Compute optimal data scaling for neural surrogates via multi-fidelity training



Paul Setinek, Fabian Paischer, Nils Thuerey, Johannes Brandstetter, Felix Koehler





Motivation

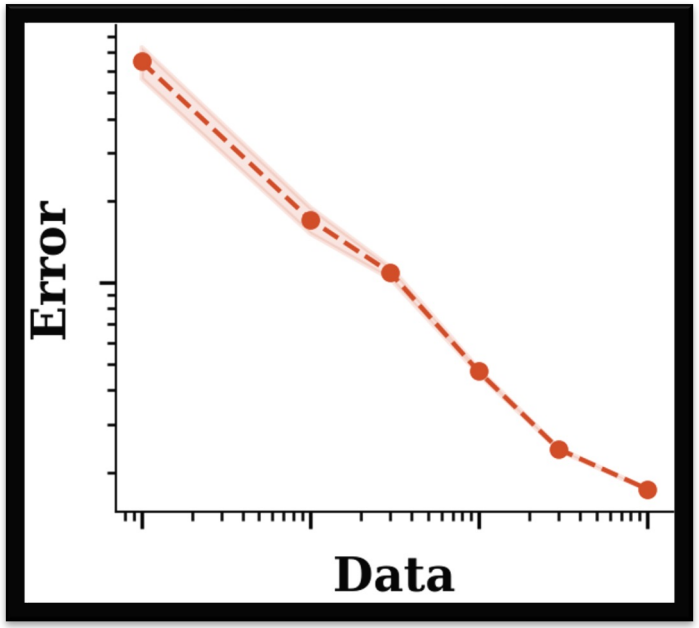

Motivation

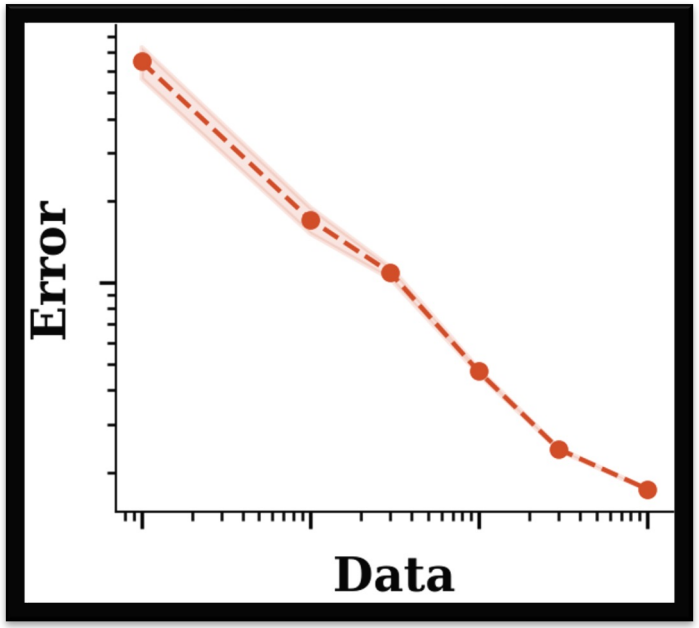

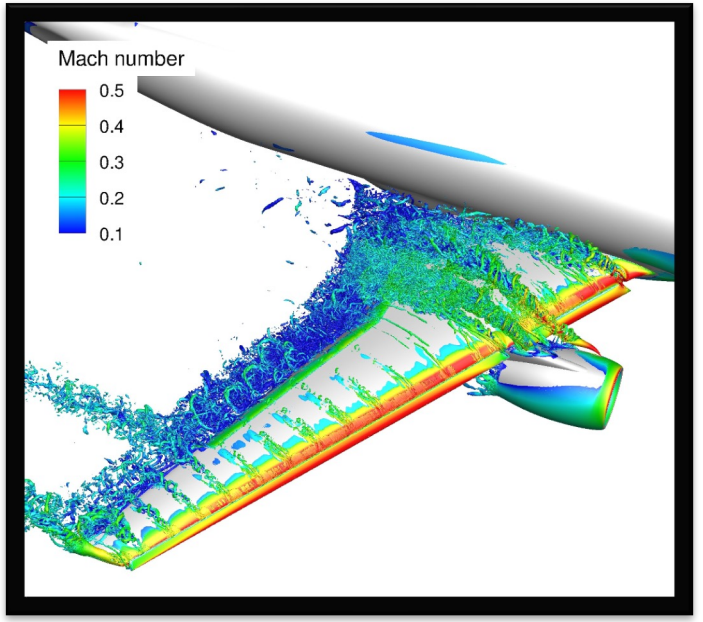

Image: DLR (CC BY-NC-ND 3.0)

~ 432 GPU hours (GB 200s) [1]

[1] Ashton et al., HiLiftAeroML: High-Fidelity CFD Dataset for High-Lift Aircraft Aerodynamics, arXiv:2605.19565, 2026.

