.2m/s)
3D UNet (RMSE=70.6m/s)
x (km)
y (km)
z (km)
Vs (m/s)

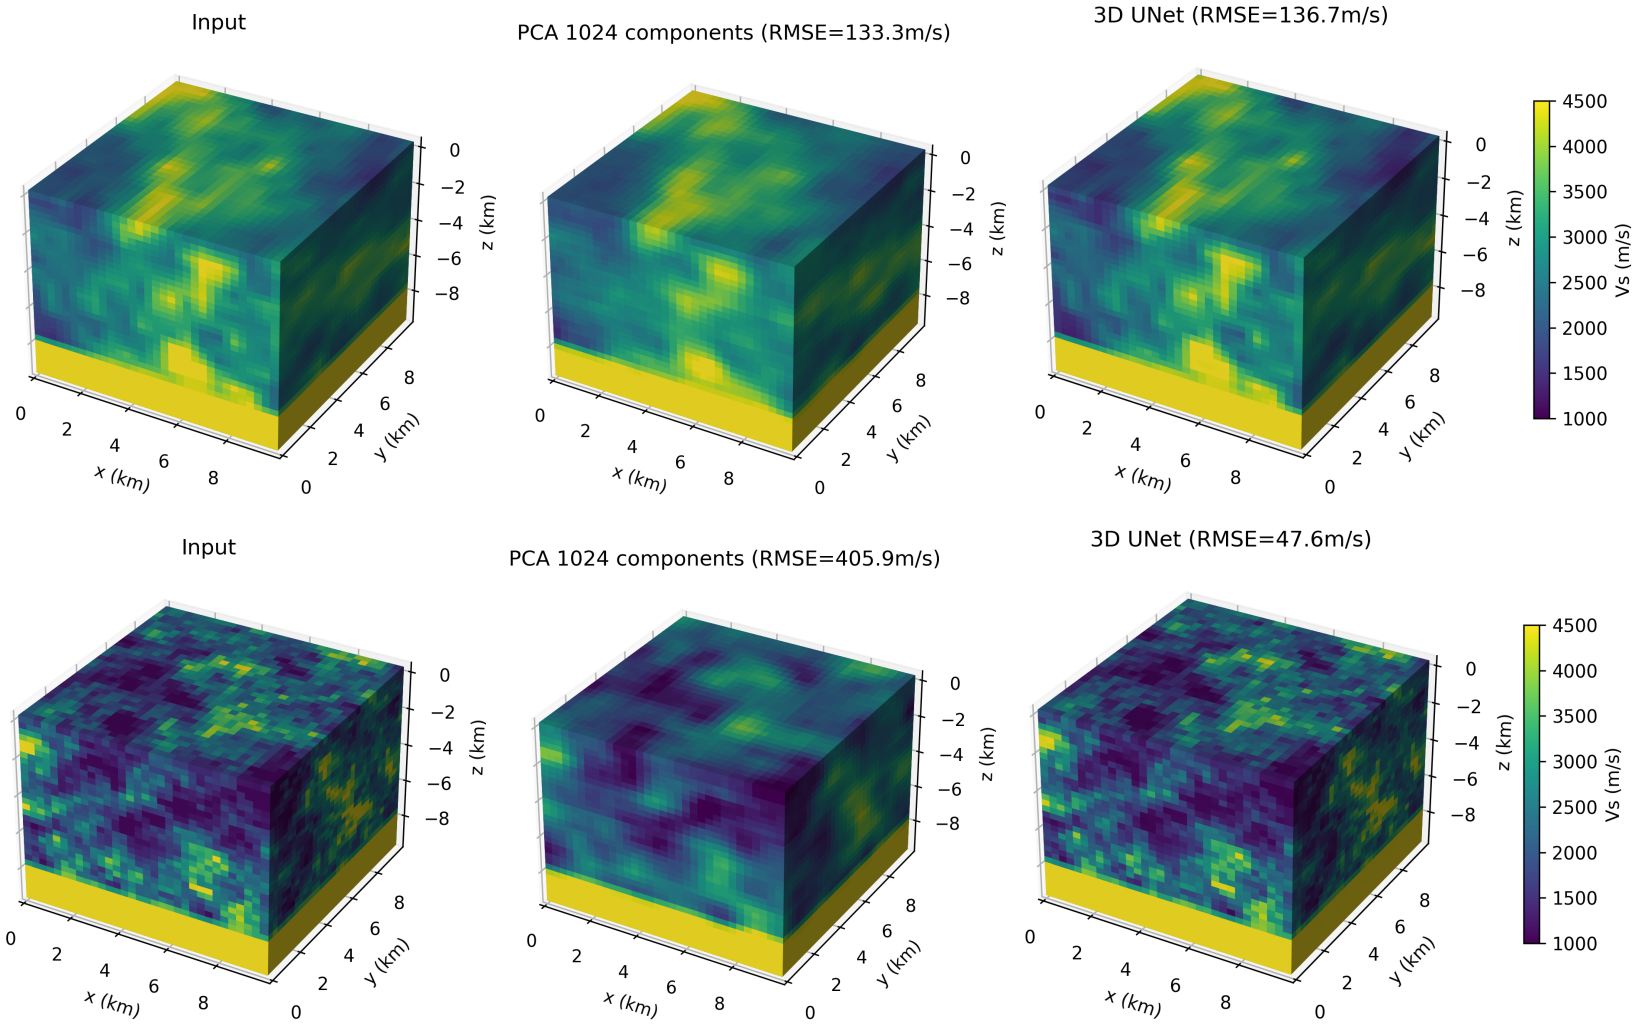
Input
PCA 1024 components (RMSE=133.3m/s)
3D UNet (RMSE=136.7m/s)
Input
PCA 1024 components (RMSE=405.9m/s)
3D UNet (RMSE=47.6m/s)
x (km)
y (km)
z (km)
Vs (m/s)

Dimensionality reduction [4]

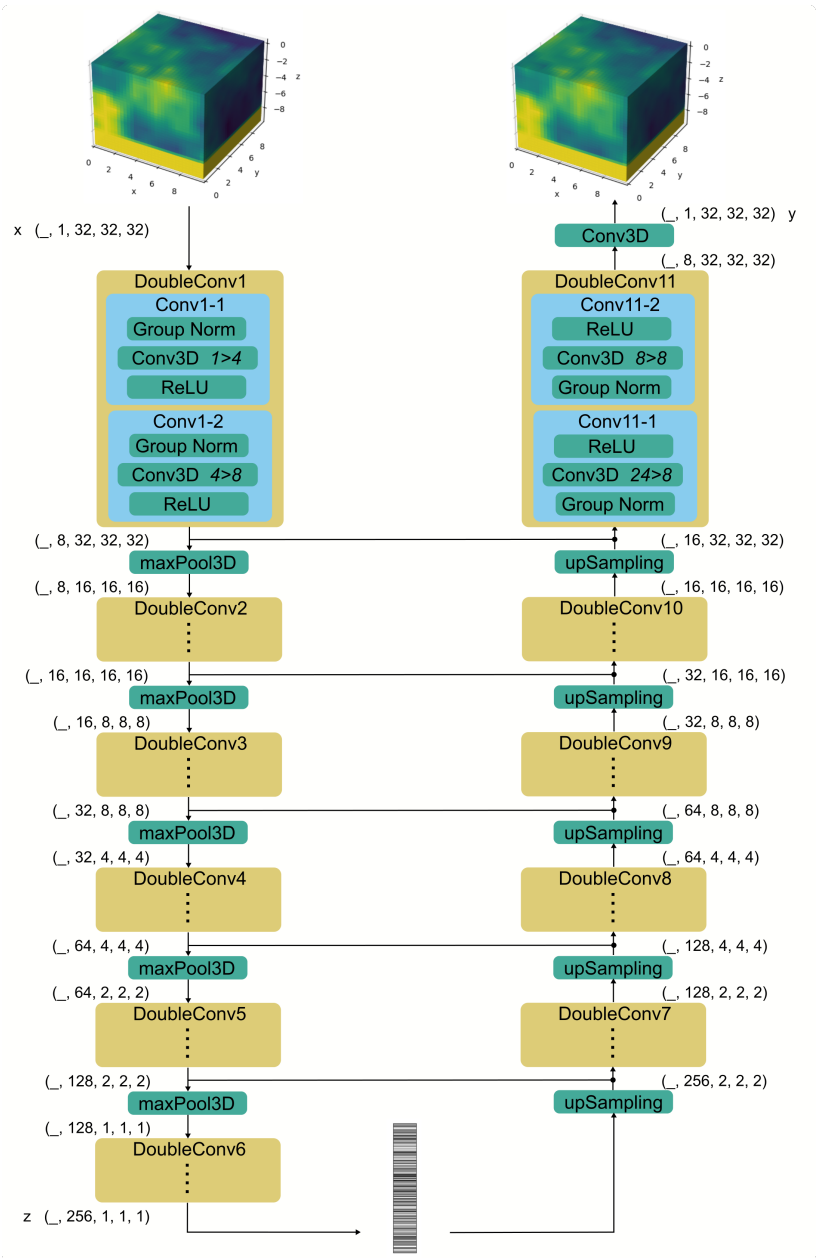

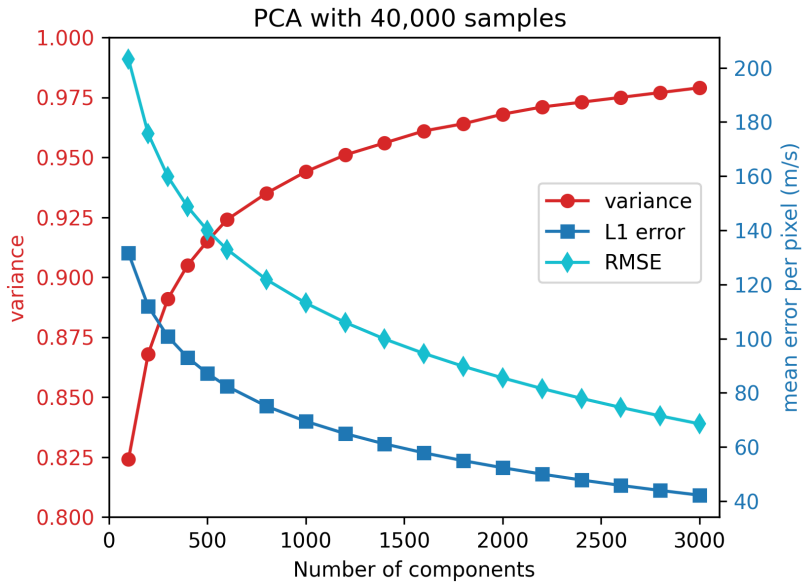

- PCA : 500 eigen-geologies explain 91.% of the variance (1000 eigen-geologies : 94.4%, 3000 eigen-geologies : 97.9%)
- Similar results with a 3D UNet of 4000000 weights trained over 10000 epochs (4 Nvidia Ampere A100 GPUs took 9.5 hours)

Dimensionality reduction [4]

