SIMSHIFT: A BENCHMARK FOR ADAPTING NEURAL SURROGATES TO DISTRIBUTION SHIFTS

Anonymous authorsPaper under double-blind review

000

001

002003004

010 011

012

013

014

016

017

018

019

021

023

025

026027028

029

031

033

034

035

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Neural surrogates for Partial Differential Equations (PDEs) often suffer significant performance degradation when evaluated on unseen problem configurations, such as new initial conditions or structural dimensions. Meanwhile, Domain Adaptation (DA) techniques have been widely used in vision and language processing to generalize from limited information about unseen configurations. In this work, we address this gap through two focused contributions. First, we introduce SIMSHIFT, a novel benchmark dataset and evaluation suite composed of four industrial simulation tasks spanning diverse processes and physics: hot rolling, sheet metal forming, electric motor design and heatsink design. Second, we extend established DA methods to state-of-the-art neural surrogates and systematically evaluate them. These approaches use parametric descriptions and ground truth simulations from multiple source configurations, together with only parametric descriptions from target configurations. The goal is to accurately predict target simulations without access to ground truth simulation data. Extensive experiments on SIMSHIFT highlight the challenges of out-of-distribution neural surrogate modeling, demonstrate the potential of DA in simulation, and reveal open problems in achieving robust neural surrogates under distribution shifts in industrially relevant scenarios.

1 Introduction

PDE simulations are essential tools for understanding and predicting physical phenomena in engineering and science (Evans, 2010). Over recent years, machine learning has emerged as a novel and promising modeling option for complex systems (Brunton & Kutz, 2020), significantly accelerating and augmenting simulation workflows across diverse applications, including weather and climate forecasting (Pathak et al., 2022; Bodnar et al., 2025), material design (Merchant et al., 2023; Zeni et al., 2025) and protein folding (Abramson et al., 2024) to name a few.

In practice, however, models are often deployed outside of their training distribution. This *distribution shift* (Quionero-Candela et al., 2009; Wang et al., 2023) often leads to a significant performance degradation (Bonnet et al., 2022; Herde et al., 2024). A well known analogue is clinical microscopy: models trained with data collected at a few hospitals often fail when deployed at others because microscopes, staining protocols, and lighting conditions differ (Tellez et al., 2019; Koh et al., 2020). For neural surrogates an analogous "instrument shift" arises from new initial conditions, such as material parameters or mesh geometries not encountered during training. Robustness to distribution shifts is crucial for industrial adoption and deployment also because it is becoming a compliance requirement, as stated by Article 15 of the EU AI Act (European Union, 2024).

While methods for increasing out-of-distribution performance have been at the center of research for a long time (Ben-David et al., 2006; Shimodaira, 2000; Sugiyama et al., 2007b), to the best of our knowledge, no benchmark systematically investigates such methods on simulation tasks. Addressing this gap is particularly relevant in scientific and industrial settings, where generating ground truth simulation data is costly, limiting the diversity of training configurations. In contrast, parametric descriptions, such as material types or structural dimensions, are often readily available or easy to generate. This problem is known as *Unsupervised Domain Adaptation (UDA)* (Ben-David et al., 2010), where parametric (input) descriptions and full simulation outputs are available for each *source* configuration, while only input descriptions are provided for *target* configurations, without corresponding outputs.

055

056

059 060 061

062

063 064 065

066

067

068

069

071

073

074

075

076

077 078

079

080

081

082

084

085

087 088

091

092

094

095

096

098

100

101

102

103

104 105

106

107

Figure 1: Schematic overview of the SIMSHIFT framework. In training, the model has access to inputs (e.g., parameters and meshes), corresponding outputs (x,y) from the source domain (left, blue), and only inputs x' from the target domain (right, yellow) are available. The neural operator g and the conditioning network ϕ are shared across domains and jointly optimized. Two loss terms are used: $\mathcal{L}_{\text{recon}}$, computed on source labels, and \mathcal{L}_{DA} , which aligns source and target ϕ features. After training, unsupervised model selection strategies choose θ_{k1} , which is expected to perform best on target domain.