Motivation

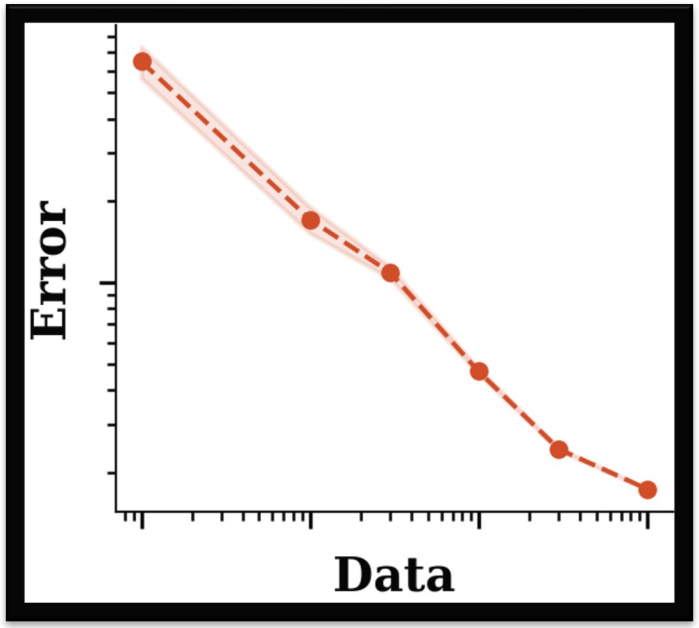

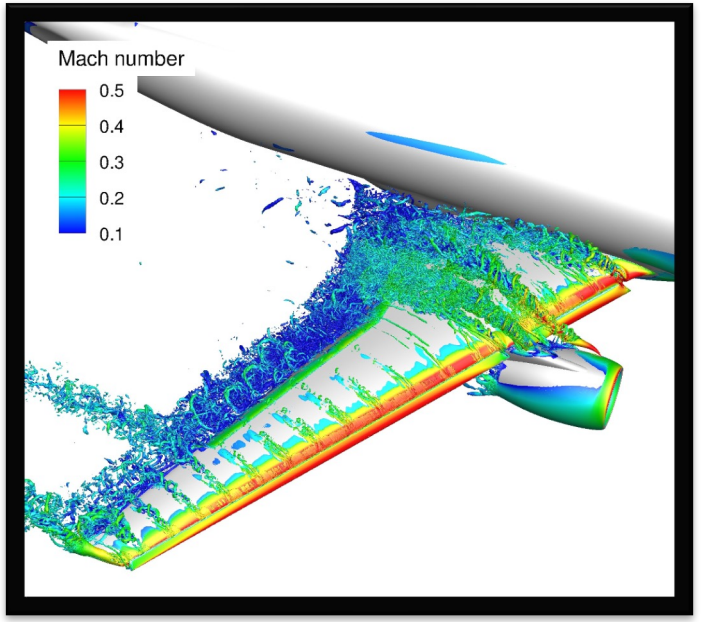

Image: DLR (CC BY-NC-ND 3.0)

~ 432 GPU hours (GB 200s) [1]

Data generation compute is the bottleneck in scaling up neural surrogates in AI4Sciene.

[1] Ashton et al., HiLiftAeroML: High-Fidelity CFD Dataset for High-Lift Aircraft Aerodynamics, arXiv:2605.19565, 2026.