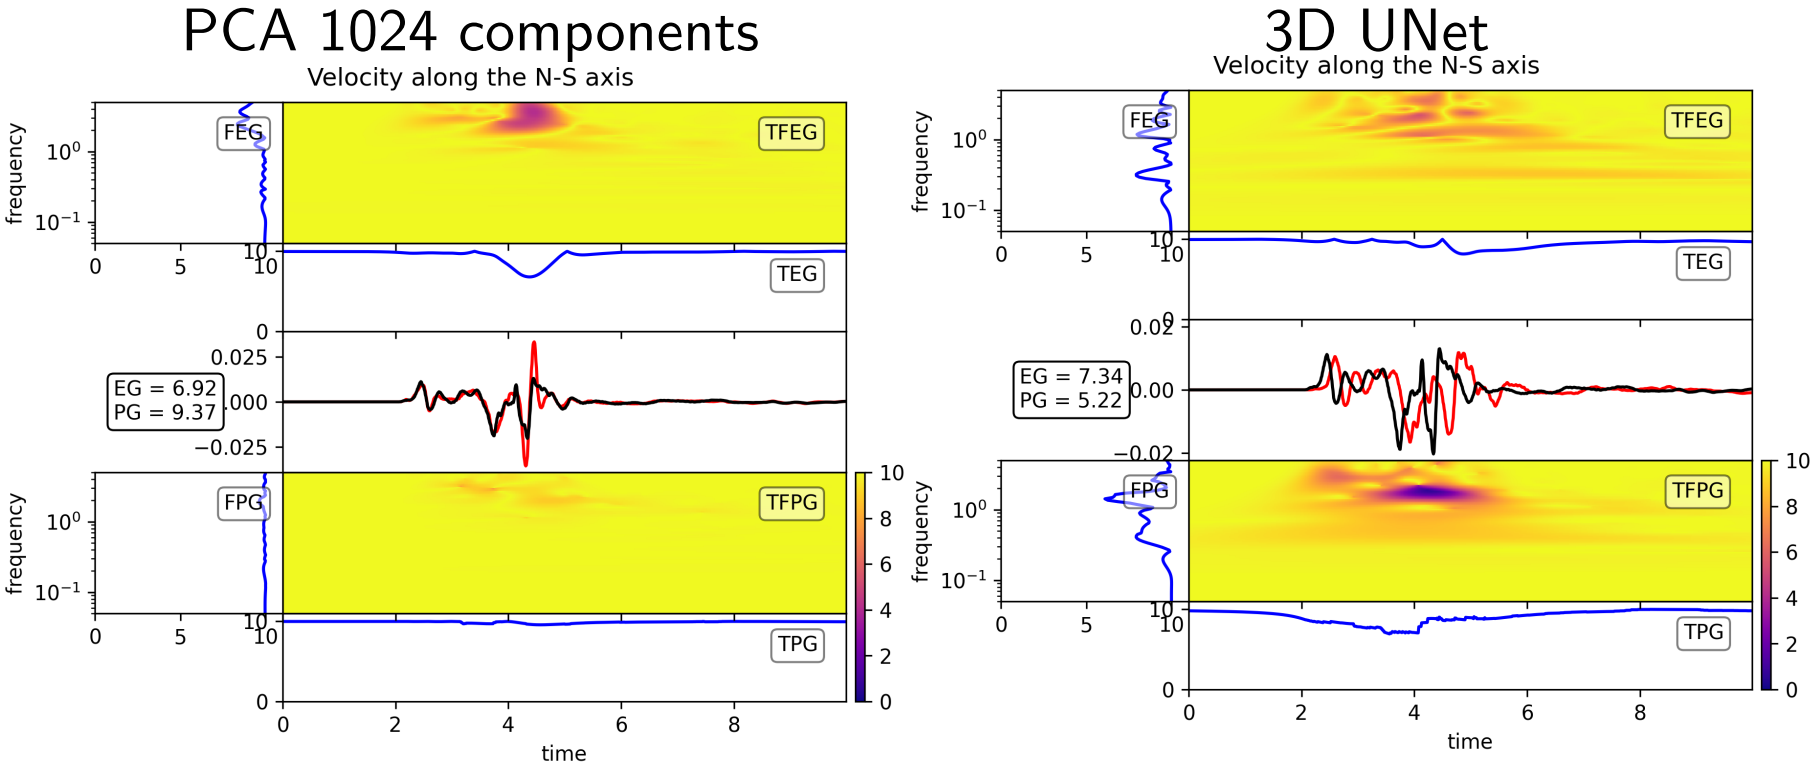

Machine learning opportunities to conduct high-fidelity earthquake simulations in multi-scale heterogeneous geology [4]

- Can we entirely replace simulations with AI ? (neural operators might...see you tomorrow)
- Can we use AI to enlarge the frequency range ?

Earthquake neural simulator

Multiple-Input FNO for source-dependent 3D elastodynamics

3D geological model

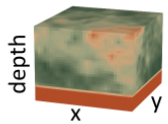

efficient computation of 3D structured grid with Factorized FNO

Source characteristics

$$\begin{pmatrix} x_s \\ y_s \\ z_s \end{pmatrix} + \begin{pmatrix} M_{xx} & M_{xy} & M_{xz} \\ & M_{yy} & M_{yz} \\ & & M_{zz} \end{pmatrix}$$

dimension increase to include in Fourier layers

MIFNO →

Surface velocity wavefields

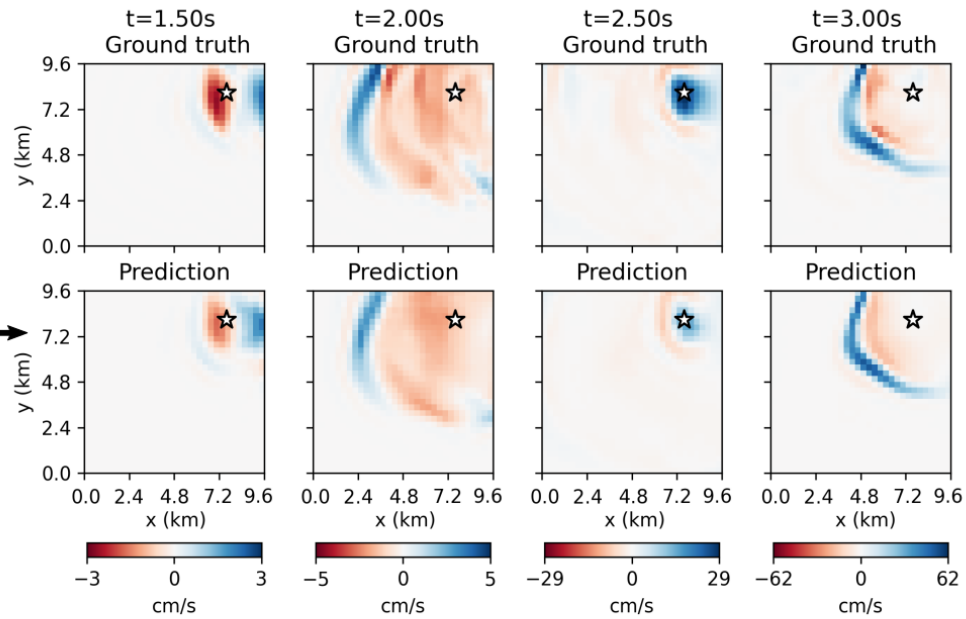

- ✓ Not autoregressive ⇒ long-term stability
- ✓ Excellent wave arrival times
- ✓ Good to excellent amplitudes

Generalization to:

- ✓ Out-of-distribution sources
- ✓ Overthrust geologies
- ✓ Higher resolution

Neural operators

HEMEW-3D (link) : 30000+ SEM3D simulations at 5 Hz (open-source !)

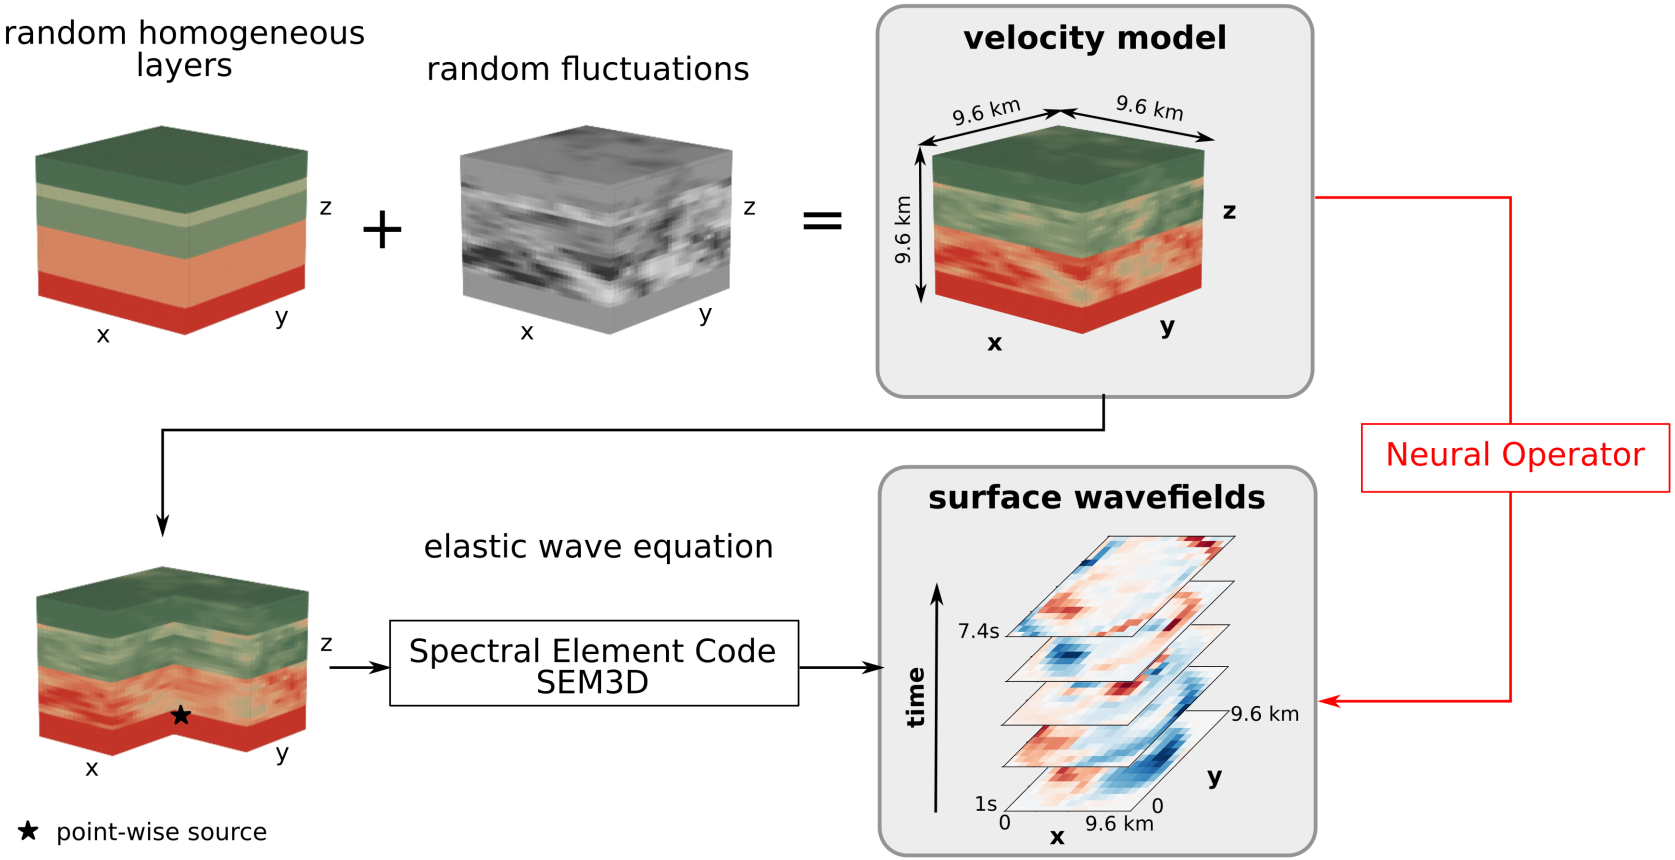

Sought non-linear mapping $\mathcal{G} : L^2(\Omega) \ni c_S \mapsto \boldsymbol{v}|_{\partial\Omega_n} \in \mathcal{S}^t$