To investigate the potential of UDA for neural surrogate modeling, we provide simulation data across a range of realistic tasks from industrial engineering design. We introduce a comprehensive benchmark that evaluates established UDA methods and neural surrogates. An overview of the framework is shown in Figure 1. Our contributions can be summarized as follows:

- We propose four practical datasets with flexible distribution shifts in *hot rolling*, *sheet metal forming*, *electric motor*, and *heatsink* design, based on realistic simulation setups.
- We present, to the best of our knowledge, the first joint study of established neural surrogate architectures and UDA on engineering simulations with unstructured meshes.
- We introduce *SIMSHIFT*, a modular benchmarking suite that complements our datasets with baseline models and algorithms. It allows easy integration of new simulations, machine learning methods, domain adaptation techniques, and model selection strategies.

2 RELATED WORK

Unsupervised Domain Adaptation. UDA research covers a wide spectrum of results from theoretical foundations (Ben-David et al., 2010; Zellinger et al., 2021a) to modern deep learning methods (Liu & Xue, 2021; Zellinger et al., 2019; Zhu et al., 2021; Long et al., 2018). A prominent class of methods, dubbed as representation learning, aims to map the data to a feature space, where source and target representations appear similar, while maintaining enough information for accurate prediction. To enforce feature similarity between domains, algorithms often employ statistical (Sun & Saenko, 2016; Gretton et al., 2006; Zhang et al., 2019; Shalit et al., 2017) or adversarial (Ganin et al., 2015; Tzeng et al., 2017) discrepancy measures. One crucial yet frequently overlooked factor in the success of UDA methods is model selection. Numerous studies underline the critical impact of hyperparameter choices on UDA algorithm performance, often overshadowing the adaptation method itself (Musgrave et al., 2021; Zellinger et al., 2021b; Dinu et al., 2023; Yang et al., 2024). Even more, since labeled data is unavailable in the target domain, standard validation approaches become infeasible. Thus, it is essential to jointly evaluate adaptation algorithms alongside their associated unsupervised model selection strategies. In this work, we focus on importance weighting strategies (Sugiyama et al., 2007a; You et al., 2019), which stand out by their general applicability, theoretical guarantees and high empirical performance.

Benchmarks for Unsupervised Domain Adaptation. Numerous different Numerous benchmark datasets and evaluation protocols have been established for UDA methods across various machine learning domains, including computer vision (Venkateswara et al., 2017; Peng et al., 2018; Arjovsky et al., 2019), natural language processing (Blitzer et al., 2007), timeseries data (Ragab et al., 2022)

and tabular data (Gardner et al., 2023). However, to the best of our knowledge, systematic UDA benchmarking for neural surrogates remains unexplored.

Benchmarks for Neural Surrogates. Recent years have seen a surge of surrogates belonging to the group of neural operators (see Appendix A), and benchmarks have grown alongside them. However, designing a robust and fair benchmark in the realm of PDEs is difficult and the current literature is not without shortcomings (Brandstetter, 2025). Many focus on solving PDEs on structured, regular grids (Gupta & Brandstetter, 2022; Takamoto et al., 2022; Ohana et al., 2024), which serve as valuable platforms for developing and testing new algorithms. However, these overlook the irregular meshes commonly used in large scale industrial simulations. In that direction, other benchmarks extend to Computational Fluid Dynamics (CFD) on irregular static meshes for airfoil simulations (Bonnet et al., 2022), aerodynamics for automotive (Elrefaie et al., 2024a;b), more academic fluid problems (Luo et al., 2023), and even particle based Smoothed Particle Hydrodynamics simulations (Toshev et al., 2023; 2024). Finally, and most closely related to our work, recent efforts have explored the application of Active Learning techniques (Cohn et al., 1996; Ren et al., 2021) to neural surrogates, introducing a benchmark specifically designed for scenarios where data is scarce (Musekamp et al., 2025). Despite these contributions, all current benchmarks often fall short when addressing a critical issue: the significant performance drop models exhibit under distribution shifts, i.e., when encountering simulation configurations beyond their training setting (Quionero-Candela et al., 2009).

3 DATASET PRESENTATION

Our datasets follow three design principles. (i) Industry relevance: They reflect practical, real-world simulation use-cases. The benchmark covers a diverse set of problems, including 2D as well as 3D cases. (ii) Parametrized conditions: The behavior of all simulations depends on the set of initial parameters only. (iii) Steady-state scenarios: We constrain them to time independent problems, being the standard use case in industry. Take for example design optimization tasks: most rely on either steady-state or time-averaged solutions rather than detailed transient dynamics. This is not just a modeling convenience, but reflects how simulation is integrated into design pipelines: numerical simulations are used to assess candidates by computing scalar objective values. This practice is well documented established various application areas, including thermal systems (Majumdar, 2021), aerodynamic shape optimization for aircrafts (Martins, 2022), wind turbine design (Martins, 2022), and car aerodynamics (Dumas, 2007). Additionally with this constraint we avoid additional complexities such as autoregressive error accumulation in neural surrogates (Lippe et al., 2023).