Simulation Pipeline

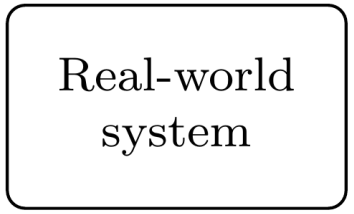

Simulation Pipeline

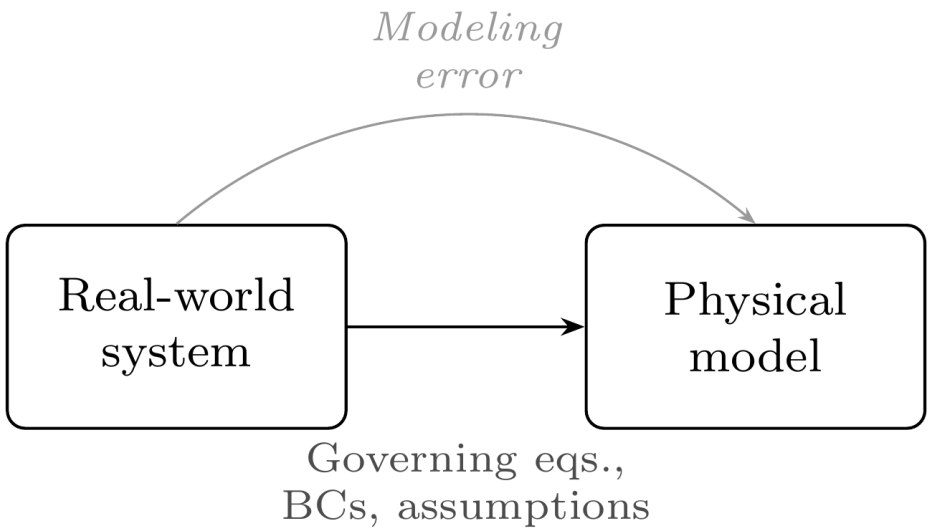

Simulation Pipeline

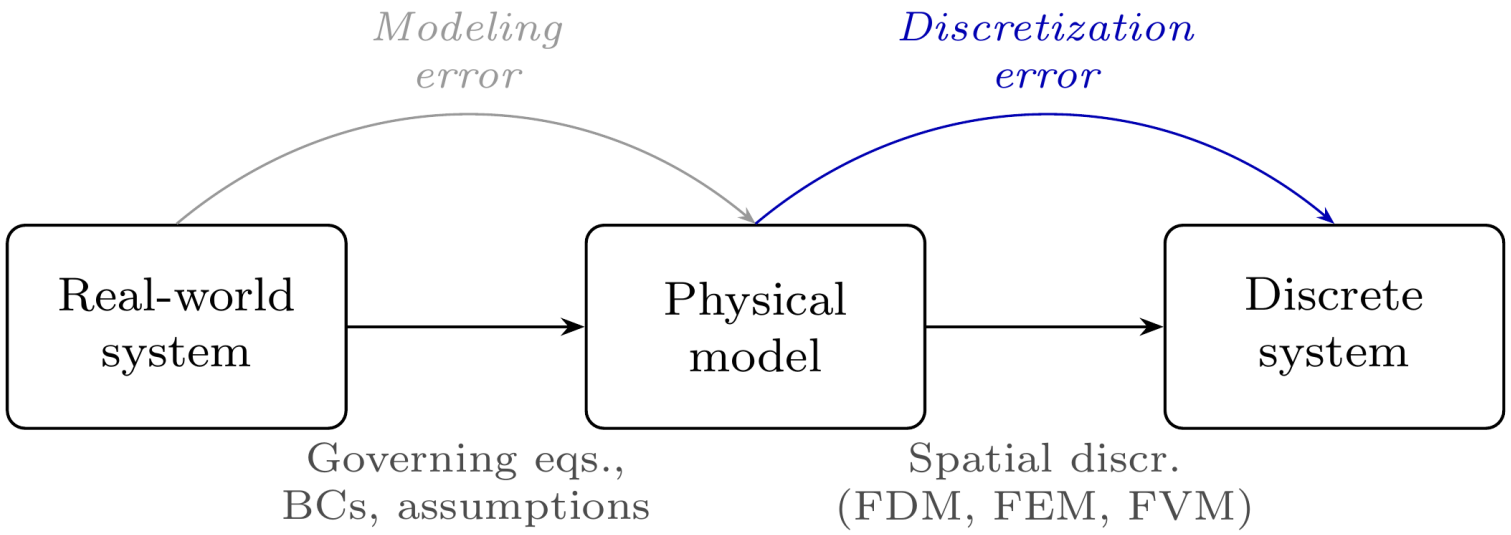

Simulation Pipeline

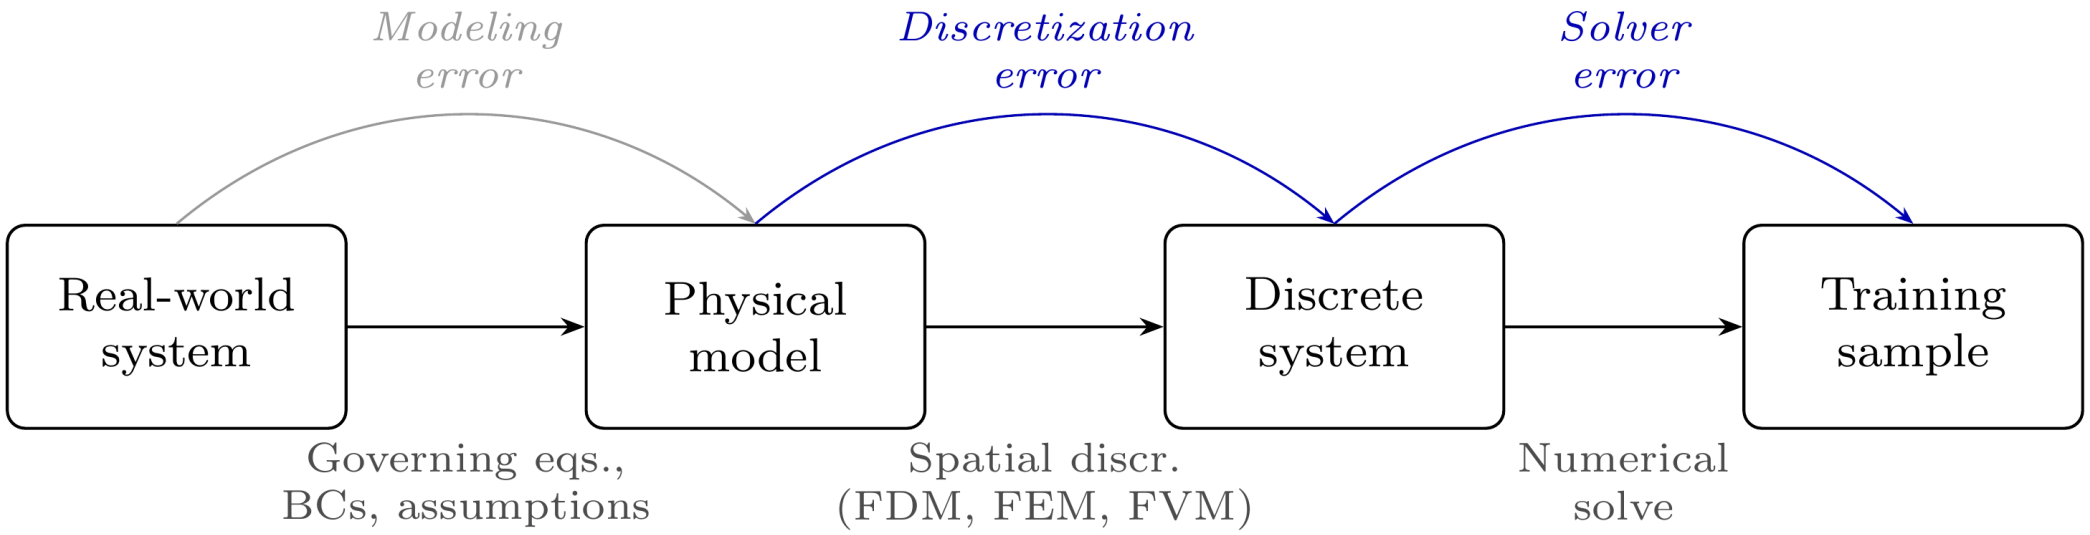

Cost Accuracy Trade-off

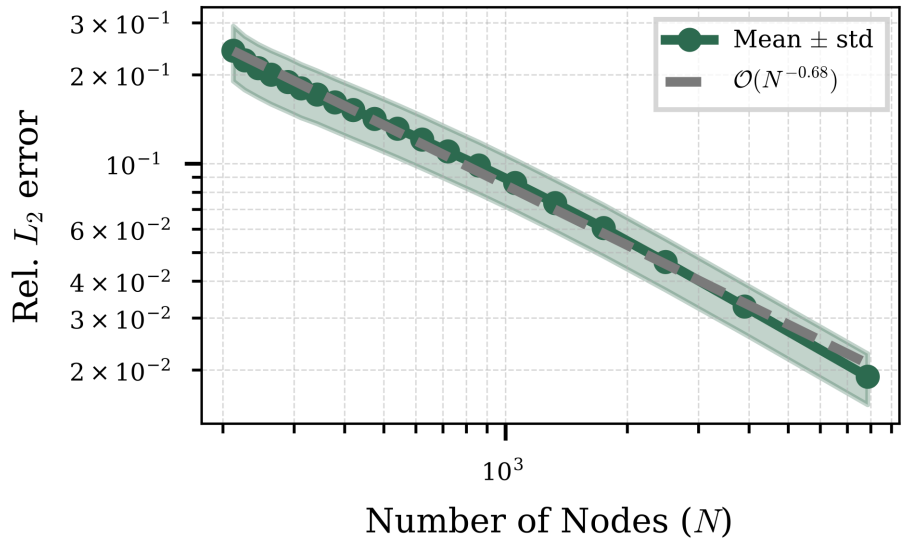

Cost Accuracy Trade-off