$\mathcal{S}^t := \left\{ \boldsymbol{v} : \Omega \times [0, T] \to \mathbb{R}^3 : \boldsymbol{v} \in H^1(\Omega)^3 \, \forall t \in [0, T] \right\}$

Neural operator : $\mathcal{G}_\theta(c_S)(\boldsymbol{x}) = \int_\Omega \mathcal{K}_\theta(\boldsymbol{x}, \boldsymbol{y}, c_S(\boldsymbol{x}), c_S(\boldsymbol{y})) \, c_S(\boldsymbol{y}) \, d\boldsymbol{y}$

Multiple-Input Fourier Neural Operator (MIFNO)

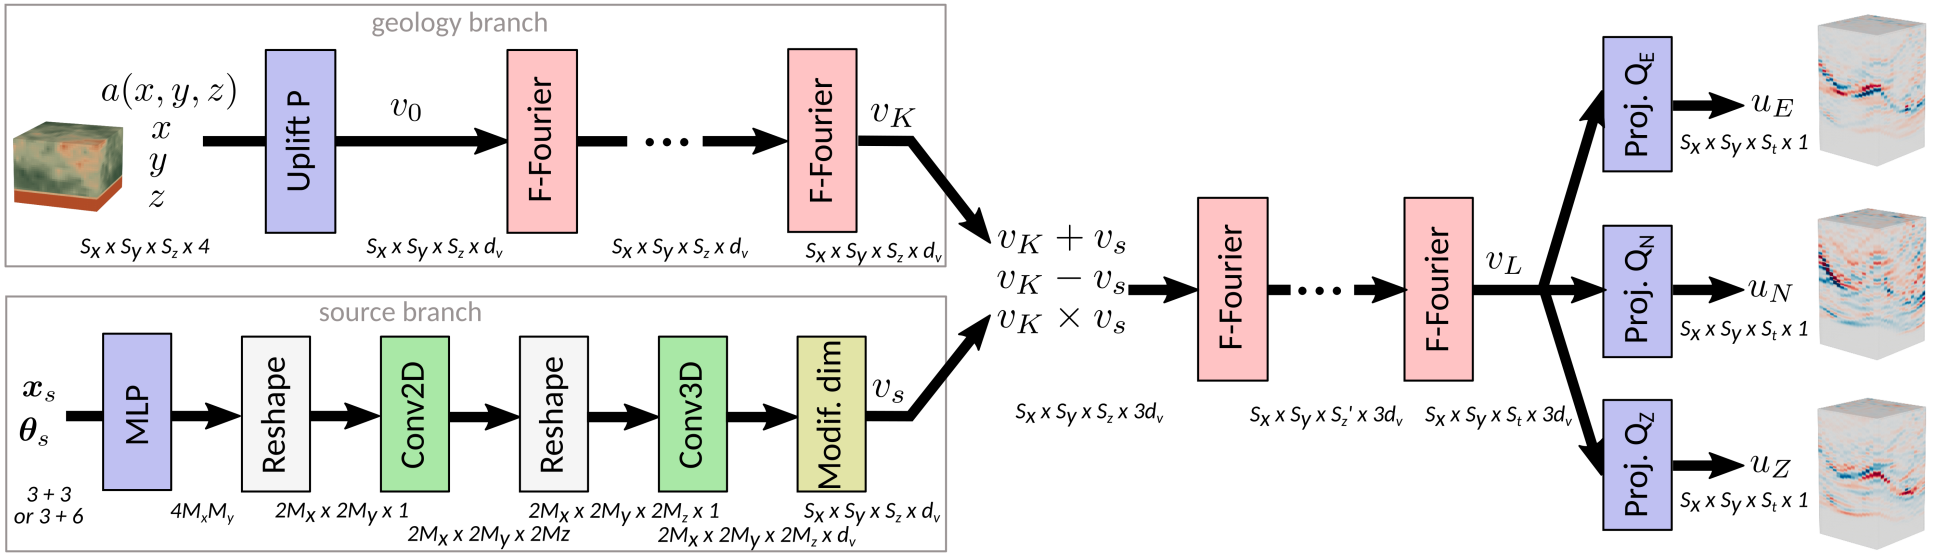

Figure – The MIFNO is made of a *geology branch* that encodes the geology with factorized Fourier (F-Fourier) layers, and a *source branch* that transforms the vector of source parameters $(\boldsymbol{x}_S, \boldsymbol{\theta}_S)$ into a 4D variable v_S matching the dimensions of the *geology branch* output v_K. Outputs of each branch are concatenated after elementary mathematical operations and the remaining factorized Fourier layers are applied. Uplift P and projection Q_E, Q_N, Q_Z blocks are the same as in the F-FNO.

MIFNO

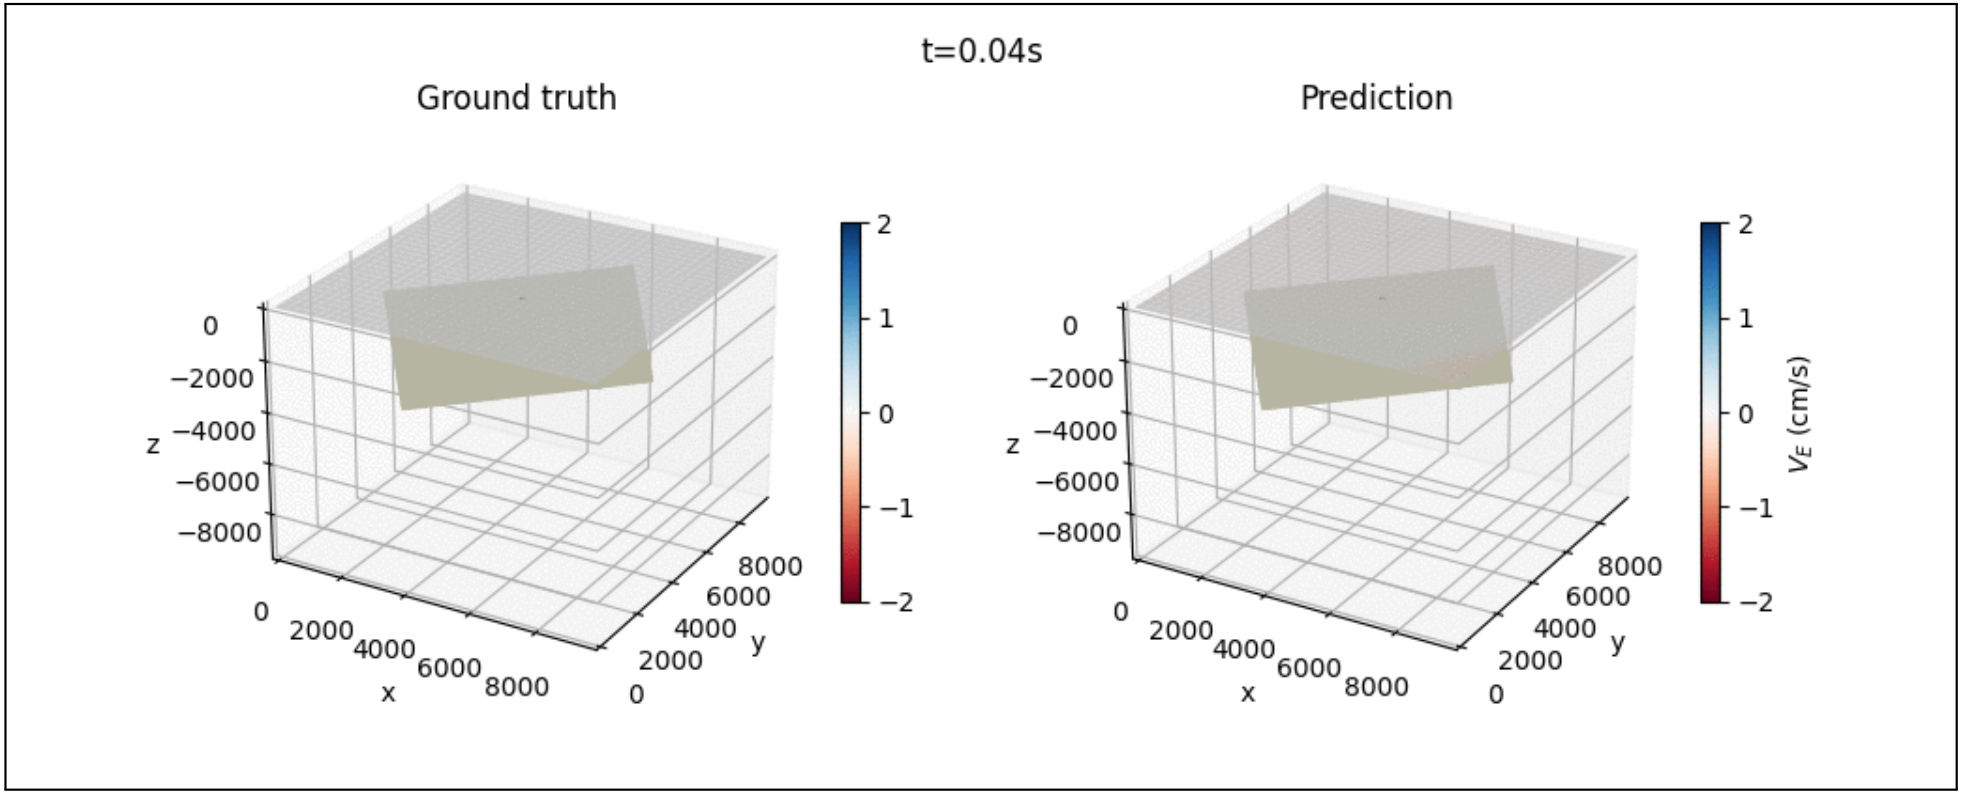

Figure – Source-to-site finite-fault simulation with MIFNO.

1. choose randomly one geology in the Le Teil database
2. for each point on the fault plane, predict the surface wavefields with the fine-tuned MIFNO
3. deconvolve all surface wavefields by the spice bench source time function
4. convolve each surface wavefield by the source time function corresponding to its associated point source
5. sum all surface wavefields