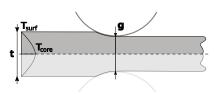
The datasets were generated using the commercial Finite Element Method (FEM) software *Abaqus*, the open-source simulation software *HOTINT* and the open-source CFD package *OpenFoam 9.* ¹ An overview of each dataset together with its most important parameters and a custom metric, motivated by engineering practice, is presented in Sections 3.1 to 3.4. Additionally, we provide detailed descriptions of the respective numerical simulations in Appendix F. Since the behavior of each simulation task is entirely determined by its input parameters, we predefine source and target domains by partitioning the parameter space into distinct, non-overlapping regions. A detailed explanation of the domain splitting strategy is provided in Section 3.5. Table 1 summarizes key characteristics of each dataset, including physical dimensionality, mesh resolution, number of conditioning parameters, and total dataset size. All datasets are publicly hosted on Hugging Face².

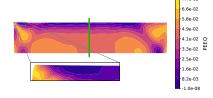
Table 1: Overview of the benchmark datasets. Heatsink meshes are subsampled to a fourth of their original size. Detailed and descriptions of the parameter sampling ranges can be found in Appendix F.

Dataset	Origin	Samples	Output channels	Avg. # nodes	Varied simulation parameters	Dim	Size (GB)
Rolling	Metallurgy	4,750	10	576	4	2D	0.5
Forming	Manufacturing	3,315	10	6,417	4	2D	4.1
Motor	Machinery	3,196	26	9,052	15	2D	13.4
Heatsink	Electronics	460	5	1,385,594	4	3D	40.8

¹Abaqus; HOTINT; OpenFoam 9.

²https://huggingface.co/datasets/simshift/SIMSHIFT_data





(a) Illustration of the simulation setup. The parameters correspond to those in Table 14. We use symmetry constraints and only simulate half of the slab.

(b) Metal slab after the process, showing PEEQ as a contour plot. The green line indicates the center cord, along which we measure the custom metric.

Figure 2: Overview of the *hot rolling* simulation scenario.

3.1 Hot Rolling

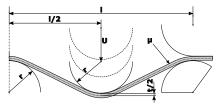
Problem Description. The *hot rolling* process plastically deforms a metal slab into a sheet metal product, as visualized in Figure 2. This complex thermo-mechanical operation involves coupled elasto-plastic deformation and heat transfer phenomena (Gupta, 2021; Galantucci & Tricarico, 1999; Jo et al., 2023). The Finite Element (FE) simulation models the progressive thickness reduction and thermal evolution of the material as it passes through a rolling gap, incorporating temperature-dependent material properties and contact between the slab and the rolls. Among the output fields, the key quantity is Equivalent Plastic Strain (PEEQ), representing the material's plastic deformation, visualized in Figure 2b. The custom metric measures the relative error of the PEEQ profile along the slab's vertical center cord (green line in Figure 2b).

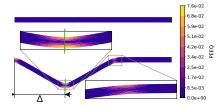
Input parameters are the initial slab thickness t, temperature characteristics $T_{\rm core}$ and $T_{\rm surf}$ of the slab, as well as the geometry of the roll gap. To vary the slab deformation we define the thickness reduction as a percentage of the initial thickness: reduction $=\frac{t-g}{t}$, where g is the rolling gap distance. Table 14 in Appendix F.1 shows a detailed overview of the parameter values together with their sampling ranges used to generate the dataset.

3.2 Sheet Metal Forming

Problem Description. The *sheet metal forming* process is a critical manufacturing operation widely used across industries such as automotive and aerospace. FEM simulations are commonly employed to estimate critical quantities such as thinning, local plastic deformation and residual stress distribution (Tekkaya, 2000; Ablat & Qattawi, 2017; Folle et al., 2024). The simulation setup consists of a symmetrical workpiece supported at the ends and center, a holder and a punch that deforms the sheet by applying a displacement (U in Figure 3a). The 2D simulation predicts the sheet's elasto-plastic deformation, providing quantities such as stress, elastic and plastic strain distributions (shown in Figure 3b). An essential engineering metric used in practice is the transverse stress (xx-component) distribution along the vertical center cord (green line in Figure 3).