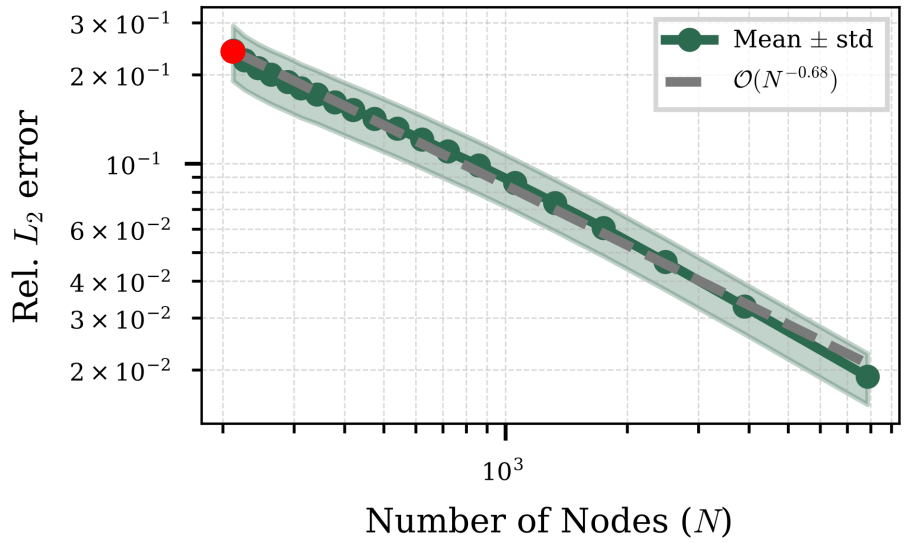

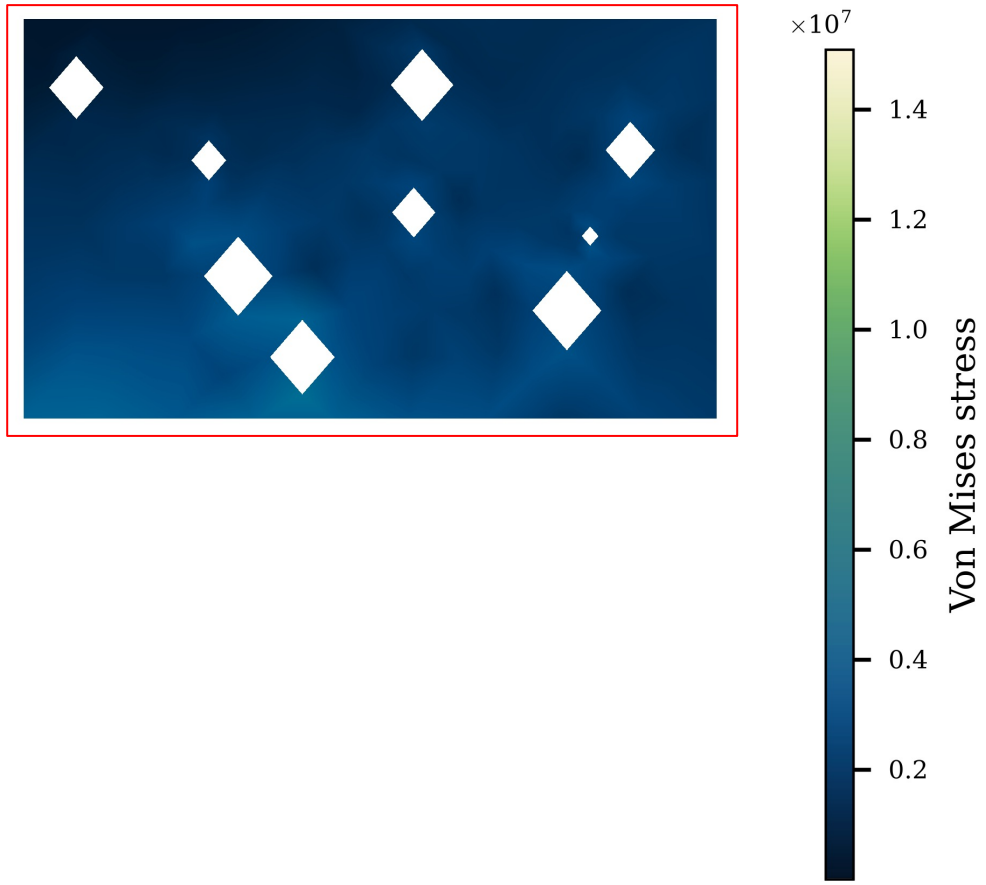

Cost Accuracy Trade-off

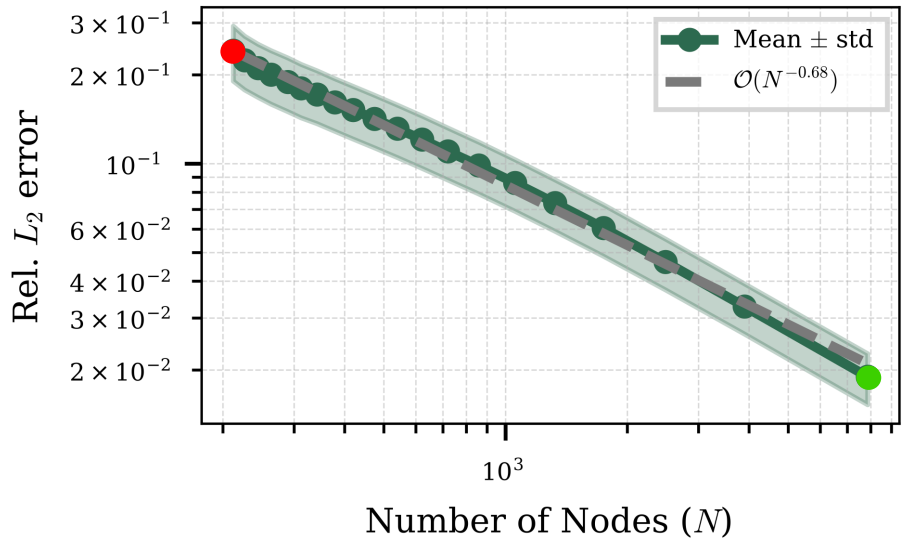

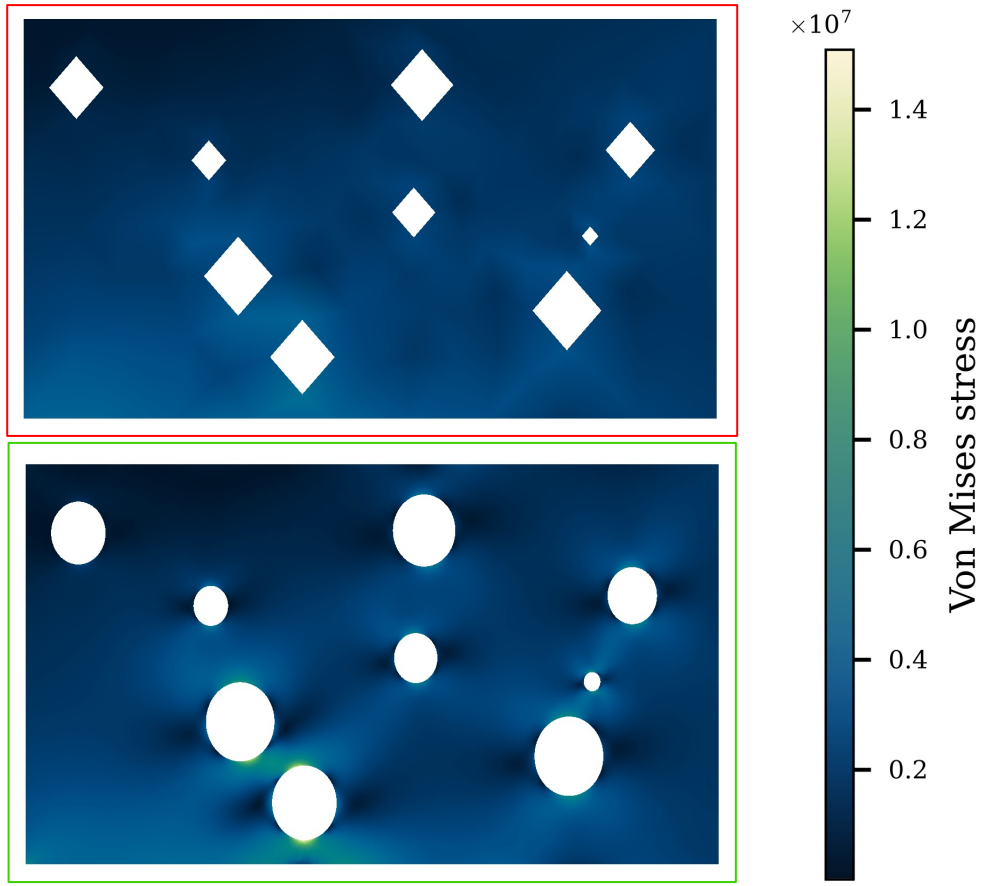

Multi-fidelity Dataset Composition

$$p(\hat{\xi}_i; \lambda, \Xi) = \frac{\exp(\lambda\, \hat{\xi}_i)}{\sum_{j=1}^{|\Xi|} \exp(\lambda\, \hat{\xi}_j)}$$