Niigata earthquake 2007

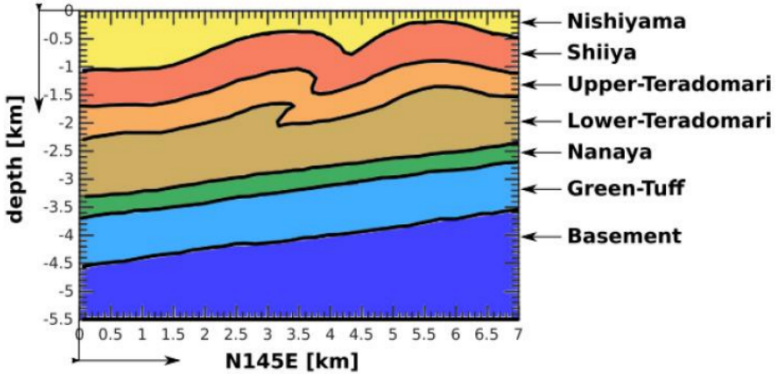

East-West velocity field

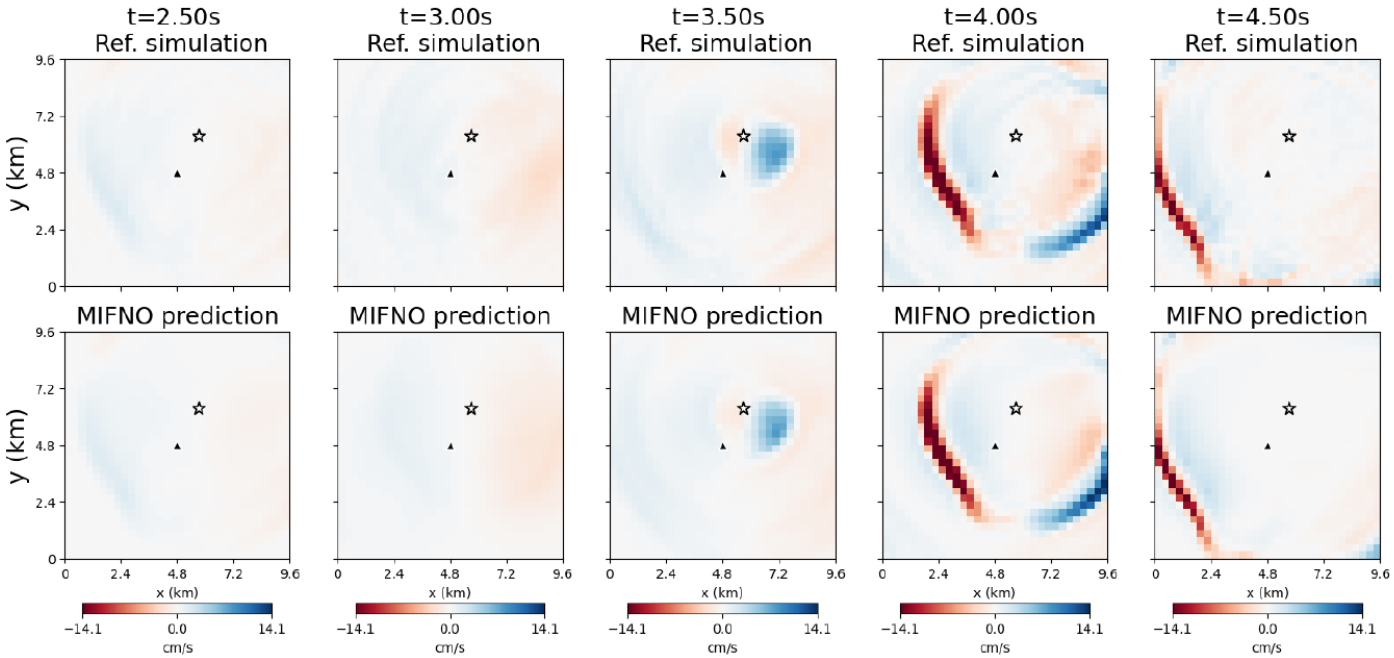

Zero-shot MIFNO

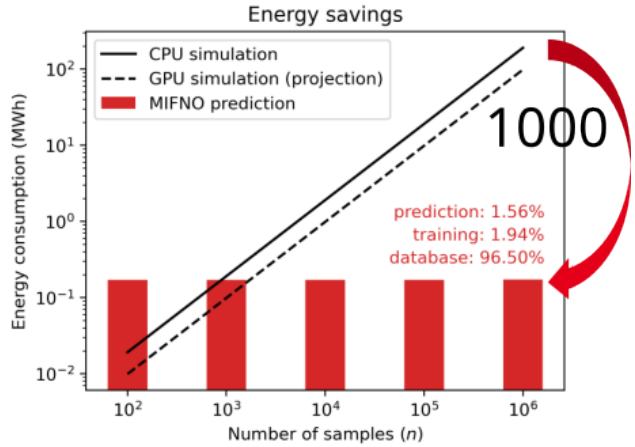

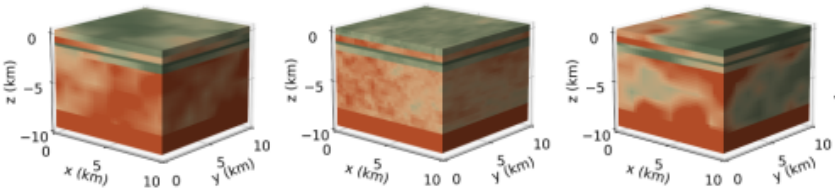

✗ Some inputs may be incompatible with reality

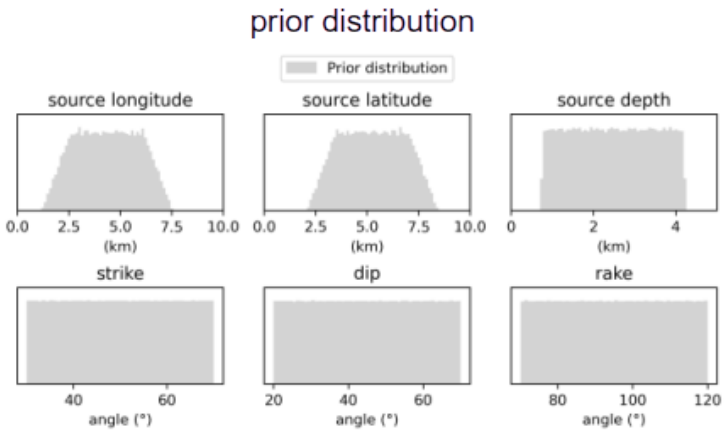

Fast predictions

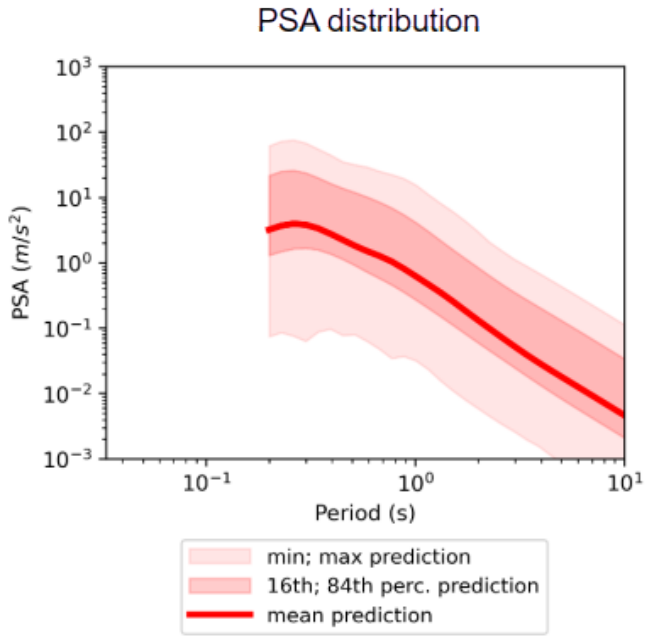

✗ PSA range is very large

Best-matching geologies

A small-magnitude event with PSA has been observed

Le Teil aftershock November 23

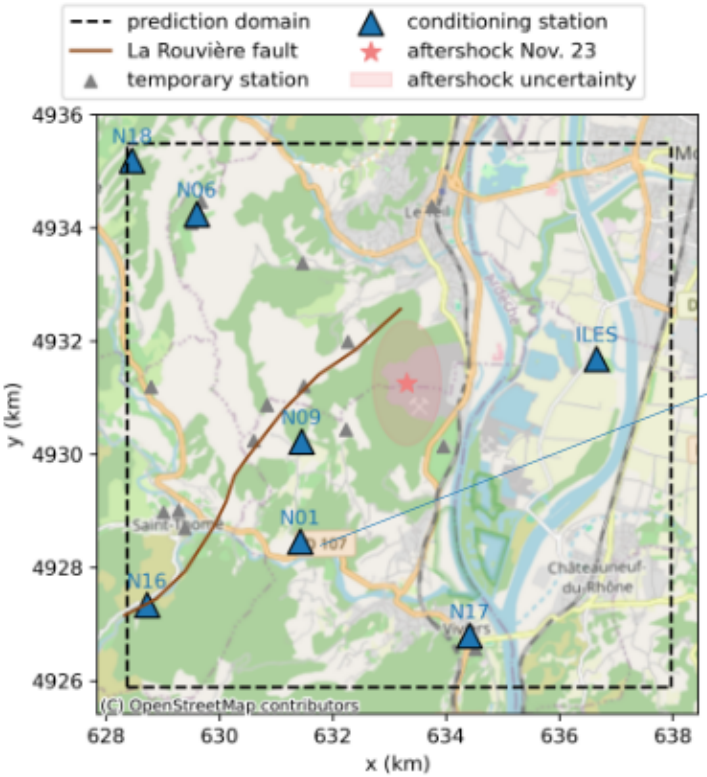

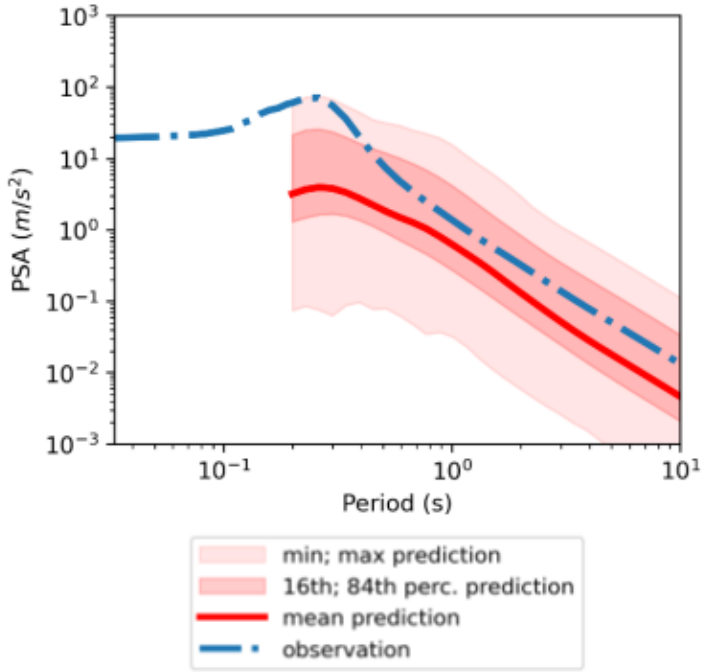

Best-matching geologies (cont'd)

posterior distribution

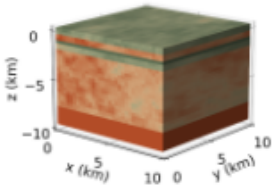

✓ We have selected a set of admissible geologies

posterior distribution

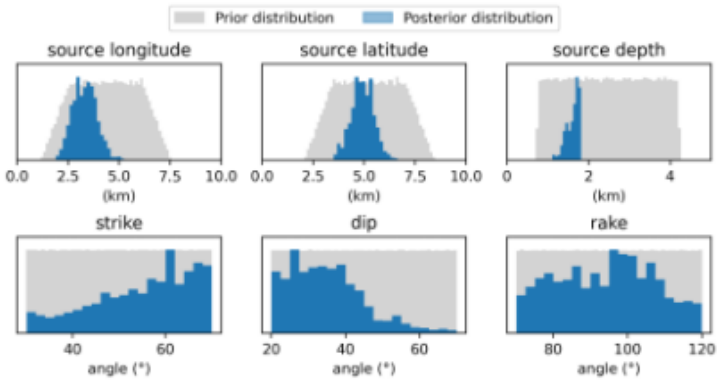

✓ We are able to constrain the source parameters

Select inputs that satisfy the observation

PSA distribution

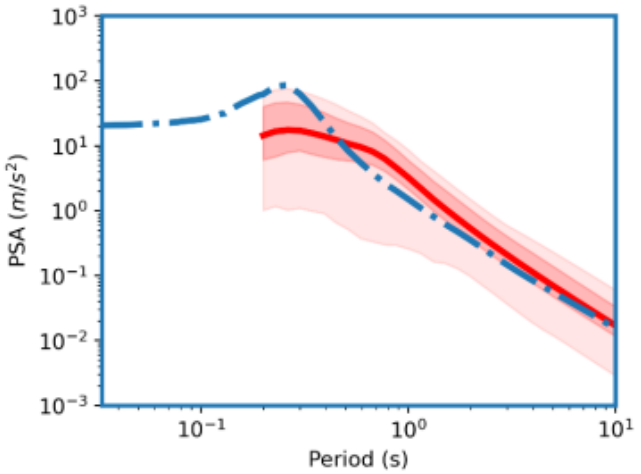

Best-matching sources

« Free » Bayesian solution to the inverse problem coherent with conventional inversion techniques