Input parameters include the deformed sheet length l, the sheet thickness t, friction coefficient μ and the radii of the holder, punch, and supports r. Table 15 in Appendix F.2 provides the sampling ranges for data generation.





(a) Illustration of the simulation setup. The parameters correspond to those listed in Table 15.

(b) Material before (top) and after (bottom) the process, shown as PEEQ contours. $\Delta=l/2$

Figure 3: Overview of the *sheet metal forming* simulation scenario.

3.3 ELECTRIC MOTOR DESIGN

Problem Description. The *electric motor design* dataset encompasses a structural FEM simulation of a rotor in electric machinery, subjected to mechanical loading at burst speed. It is motivated by the conflicting design objectives in rotor development: while magnetic performance favors certain rotor topologies to optimize flux paths and torque generation, structural integrity requires designs capable of withstanding centrifugal loads without plastic deformation (Gerlach et al., 2021; Dorninger et al., 2021). The 2D simulation predicts stress and deformation responses due to assembly pressing forces and centrifugal loads, accounting for the rotor's topology, material properties, and rotation speed. The custom metric measures the relative error in Mises stress along the cord shown in green Figure 4.

Input Parameters together with their variations and a detailed technical drawing are omitted from the main body since this case is more complex than the preceding datasets. They are provided in Figure 31 and Table 17, both in Appendix F.3.

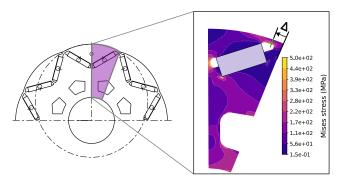


Figure 4: The *electric motor design* simulation scenario, with a schematic sketch of the motor (left) and zoomed-in detail from the simulated radial portion (right). Mises stress field contour plot is shown. The custom error metric is measured along the green line at $\Delta = \frac{t_{rsb1}}{2} + 1.1 * r_{r2}$.

3.4 HEATSINK DESIGN

Problem Description. The *heatsink design* dataset represents a CFD simulation focused on the thermal performance of heat sinks, commonly used in electronic cooling applications (Arularasan & Velraj, 2010; Rahman et al., 2024). It models the convective heat transfer from a heated base through an array of fins to the surrounding air. The simulation captures how geometric fin characteristics, specifically, the number, height, and thickness of fins, affect the overall heat dissipation, along with the temperature of the heat sink. Outputs include steady state temperature, velocity and pressure fields, enabling the assessment of design efficiency and thermal resistance under varying configurations. The main engineering metric measures the relative error in the temperature distribution along the dashed green line in Figure 5.

Input Parameters and their variations as well as an overview of the setup are provided in Appendix F.4.

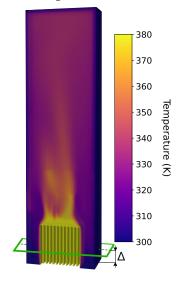


Figure 5: Slice of the *heatsink* 3D temperature field. Custom metric along dashed cord at $\Delta = 0.0025$.

3.5 DISTRIBUTION SHIFTS

SIMSHIFT's functionality allows to generate arbitrary n-dimensional shifts on the parameter spaces of each problem, ensuring flexibility and extensibility. For benchmarking, each dataset includes three predefined distribution shifts: *easy*, *medium* and *hard*, which reflect increasing distributional distance in the respective input spaces (see Table 11 for parameter ranges). The source and target domains are defined by shifting along the dominant input parameter for each simulation scenario, as suggested by domain experts. To validate these choices, we train models across the full parameter

ranges and perform a cluster analysis of their latent representations as the input conditions are varied. The resulting clusters consistently align with the parameter selected by the experts, indicating that the chosen parameters dominate latent space variation (see visualizations in Figures 25 to 28, Appendix C). Besides benchmarking on the predefined shifts, we confirm consistent behavior of models and algorithms in an ablation under a two-dimensional shift on the *electric motor design* dataset in Appendix G.2.

4 BENCHMARK SETUP

This section outlines the learning problem (Section 4.1), the UDA algorithms considered (Section 4.2), the unsupervised model selection strategies (Section 4.3), and the baseline models used (Section 4.4). Finally, we describe the experimental setup and evaluation metrics in Section 4.5.