Multi-fidelity Dataset Composition

$$p(\hat{\xi}_i; \lambda, \Xi) = \frac{\exp(\lambda\, \hat{\xi}_i)}{\sum_{j=1}^{|\Xi|} \exp(\lambda\, \hat{\xi}_j)}$$

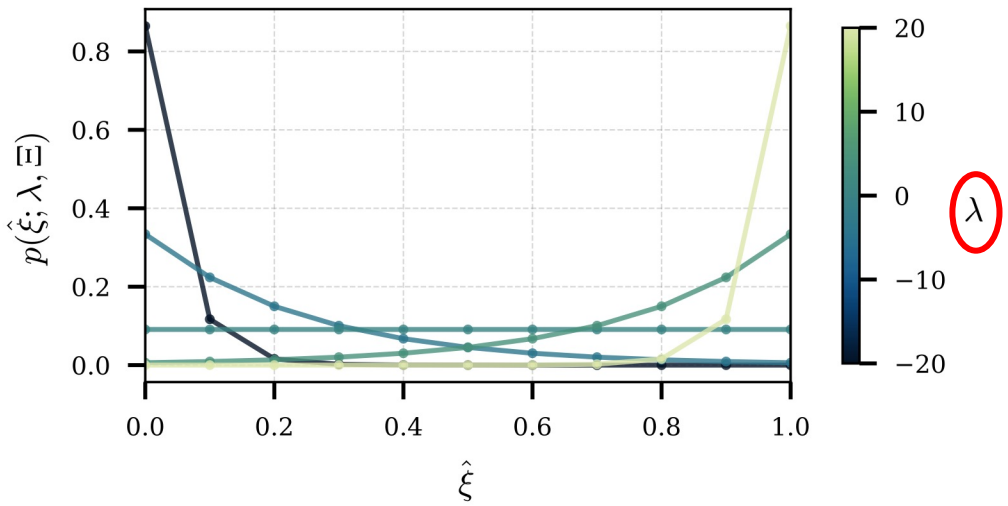

Multi-fidelity Dataset Composition

$$p(\hat{\xi}_i; \lambda, \Xi) = \frac{\exp(\lambda \hat{\xi}_i)}{\sum_{j=1}^{|\Xi|} \exp(\lambda \hat{\xi}_j)}$$