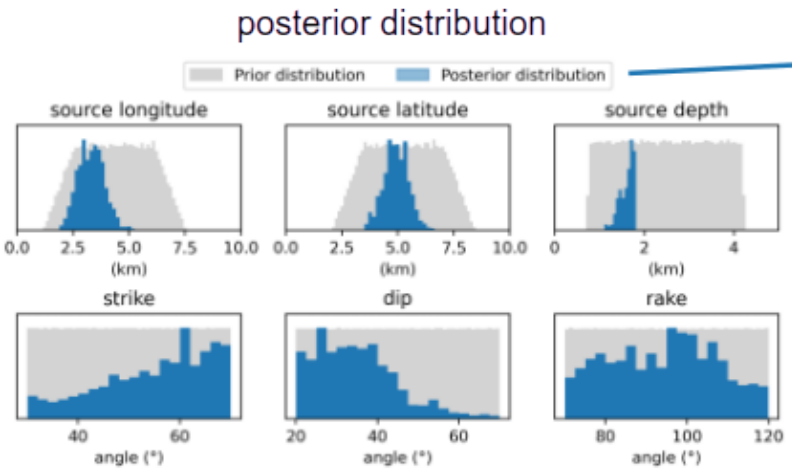

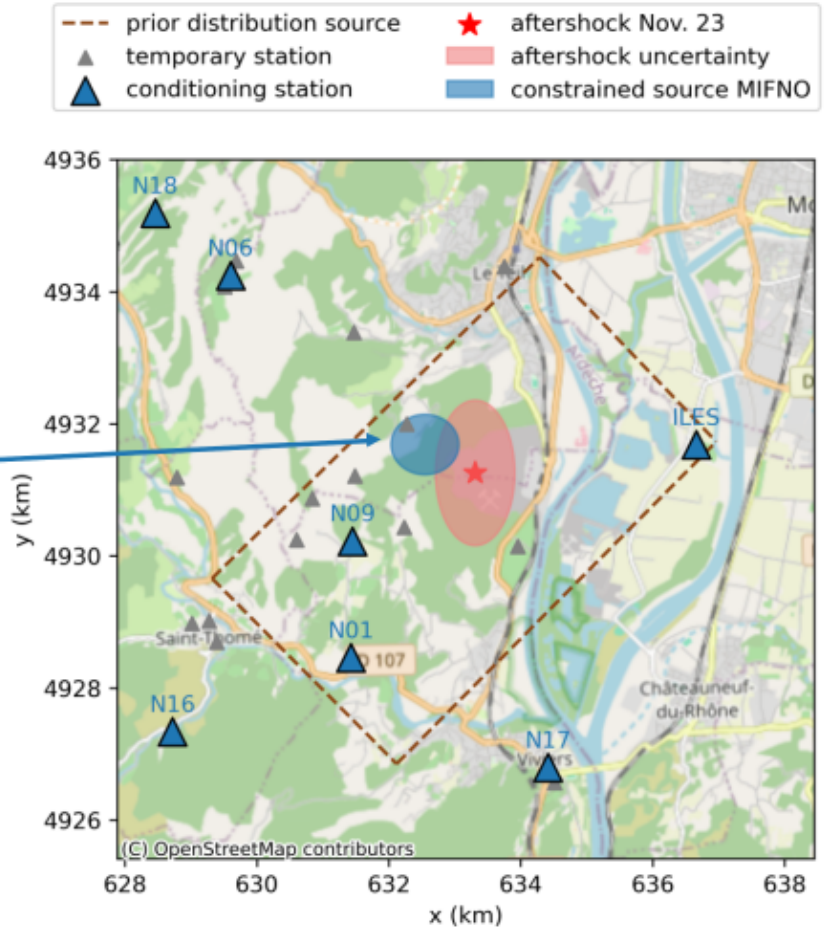

Data-tuned MIFNO

1000 conditioned geologies

Fault source model

1000 MIFNO predictions (0.01 core.hour)

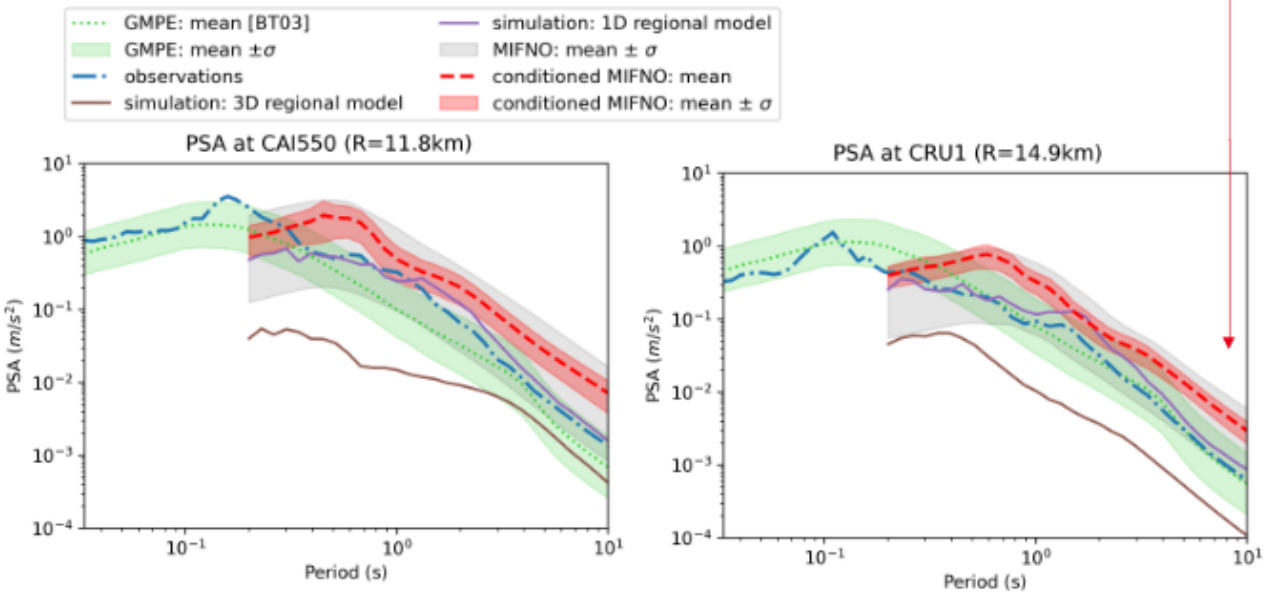

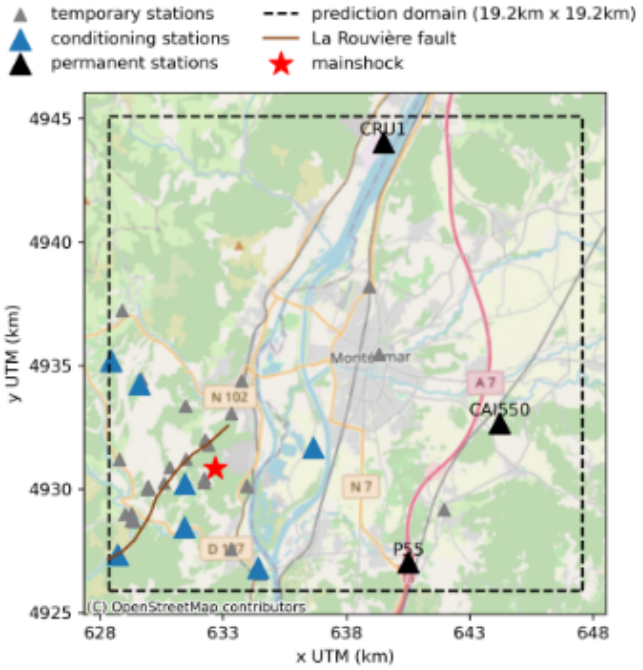

- ✓ Conditioned distributions
- ✓ For a strong event
- ✓ On permanent stations

Earthquake neural simulator

Conditioned denoising diffusion on seismic records

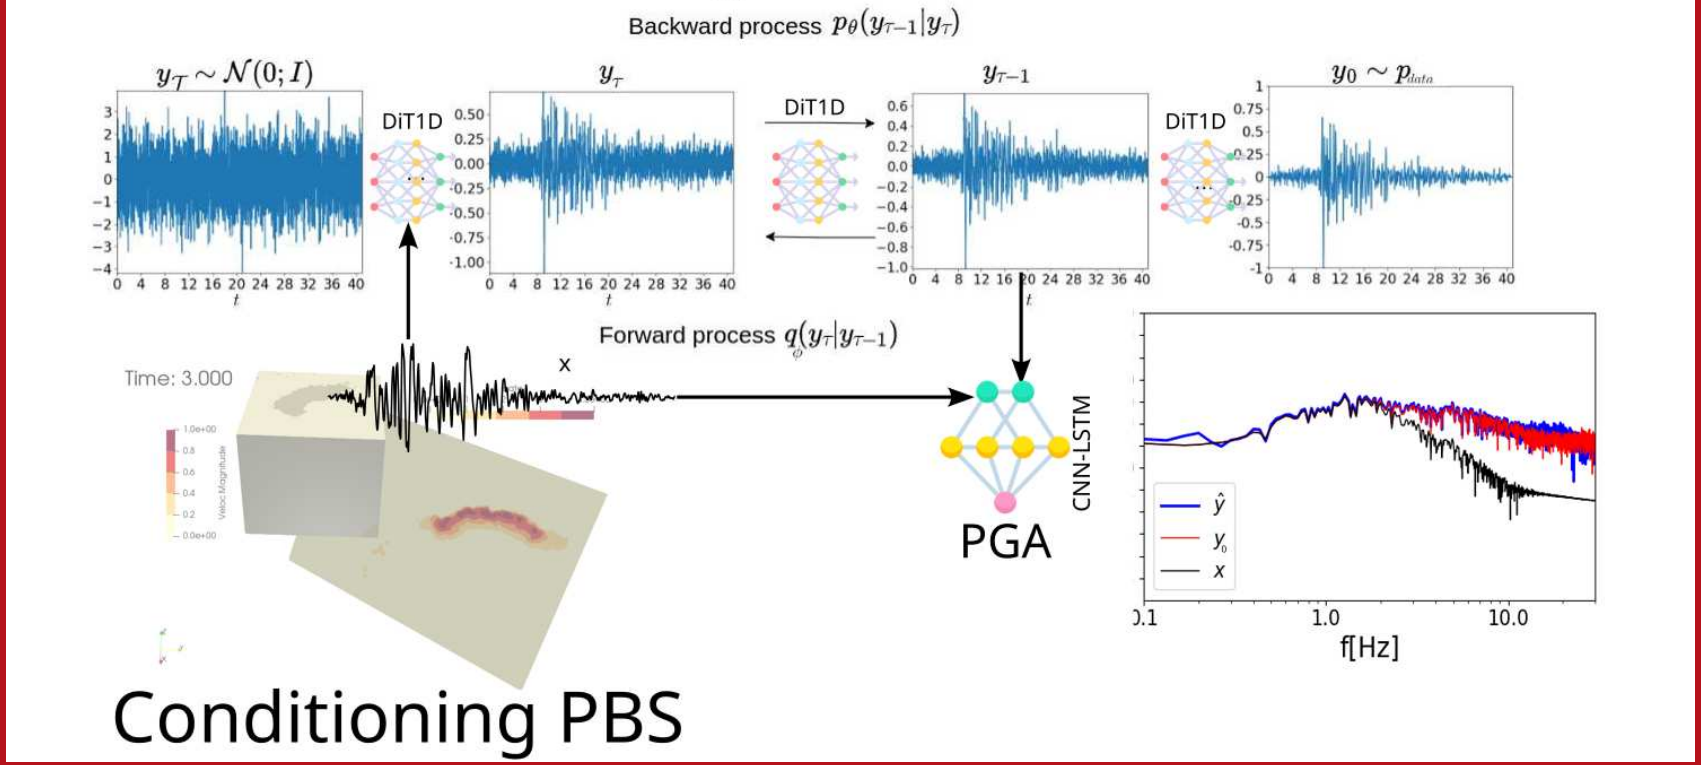

The “plague” of spectral bias

Model	rRMSE	rFFT_{low}	rFFT_{mid}	$\text{rFFT}_{\text{high}}$
FNO	**0.23 ± 0.088**	**-0.17 ± 0.18**	**-0.30 ± 0.24**	**-0.44 ± 0.30**
U-NO	0.22 ± 0.088	-0.16 ± 0.16	-0.28 ± 0.22	-0.44 ± 0.28
G-FNO	**0.26 ± 0.073**	**-0.53** ± 0.12	**-0.76 ± 0.11**	**-0.94 ± 0.04**
F-FNO $N^{(\ell)}=4$	0.23 ± 0.081	-0.29 ± 0.16	-0.42 ± 0.21	-0.55 ± 0.26
F-FNO $N^{(\ell)}=8$	0.21 ± 0.083	-0.15 ± 0.14	-0.26 ± 0.19	-0.37 ± 0.26