4.1 LEARNING PROBLEM

Let $\mathcal X$ be an input space containing geometries and conditioning parameters (e.g., thickness and temperatures in Figure 2a) and $\mathcal Y$ be an output space containing ground truth solution fields, obtained from a numerical solver (e.g., PEEQ field in Figure 2b). Following (Ben-David et al., 2010), a domain is represented by a probability density function p on $\mathcal X \times \mathcal Y$ (e.g., describing the probability of observing an input-output pair corresponding to the parameter range $r \in [0.01, 0.115)$ in Table 11). UDA has been formulated as follows: Given a source dataset $(x_1, y_1), ..., (x_n, y_n)$ drawn from a source domain p_S together with an unlabeled target dataset $x_1', ..., x_m'$ drawn from the ($\mathcal X$ -marginal) of a target domain p_T , the problem is to find a model $f: \mathcal X \to \mathcal Y$ that has small expected risk on the target domain:

$$\mathbb{E}_{(x,y)\sim p_T}[\ell(f(x),y)],\tag{1}$$

with $\ell: \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ being some loss function. For example, consider the square loss $\ell(f(x), y) = (f(x) - y)^2$. In our setup $f(x) = g(x, \phi(x))$ is composed of a conditioning network ϕ and a surrogate g (see Figure 1).

4.2 Unsupervised Domain Adaptation Algorithms

Our UDA baseline algorithms are from the class of *domain-invariant representation learning* methods. These methods are strong baselines, in the sense that their performance typically lies within the standard deviation of the winning algorithms in large scale empirical evaluations (i.e., no significant outperformance is observed), see CMD, Deep CORAL and DANN in (Dinu et al., 2023, Tables 12–14), M3SDA in (Peng et al., 2019), MMDA and HoMM in (Ragab et al., 2022).

Following Johansson et al. (2019) and Zellinger et al. (2021b), we express the objective of domain-invariant learning using two learning models: a representation mapping $\phi \in \Phi \subset \{\phi: \mathcal{X} \to \mathcal{R}\}$, which in our case corresponds to the conditioning network that maps simulation parameters into some representation space $\mathcal{R} \subset \mathbb{R}^k$ and a regressor $g \in \mathcal{G} \subset \{g: \mathcal{X} \times \mathcal{R} \to \mathcal{Y}\}$, which is realized by a neural surrogate. The goal is to find a mapping ϕ under which the source representations $\phi(\mathbf{x}) := (\phi(x_1), \dots, \phi(x_m))$ and the target representations $\phi(\mathbf{x}') := (\phi(x_1'), \dots, \phi(x_m'))$ appear similar, and, at the same time, enough information is preserved for prediction by g, see (Quionero-Candela et al., 2009). This is realized by estimating objectives of the form

$$\min_{g \in \mathcal{G}, \phi \in \Phi} \mathbb{E}_{(x,y) \sim p_T} [\ell(g(x,\phi(x)), y)] + \lambda \cdot d(\phi(\mathbf{x}), \phi(\mathbf{x}')), \tag{2}$$

where d is a distance between source and target representations and λ is a regularization parameter. Good choices for d in Equation (2) have been found to be the Wasserstein distance (Courty et al., 2017), the Maximum Mean Discrepancy (Baktashmotlagh et al., 2013), moment distances (Sun & Saenko, 2016; Zellinger et al., 2019), adversarially learned distances (Ganin et al., 2015) and other measures of divergence (Johansson et al., 2019; Zhang et al., 2019). Appropriately choosing λ is crucial for high performance (Musgrave et al., 2021; Dinu et al., 2023; Yang et al., 2024), making model selection necessary.

4.3 Unsupervised Model Selection Strategies

Among all algorithm design choices in UDA, model selection has been repeatedly recognized as one of the most crucial (Musgrave et al., 2021; Yang et al., 2024), with sub-optimal choices potentially leading to *negative transfer* (Pan & Yang, 2010). However, classical approaches (e.g., validation set, cross-validation, information criterion) cannot be used due to missing labels and distribution shifts. It is therefore a natural benchmark requirement for UDA to provide also unified model selection strategies in addition to UDA algorithms.