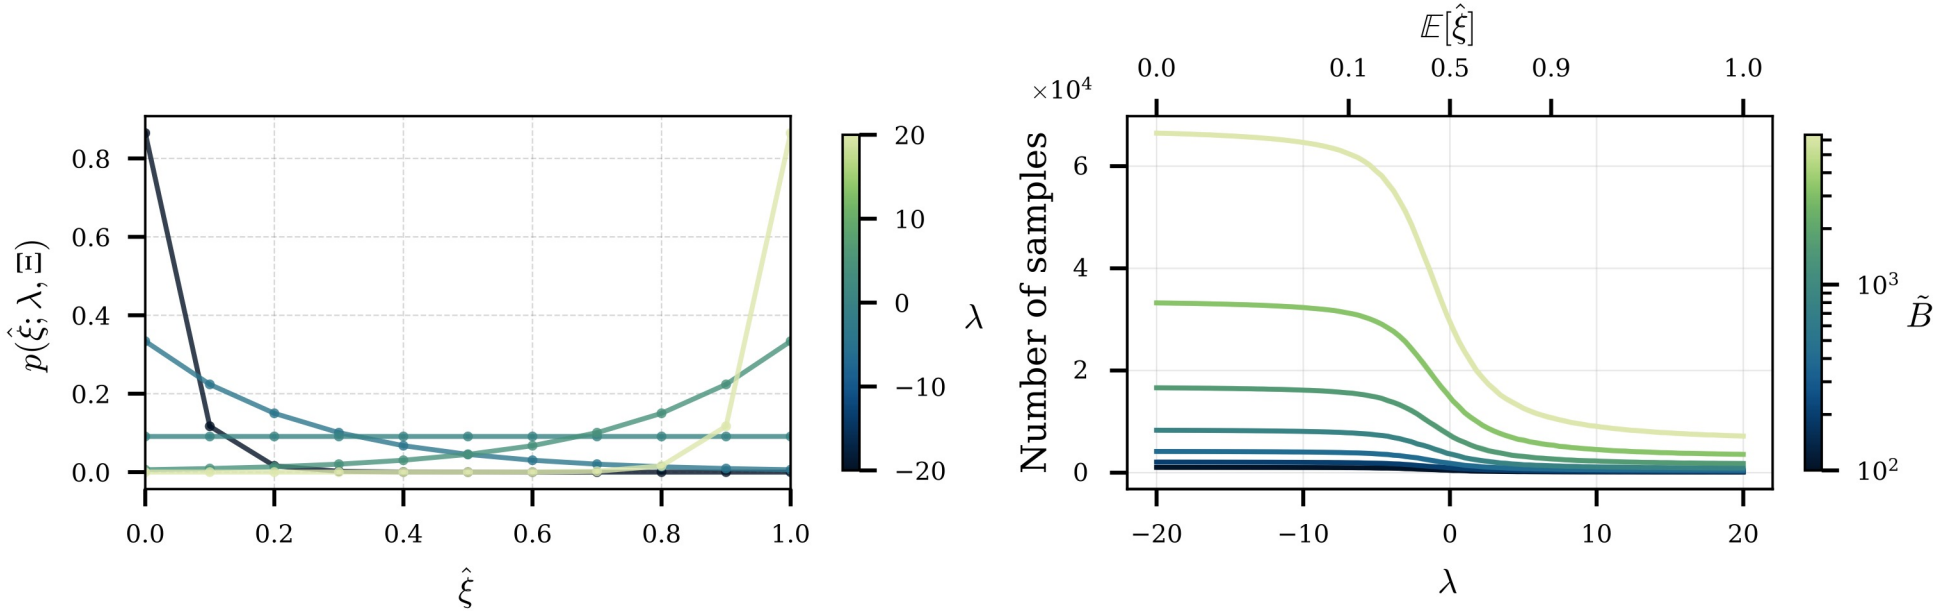

Learning Objective & Test Measure

Training

$$\hat{\theta}(\lambda, B) = \arg\min_{\theta} \; \mathbb{E}_{(\mathbf{f}_h, \xi) \sim D(\lambda, B)} \left[\zeta \left(\mathcal{S}_\theta \left(\mathbf{f}_h; \hat{\xi}(\xi) \right), \; \mathcal{P}_h(\mathbf{f}_h; \xi) \right) \right]$$

Learning Objective & Test Measure

Training

$$\hat{\theta}(\lambda, B) = \arg\min_{\theta} \; \mathbb{E}_{(\mathbf{f}_h, \xi) \sim D(\lambda, B)} \left[\zeta \left(\mathcal{S}_{\theta} \left(\mathbf{f}_h; \hat{\xi}(\xi) \right), \; \mathcal{P}_h(\mathbf{f}_h; \xi) \right) \right]$$

Testing

$$\mathcal{E}(\lambda, B) = \mathbb{E}_{\mathbf{f}_h \sim D_{\text{test}}} \left[\tilde{\zeta} \left(\mathcal{S}_{\hat{\theta}(\lambda, B)} (\mathbf{f}_h; \hat{\xi}(\xi_{\max}), \; \mathcal{P}_h(\mathbf{f}_h; \xi_{\max}) \right) \right]$$