F-FNO $N^{(\ell)}=28$	**0.17 ± 0.081**	**-0.08 ± 0.10**	**-0.14 ± 0.15**	**-0.23 ± 0.22**

Table – Mean and standard deviation of relative biases for 1000 validation samples. rFFT_{low}, rFFT_{mid}, $\text{rFFT}_{\text{high}}$: frequency biases/deviations [5].

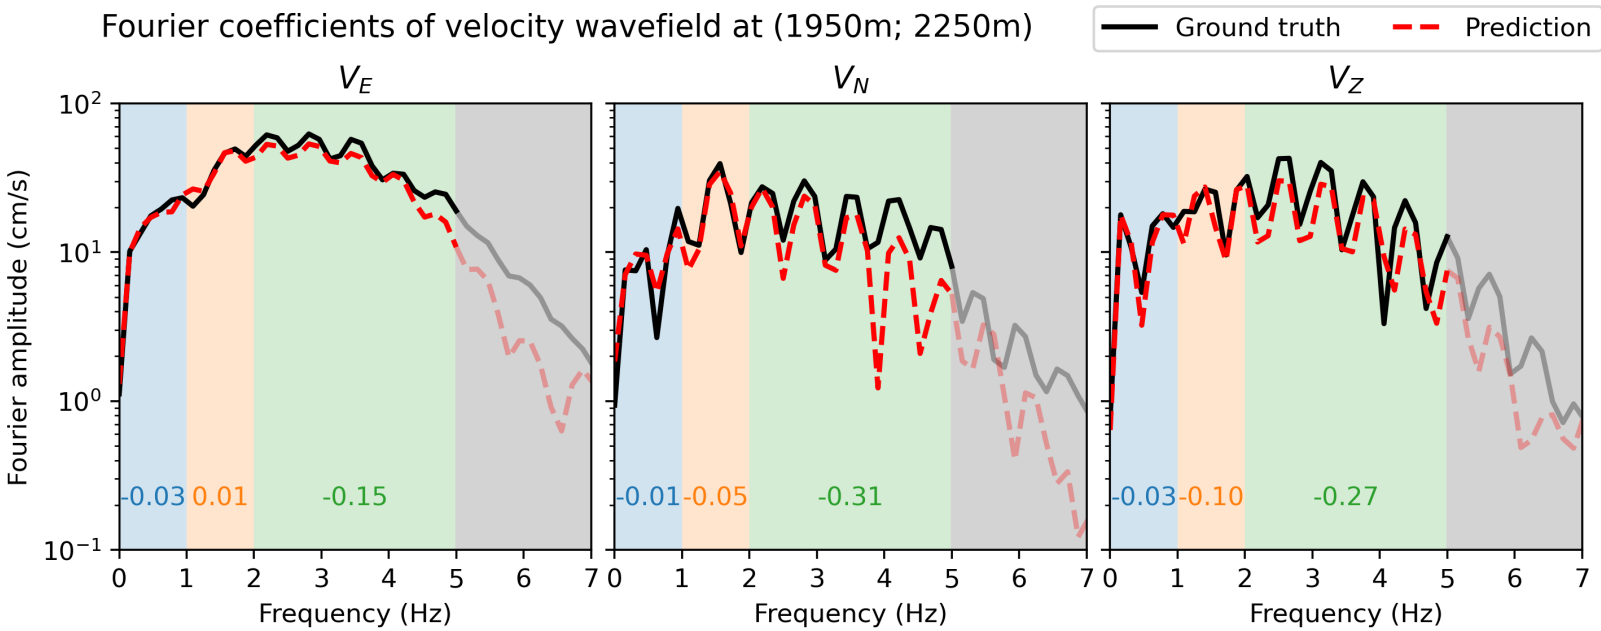

FreqNO-DPS : fuse a diffusion prior, sensors, and a frozen NO

$$\nabla_{\boldsymbol{u}_\tau} \log p_\tau(\boldsymbol{u}_\tau \mid \boldsymbol{y}, \boldsymbol{u}_{\mathrm{NO}}) \approx \underbrace{\frac{\bar{\boldsymbol{u}}_0 - \boldsymbol{u}_\tau}{\sigma_\tau^2}}_{\text{prior (Tweedie)}} + \underbrace{\nabla_{\boldsymbol{u}_\tau} \log p(\boldsymbol{y} \mid \boldsymbol{u}_\tau)}_{\text{sensor anchor}} + \underbrace{\nabla_{\boldsymbol{u}_\tau} \log p(\boldsymbol{u}_{\mathrm{NO}} \mid \boldsymbol{u}_\tau)}_{\text{NO guidance}}$$

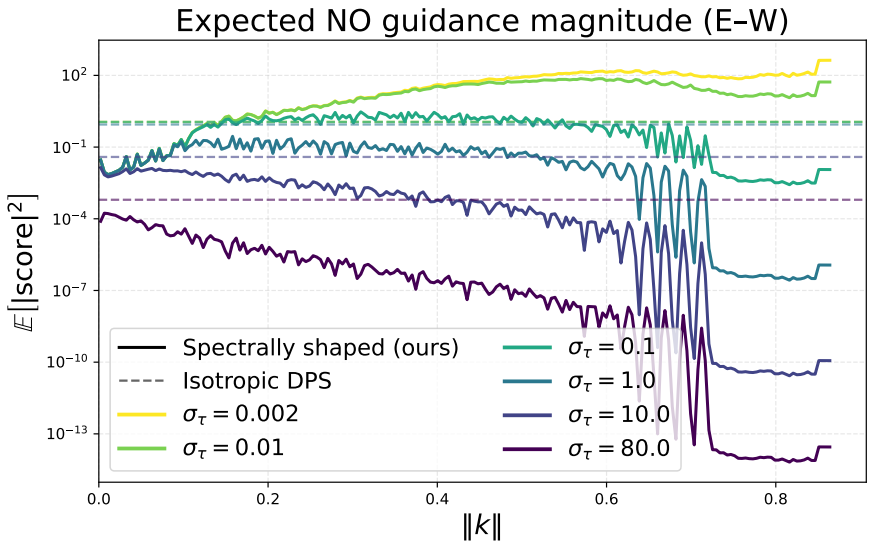

Naive (isotropic) NO guidance reimports the surrogate's bias.