In this work, we rely on Importance Weighted Validation (IWV) (Sugiyama et al., 2007a) and Deep Embedded Validation (DEV) (You et al., 2019) to overcome the two challenges: (i) distribution shift and (ii) missing target labels. These methods rely on the Radon-Nikodým derivative and the covariate shift assumption $p_S(y|x) = p_T(y|x)$ to obtain

$$\mathbb{E}_{(x,y)\sim p_T}[\ell(f(x),y)] = \mathbb{E}_{(x,y)\sim p_S}\left[\frac{p_T(x)p_T(y|x)}{p_S(x)p_S(y|x)}\ell(f(x),y)\right] = \mathbb{E}_{(x,y)\sim p_S}[\beta(x)\ell(f(x),y)]. \tag{3}$$

Equation (3) motivates to estimate the target error by a two step procedure: First, approaching challenge (i) by estimating the density ratio $\beta(x) = \frac{p_T(x)}{p_S(x)}$ from the input data only, and, approaching challenge (ii) by estimating target error by the weighted source error using *labeled* source data.

4.4 BASELINE MODELS

We provide a comprehensive range of machine learning methods, adapted to our conditioned simulation task, organized by their capacity to model interactions across different spatial scales:

Global context models such as PointNet (Qi et al., 2017) incorporate global information into local Multi-Layer Perceptrons (MLPs) by summarizing features of all input points by aggregation into a global representation, which is then shared among nodes. Recognizing the necessity of *local information* when dealing with complex meshes and structures, we include GraphSAGE (Hamilton et al., 2017), a proven Graph Neural Network (GNN) architecture (Scarselli et al., 2009; Battaglia et al., 2018) already used in other mesh based tasks (Pfaff et al., 2020; Bonnet et al., 2022). However, large scale applications of GNNs are challenging due to computational expense (Alkin et al., 2024a) and issues like oversmoothing (Rusch et al., 2023). Finally, to overcome these limitations, we employ *attention based models* (Vaswani et al., 2017). These models typically scale better with the number of points, and integrate both global and local information enabling stronger long-range interactions and greater expressivity. We include Transolver (Wu et al., 2024), a modern neural operator Transformer.

As an alternative categorization, baselines can also be classified by input-output pairings into *point-to-point* and *latent* approaches. The former explicitly encodes nodes, while the latter represents the underlying fields in a latent space and requires queries to retrieve nodes. While all previously mentioned models are *point-to-point*, we also include Universal Physics Transformer (UPT) (Alkin et al., 2024a; Fürst et al., 2025), as an example of a latent field method. UPTs are designed for large problems and offer favorable scaling on big meshes through latent field modeling; however they are better suited for static-mesh scenarios, as they are lacking the notion of point and don't handle deformations out-of-the-box. Therefore we benchmark UPT only on the *heatsink design* dataset.

We provide detailed explanations of all implemented architectures in Appendix D. Our framework explicitly conditions neural operators on configuration parameters. We first embed them using a sinusoidal (sin–cos) encoding and a shallow MLP ϕ to produce a latent representation and then condition the neural operator g by using either concatenation of the latent conditioning vector, FiLM (Perez et al., 2018) or DiT conditioning layers (Peebles & Xie, 2023). As an alternative, we also evaluate replacing ϕ with a geometric mesh encoder that derives the latent representation directly from the input geometry. On the *electric motor design* dataset, this variant performs worse (see Appendix G.1), supporting our design choice.

4.5 EXPERIMENTS AND EVALUATION

Experimental Setup. We benchmark three prominent UDA algorithms (Deep Coral (Sun & Saenko, 2016), CMD (Zellinger et al., 2019) and DANN (Ganin et al., 2015)) in combination with the following four unsupervised model selection strategies: IWV (Sugiyama et al., 2007a),

DEV (You et al., 2019), Source Best (SB) (selecting models based on source domain validation performance) and Target Best (TB) (selecting models based on target simulation data, which is not available in UDA but serves as a lower bound for perfect model selection).

For the baseline neural surrogate models, we evaluate PointNet, GraphSAGE, and Transolver on the *hot rolling*, *sheet metal forming*, and *electric motor design* datasets. Due to memory and runtime constraints on the large scale *heatsink design* dataset, we omit GraphSAGE and instead benchmark UPT alongside PointNet and Transolver.

Experimental Scale. In total, this results in $3_{\text{models}} \times 3_{\text{UDA algorithms}} \times 4_{\text{selection algorithms}} + 3_{\text{unregularized models}} = 39$ configurations per dataset. We perform an extensive sweep over the critical UDA parameter λ and average across four seeds, resulting in 1,200 training runs. Details on architectures, hyperparameters, training setup and normalization, as well as a breakdown of training times are included in Appendices D and E.