Mixing fidelities beats high-fidelity only

Poisson 3D

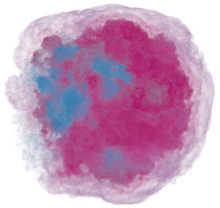

Airfoil



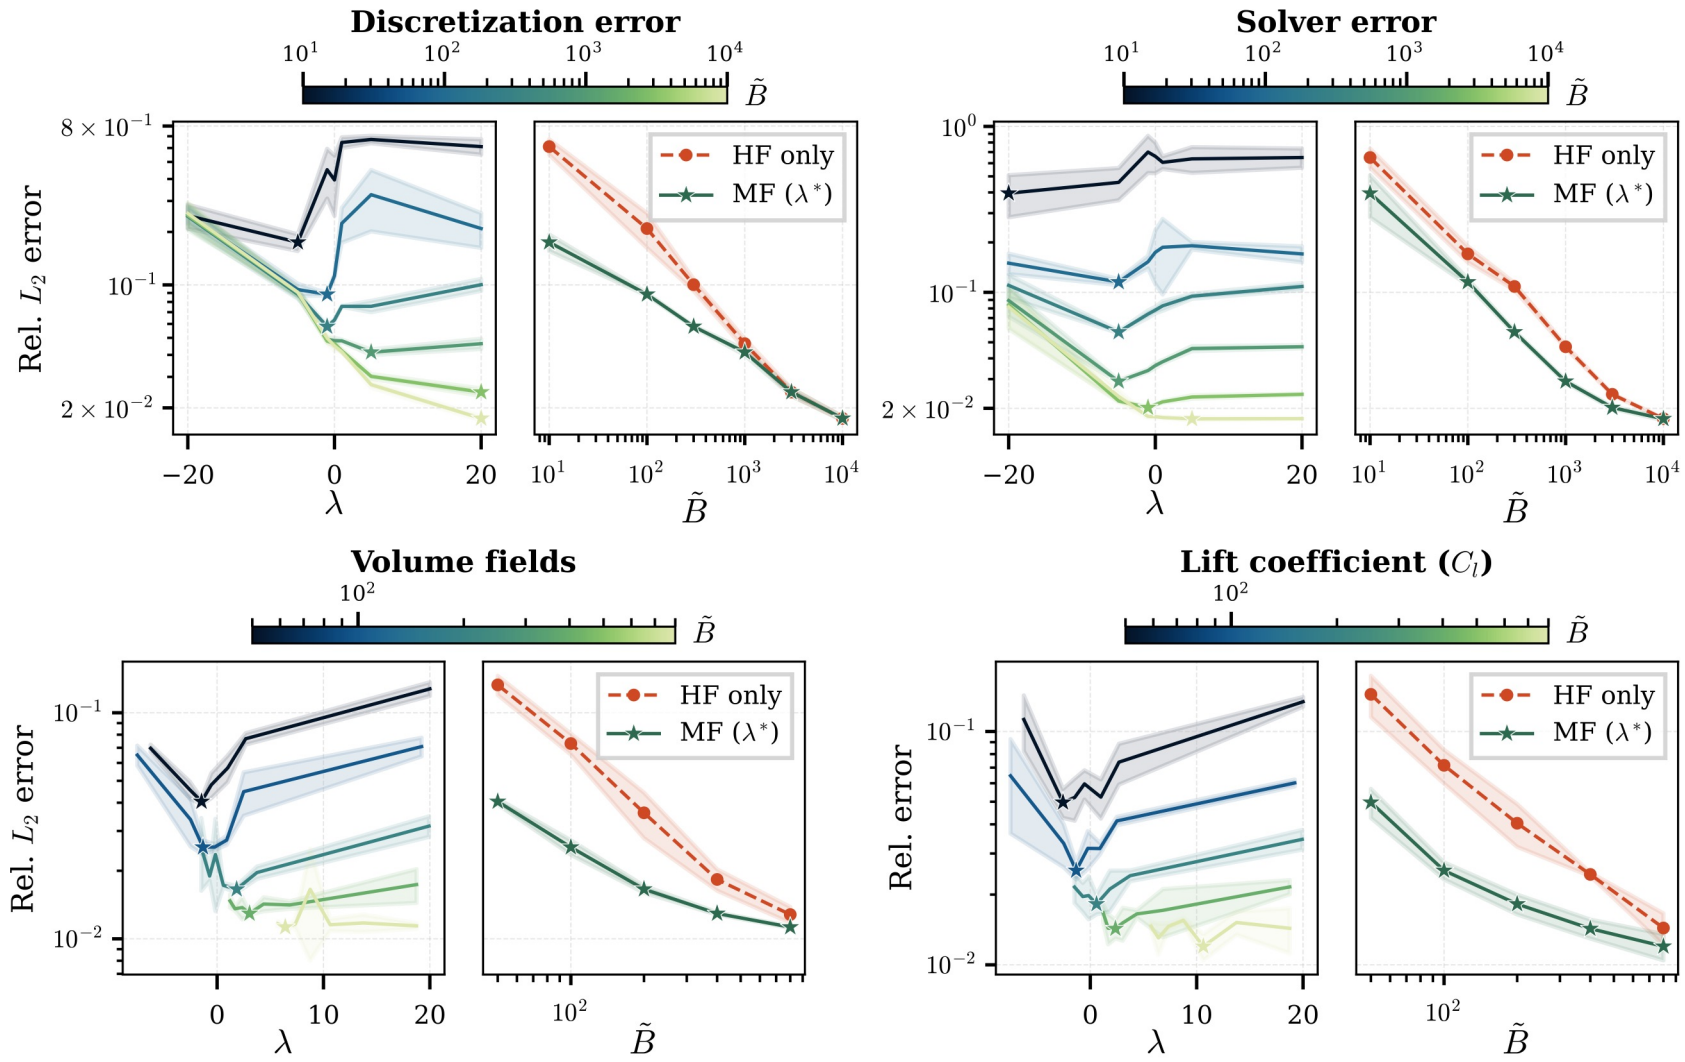

JOHANNES KEPLER UNIVERSITY LINZ

Mixing fidelities beats high-fidelity only

Poisson 3D

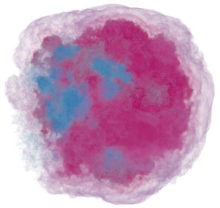

Airfoil



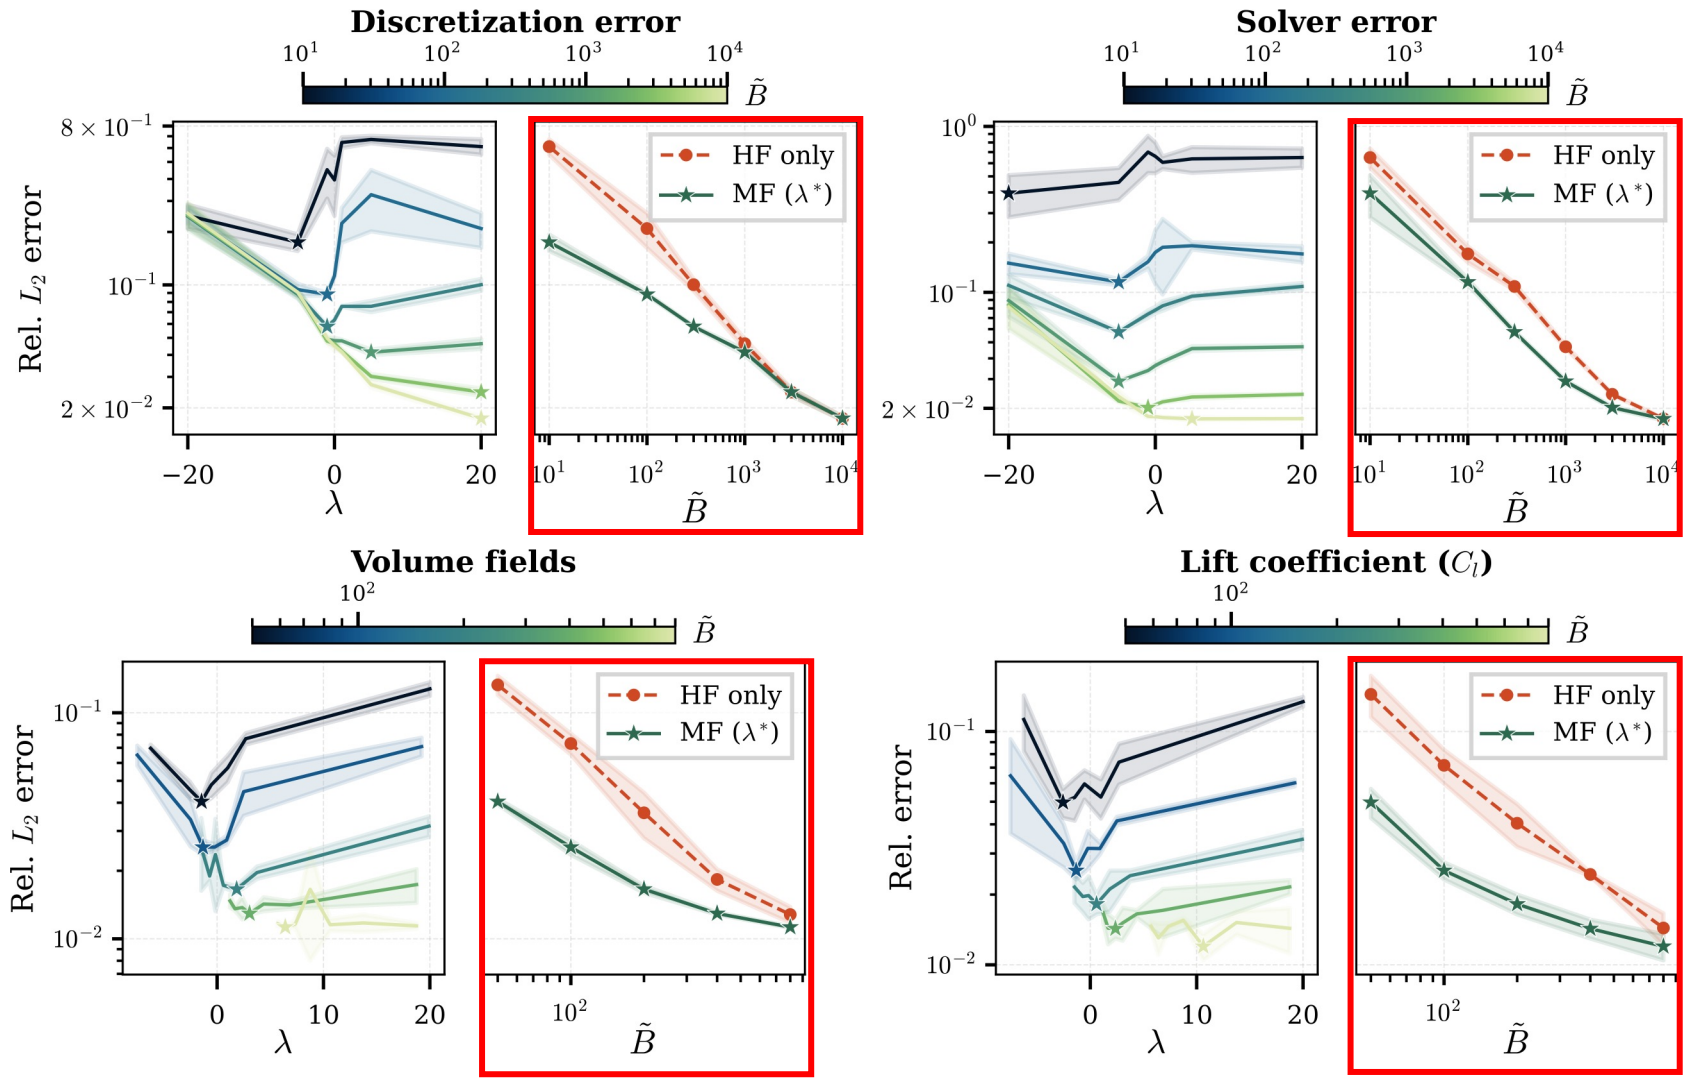

JOHANNES KEPLER UNIVERSITY LINZ