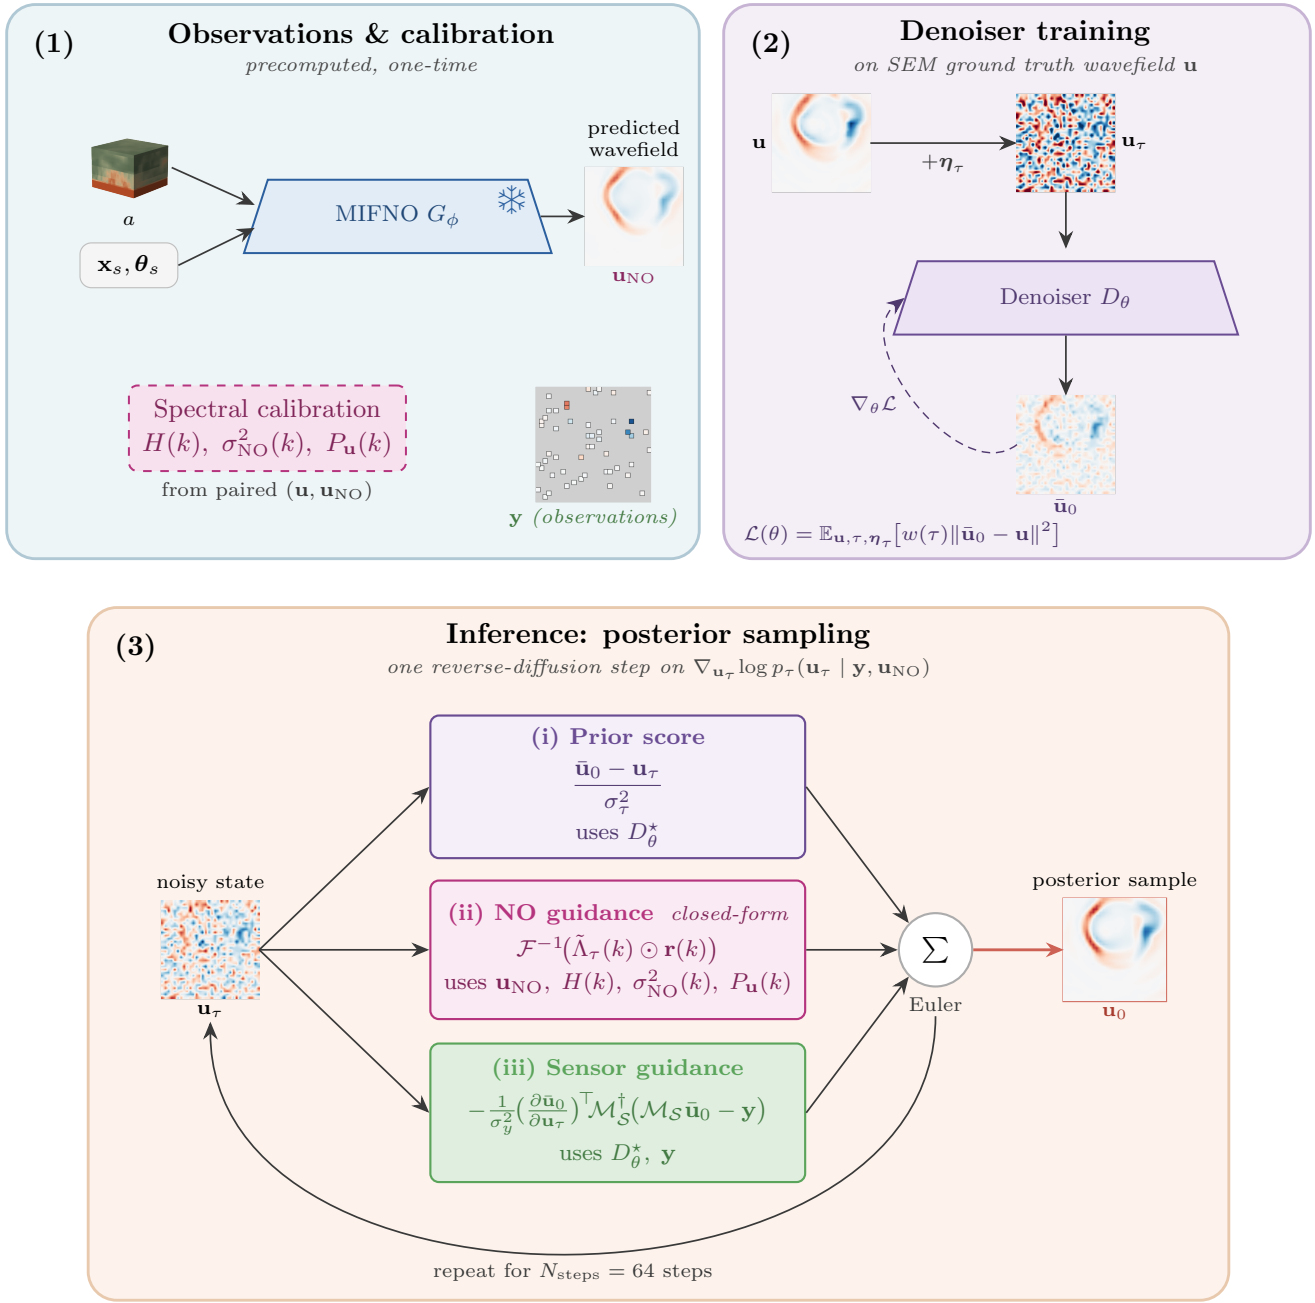
(1)
Observations & calibration
precomputed, one-time
predicted wavefield
MIFNO G_ϕ
a
$\mathbf{x}_s, \boldsymbol{\theta}_s$
$\mathbf{u}_{\mathrm{NO}}$
Spectral calibration
$H(k),\ \sigma^2_{\mathrm{NO}}(k),\ P_{\mathbf{u}}(k)$
from paired $(\mathbf{u}, \mathbf{u}_{\mathrm{NO}})$
$\mathbf{y}$ (observations)
(2)
Denoiser training
on SEM ground truth wavefield $\mathbf{u}$
$\mathbf{u}$
$+\boldsymbol{\eta}_\tau$
$\mathbf{u}_\tau$
Denoiser D_θ
$\nabla_\theta \mathcal{L}$
$\bar{\mathbf{u}}_0$
$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{u},\tau,\boldsymbol{\eta}_\tau}[w(\tau)\|\bar{\mathbf{u}}_0 - \mathbf{u}\|^2]$
(3)
Inference: posterior sampling
one reverse-diffusion step on $\nabla_{\mathbf{u}_\tau} \log p_\tau(\mathbf{u}_\tau \mid \mathbf{y}, \mathbf{u}_{\mathrm{NO}})$
(i) Prior score
$\frac{\bar{\mathbf{u}}_0 - \mathbf{u}_\tau}{\sigma^2_\tau}$
uses $D^\star_\theta$
(ii) NO guidance closed-form
$\mathcal{F}^{-1}\big(\tilde{\Lambda}_\tau(k) \odot \mathbf{r}(k)\big)$
uses $\mathbf{u}_{\mathrm{NO}},\ H(k),\ \sigma^2_{\mathrm{NO}}(k),\ P_{\mathbf{u}}(k)$
(iii) Sensor guidance
$-\frac{1}{\sigma^2_y}\big(\frac{\partial \bar{\mathbf{u}}_0}{\partial \mathbf{u}_\tau}\big)^\top \mathcal{M}^\dagger_{\mathcal{S}}(\mathcal{M}_{\mathcal{S}}\bar{\mathbf{u}}_0 - \mathbf{y})$
uses $D^\star_\theta,\ \mathbf{y}$
noisy state
$\mathbf{u}_\tau$
Σ
Euler
posterior sample
$\mathbf{u}_0$
repeat for $N_{\mathrm{steps}} = 64$ steps

Spoiler alert !

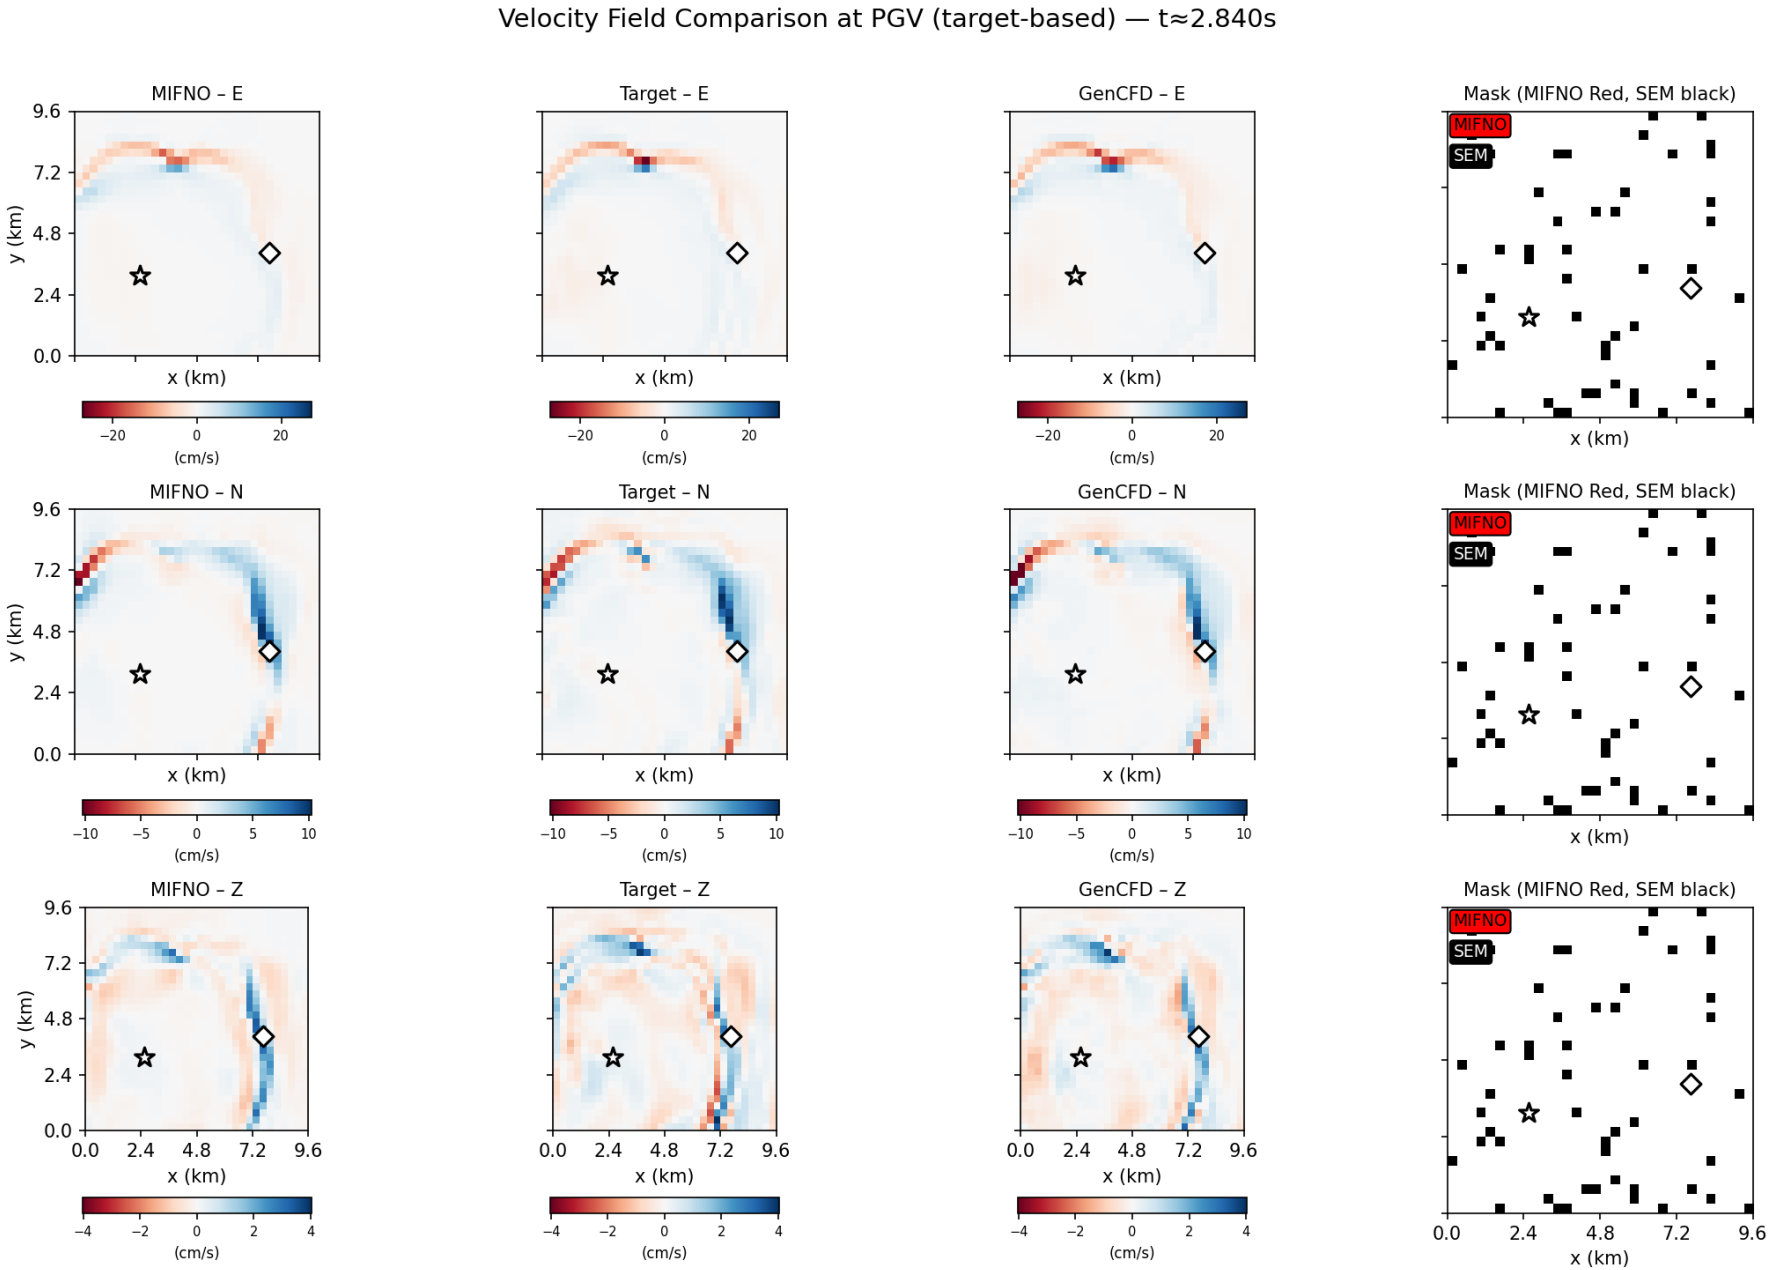

The key idea : spectrally shaped NO guidance

Per-mode error model, calibrated from paired data :

$$\mathcal{F}(\boldsymbol{u}_{\mathrm{NO}})(\boldsymbol{k})_c = H(\boldsymbol{k})\,\mathcal{F}(\boldsymbol{u})(\boldsymbol{k})_c + \hat{\eta}_{\mathrm{NO}}(\boldsymbol{k}), \quad \mathbb{E}|\hat{\eta}_{\mathrm{NO}}|^2 = \sigma^2_{\mathrm{NO}}(\boldsymbol{k}).$$