Evaluation Metrics. For each dataset, we report the Normalized Root Mean Squared Error (NRMSE) averaged over all output fields, as well as the per field Root Mean Squared Error (RMSE) values (computed on denormalized data), the Euclidean error for deformation predictions and the custom error metrics described in Sections 3.1 to 3.4. Detailed metric definitions are provided in Appendix E.2.

5 BENCHMARKING RESULTS

Table 2 overviews our benchmarking results, showing the best UDA and selection combination per model. Across datasets and architectures, UDA applied together with unsupervised model selection generally leads to a target error reduction, measured by NRMSE averaged across all fields. However, when examining the dataset-specific custom metrics introduced in Sections 3.1 to 3.4 and the individual fields, gains are not uniform, and some methods improve the global loss, while performance on particular fields or regions deteriorates. This pattern suggests that standard methods are a good starting point, but specialized algorithms tailored to high-dimensional regression tasks are needed. Furthermore, the gap between the best UDA + selection configurations and the TB oracle (lower bound on error) indicates that current unsupervised model selection strategies also leave room for improvement. Despite the clear benefits of UDA, no single UDA algorithm or unsupervised selection strategy dominates across all datasets. In addition to this summary, we report full source and target domain metrics across architectures, algorithms, and selection strategies in Tables 3 to 9 in Appendix B.

Table 2: Best performing combination of UDA algorithm and unsupervised model selection for each dataset (*medium* difficulty) and architecture. We also report an oracle with target best (TB) selection, which provides a lower bound on the selection error. Entries show the target domain (N)RMSE. Promotion on the unregularized baseline are shown in parentheses, with improvement indicated as negative values and asterisks marking unstable unregularized baselines. For each dataset, the best configuration (green row) is chosen by the lowest NRMSE across all fields (bold).

Dataset	Model	Best UDA Method + Model Selection	All Fields Normalized Avg (-)	Rel Custom Error	Mises Stress (MPa)	$\begin{array}{c} \text{PEEQ} \\ (\times 10^{-2}) \end{array}$	Deformation (mm)
	PointNet	CMD + SB	0.387 (-0.082)	0.261 (-0.424)	27.922 (+0.311)	2.511 (-0.008)	11.327 (-0.146)
Rolling	GraphSAGE	CMD + IWV	0.228 (-0.137)	0.170 (-0.022)	14.494 (-5.296)	1.563 (-0.550)	4.622 (-1.093)
Koning	Transolver	CMD + SB	0.899 (*)	0.594 (*)	77.712 (*)	5.798 (*)	13.867 (*)
	Oracle (GraphSAGE)	Deep Coral + TB	0.212 (-0.152)	0.158 (-0.035)	13.829 (-5.960)	1.427 (-0.686)	4.547 (-1.168)
	PointNet	Deep Coral + SB	0.182 (-0.044)	1.154 (+0.273)	31.345 (-0.090)	0.150 (-0.009)	2.555 (-0.001)
Forming	GraphSAGE	DANN + IWV	0.329 (-0.047)	5.599 (+3.104)	52.401 (+6.304)	0.269 (-0.003)	2.095 (+0.156)
roming	Transolver	Deep Coral + DEV	0.155 (-0.013)	0.803 (+0.194)	24.395 (+1.381)	0.143 (+0.016)	1.356 (+0.167)
	Oracle (Transolver)	CMD + TB	0.131 (-0.037)	0.794 (+0.185)	20.285 (-2.730)	0.117 (-0.010)	1.023 (-0.165)
	PointNet	Deep Coral + SB	0.313 (-0.084)	0.197 (-0.147)	26.229 (-4.425)	_	1.528 (-0.057)
Motor	GraphSAGE	CMD + SB	0.344 (-0.032)	0.326 (-0.101)	28.918 (-0.540)	-	1.305 (-0.186)
Motor	Transolver	Deep Coral + SB	0.103 (-0.018)	0.091 (-0.023)	7.678 (-0.650)	-	1.296 (-0.201)
	Oracle (Transolver)	Deep Coral + TB	0.102 (-0.019)	0.092 (-0.023)	7.594 (-0.734)	-	1.254 (-0.244)
			All Fields Normalized Avg (-)	Rel Custom Error (-)	Temperature (K)	Velocity (m/s)	Pressure (kPa)
	PointNet	Deep Coral + SB	0.394 (-0.175)	0.031 (-0.020)	17.428 (-3.699)	0.044 (+0.000)	1.037 (-0.842)
	Transolver	Deep Coral + DEV	0.310 (-0.135)	0.009 (-0.001)	8.717 (-1.003)	0.037 (-0.001)	1.334 (-0.353)
Heatsink	UPT	Deep Coral + SB	0.328 (-0.113)	0.014 (+0.001)	12.610 (-0.424)	0.040 (-0.000)	0.866 (+0.049)
	Oracle (Transolver)	Deep Coral + TB	0.310 (-0.135)	0.009 (-0.001)	8.717 (-1.003)	0.037 (-0.001)	1.334 (-0.353)