Savings Overview

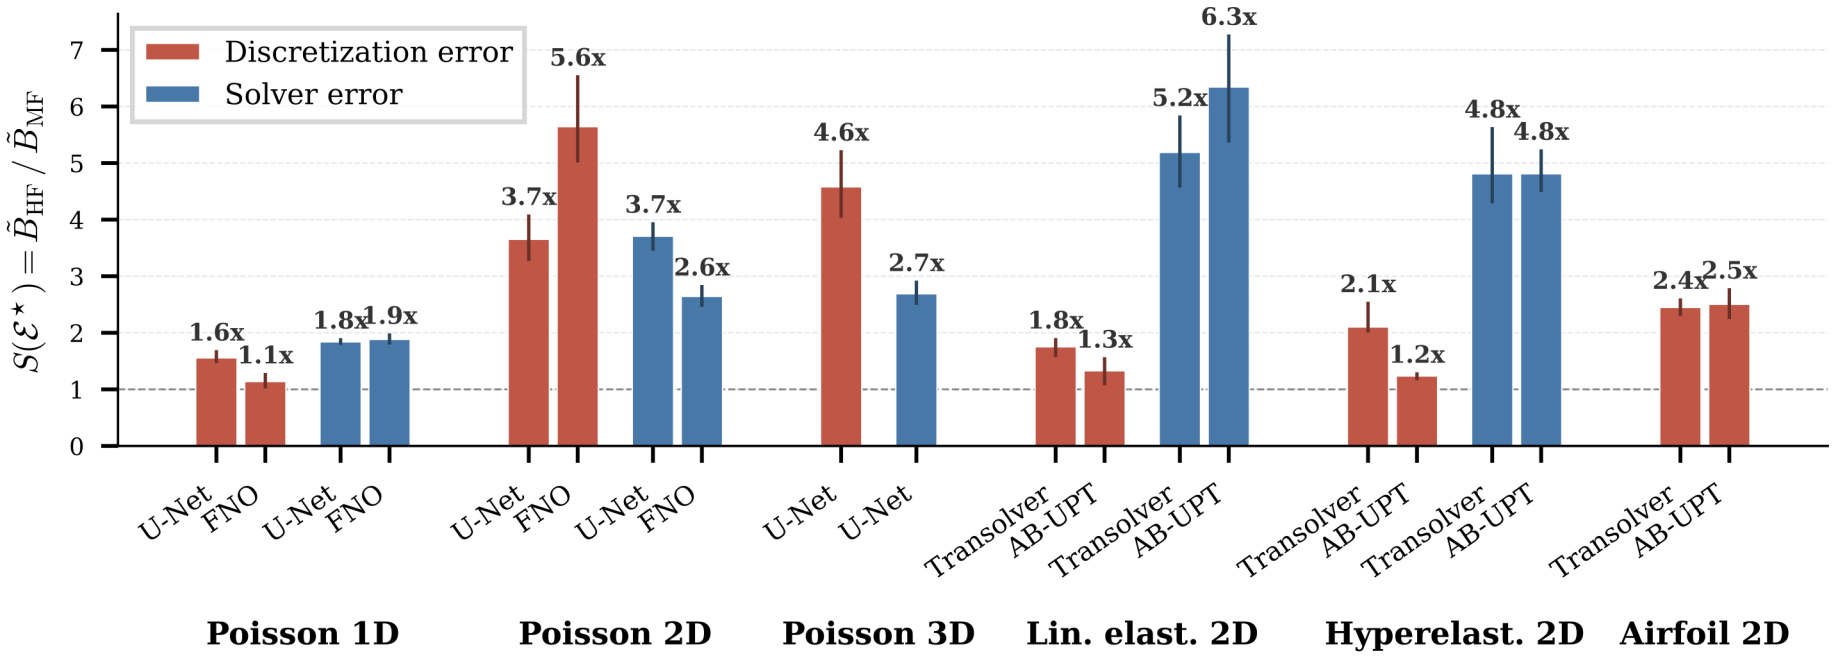

Conclusions

- Multi-fidelity training matches high-fidelity-only errors while enabling budget savings
- Gains are largest in **compute restricted regimes** (~100-500 HF sample equivalent)
- **Solver axis:** large savings and reliable

Thu, Dec 8 5:00 – 8:00 PM
Hall 1-4

Project Page

GitHub

Paper

Data