Two-step LMMSE, with $X = \mathcal{F}(\boldsymbol{u})_c,\ Y = \mathcal{F}(\boldsymbol{u}_\tau)_c,\ Z = \mathcal{F}(\boldsymbol{u}_{\mathrm{NO}})_c$:

Step 1 — diffusion channel : estimate clean X from the noisy state Y

Wiener filter α (shrinks Y toward the prior mean) and its error σ^2_L :

$$\alpha(\boldsymbol{k}) = \frac{P_{\boldsymbol{u}}(\boldsymbol{k})}{\sigma^2_\tau + P_{\boldsymbol{u}}(\boldsymbol{k})}, \qquad \sigma^2_L(\boldsymbol{k}) = \sigma^2_\tau\,\alpha(\boldsymbol{k}) \ \leq\ \min\big(P_{\boldsymbol{u}}, \sigma^2_\tau\big).$$

Step 2 — NO channel : refine with the surrogate Z

Innovation $r(\boldsymbol{k}) = Z - H(\boldsymbol{k})\,\alpha(\boldsymbol{k})\,Y$ has variance

$$\lambda_\tau(\boldsymbol{k}) = |H(\boldsymbol{k})|^2\,\sigma^2_L(\boldsymbol{k}) + \sigma^2_{\mathrm{NO}}(\boldsymbol{k}).$$

Result : near-zero spectral bias at 2–5% coverage

Closed-form score — two FFTs, no denoiser backprop :

$$\nabla_{\boldsymbol{u}_\tau} \log p(\boldsymbol{u}_{\mathrm{NO}} \mid \boldsymbol{u}_\tau) = \mathcal{F}^{-1}\big(\tilde{\lambda}_\tau \odot r\big), \quad \tilde{\lambda}_\tau(\boldsymbol{k}) = \frac{2\,H(\boldsymbol{k})\,\alpha(\boldsymbol{k})}{\lambda_\tau(\boldsymbol{k})} \xrightarrow[\text{high } \|\boldsymbol{k}\|]{} 0.$$

Method ($\rho = 5\%$)	rMAE ↓	$\mathrm{rFFT}_{\mathrm{high}} \to 0$
MIFNO (surrogate)	0.133	-0.239
DPS (sensors only)	0.113	-0.232
DPS + NO (iso)	**0.099**	-0.210
FreqNO-DPS	0.100	$+0.002$

Isotropic guidance improves pointwise error *but keeps the bias*. FreqNO-DPS reaches near-zero bias at no pointwise cost ($+0.047$ even at $\rho = 2\%$).

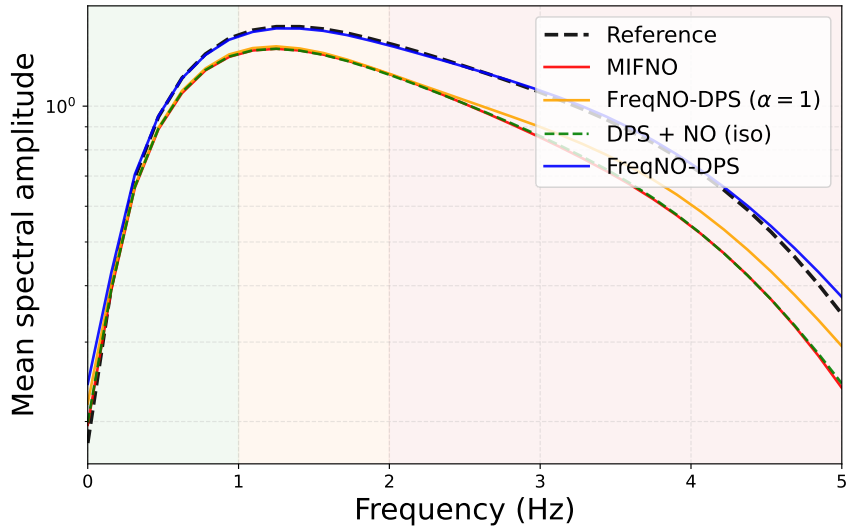

Ensemble-averaged spectrum, $\rho = 5\%$. Only FreqNO-DPS (blue) the reference across all bands.

Guiding regimes

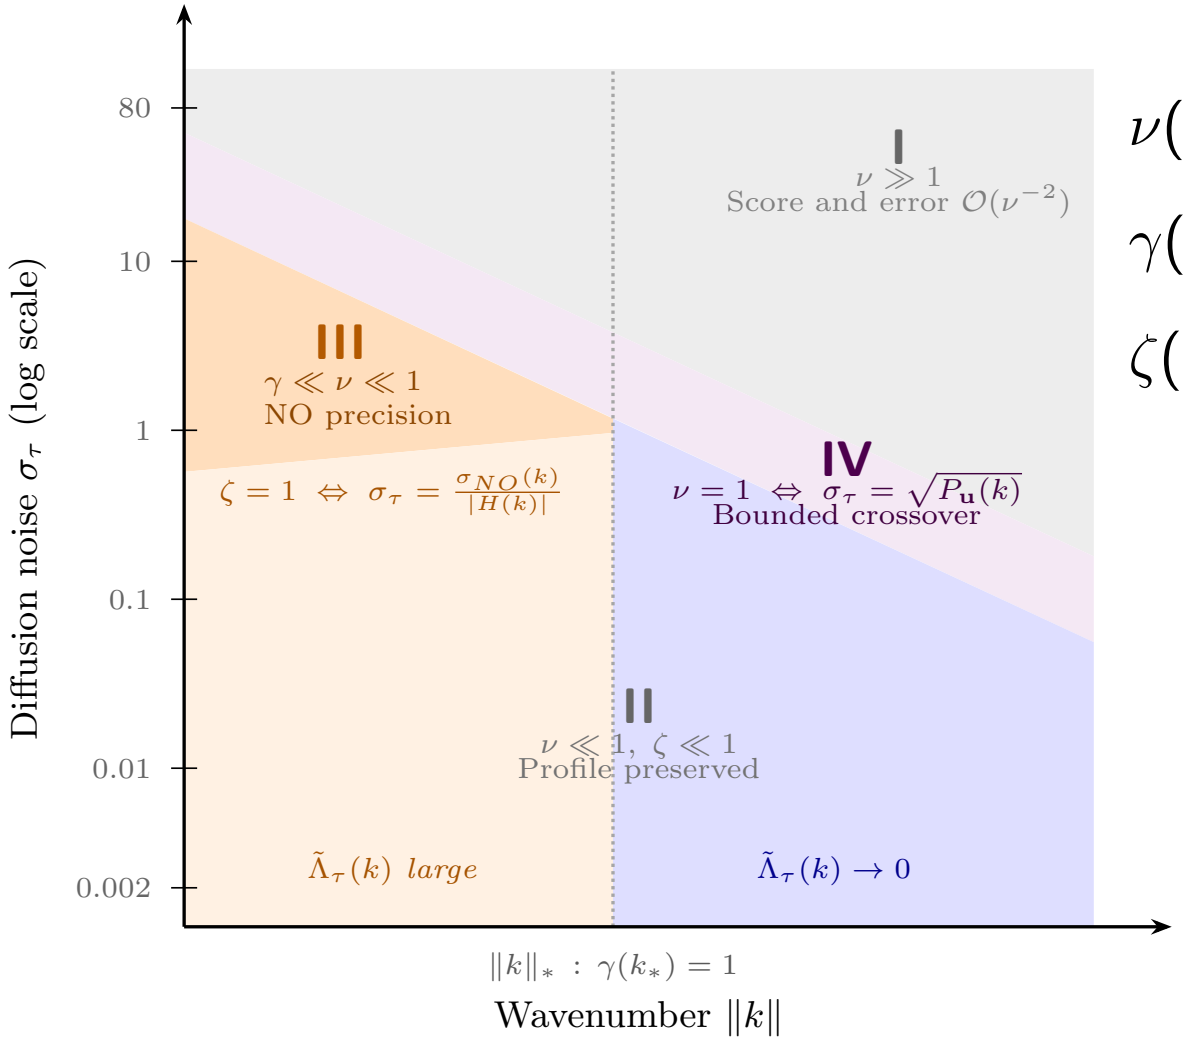

$$\nu(\boldsymbol{k},\tau) := \frac{\sigma_\tau^2}{P_{\boldsymbol{u}}(\boldsymbol{k})}$$

$$\gamma(\boldsymbol{k}) := \frac{\sigma_{\mathrm{NO}}^2(\boldsymbol{k})}{|H(\boldsymbol{k})|^2\, P_{\boldsymbol{u}}(\boldsymbol{k})}$$

$$\zeta(\boldsymbol{k},\tau) := \frac{|H(\boldsymbol{k})|^2\, \sigma_L^2(\boldsymbol{k})}{\sigma_{\mathrm{NO}}^2(\boldsymbol{k})} = \frac{\nu}{\gamma(1+\nu)}$$

NO guidance ($\gamma < 1$) Sensors + prior ($\gamma > 1$) Prior only ($\nu \gg 1$)

Takeaways

- Frequency-dependent calibration is essential, not merely beneficial : it is what separates useful NO guidance from a faithful reimport of its bias.
- **Cheap and modular** : closed form, two FFTs per step, no denoiser backprop ; wraps *any* frozen surrogate.
- **Distribution-agnostic** : an exact score identity plus error bounds across the frequency–diffusion-time plane preserve the guidance's spectral shape regardless of distributional assumptions.
- **Lightweight prerequisites** : only paired surrogate/reference data for calibration, plus a coherence diagnostic to verify applicability to a new surrogate.

Applies broadly to FNO-family surrogates across wave propagation, fluids, and other PDE settings.

Thank you !

Contact : filippo.gatti@centralesupelec.fr

Bibliography

Bibliography I

- **CEA and CentraleSupélec and IPGP and CNRS**. "SEM3D Ver 2017.04 Registered at French Agency for Protection of Programs (Dépôt APP)". IDDN.FR.001.400009.000.S.P.2018.000.31235. 2017.
- **Gatti, F. et al.** "Ground Motion Relations for the Geometric Mean Horizontal Component of Peak and Spectral Ground Motion Parameters". en. In : *9th european conference on numerical methods in geotechnical engineering, 25-27 june 2018, porto (portugal)*. tex.owner : filippo tex.timestamp : 2018.01.23. CRC Press, juin 2018, p. 829-835. ISBN : 978-1-351-00362-9. URL : https://www.taylorfrancis.com/books/9781351003612/chapters/10.1201/9781351003629-104.

Bibliography II

- **Korres, M. et al.** "Enhanced seismic response prediction of critical structures via 3D regional scale physics-based earthquake simulation". en. In : *Journal of Earthquake Engineering* 27.3 (oct. 2022). abstract : Realistic physics-based 3D earthquake simulation for source-to-structure wave propagation consists of a powerful numerical tool for seismic response prediction of critical structures submitted to high safety standards. Structural response considering soil-structure interaction (SSI) is usually estimated by Finite Element Method (FEM) approach, as it is considered as the most flexible numerical approach for nonlinear structural dynamics. However, current engineering practice considers seismic input motion as vertically incident plane waves, despite the fact that this assumption excludes wave passage effects for large infrastructures and surface waves appearing from possible local basin effects. In this