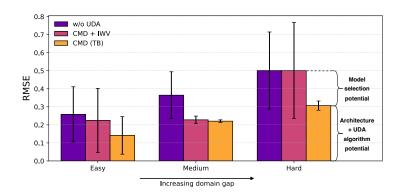


Figure 6: Error scaling with increasing domain gap. We show the averaged RMSE across all (normalized) fields for the *easy*, *medium*, and *hard* gaps on the *hot rolling* task. We compare models without UDA, the best performing UDA method with unsupervised model selection (CMD + IWV), and the theoretical lower bound (TB). Error bars indicate the standard deviation across four seeds. Furthermore, we highlight potentials of architecture, algorithm and selection improvements on the *hard* task.

Finally, since the presented tables only report performance on the *medium* difficulty setting, we additionally visualize model behavior of the best performing combination (model + UDA algorithm + selection strategy) across all difficulty levels of the *hot rolling* dataset in Figure 6. It illustrates the increase in prediction error as the domain gap widens and highlights the consistent improvements achieved by applying UDA algorithms combined with unsupervised model selection strategies on the *easy* and *medium* settings.

For the *hard* setting, however, the shown unsupervised model selection algorithm fails to identify suitable models, as the mean error matches that of the unregularized baselines with the standard deviation even increasing. Nonetheless, the theoretical lower bound (TB) remains substantially below the unregularized error. Figure 6 again highlights the two promising directions for further improvement of the presented baselines: (i) enhancement of neural surrogate architectures and UDA algorithms, and (ii) especially, improvement of unsupervised model selection strategies.

6 Discussion

We presented SIMSHIFT, a collection of industry relevant datasets paired with a benchmarking library for comparing UDA algorithms, unsupervised model selection strategies and neural operators in real word scenarios. We adapt available techniques, apply them on physical simulation data and perform extensive experiments to evaluate their performance on the presented datasets. Our findings suggest that standard UDA training methods can improve performance of neural operators to unseen parameter ranges in physical simulations, with improvement margins in line with those seen in UDA literature (Dinu et al., 2023; Ragab et al., 2022). Additionally, we find correct unsupervised model selection to be extremely important in downstream model performance on target domains, with it arguably having as much impact as the UDA training itself, which is also in agreement with other DA works (Musgrave et al., 2021).

Limitations. We acknowledge that our datasets are limited under two main aspects: (i) They only cover *steady-state* problems, which represent a large portion of industrial simulation tasks. However, an extension with *time-dependent* datasets could be valuable for certain application areas. (ii) They cover a wide range of mesh sizes, ranging from roughly $\mathcal{O}(10^2)$ up to $\mathcal{O}(10^6)$ nodes. Nevertheless, many industrial scenarios require substantially larger meshes. These two design choices prioritize benchmarking clarity and computational feasibility and leave room for future extensions.

REPRODUCIBILITY STATEMENT

The first step towards reproducibility are the datasets. We provide a download link for all dataset with our predefined domain shifts and a high–level description of each dataset and the splitting strategy in Section 3 and Appendix C. To take transparency a step further, we present detailed descriptions of the respective numerical simulations together with the employed solvers, initial/boundary conditions and the configuration details in Appendix F. Concerning the machine learning parts of the benchmark, we provide detailed descriptions of the architectures, their configurations and training hyperparameters in Appendix E.1. Additionally, we provide an anonymized codebase with pinned dependencies, fixable seeds and a comprehensive README.md along with all necessary configuration files used in our benchmarking pipeline.

REFERENCES

- Muhammad Ali Ablat and Ala Qattawi. Numerical simulation of sheet metal forming: a review. *The international journal of advanced manufacturing technology*, 89:1235–1250, 2017.
- Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, pp. 1–3, 2024.
- Benedikt Alkin, Andreas Fürst, Simon Schmid, Lukas Gruber, Markus Holzleitner, and Johannes Brandstetter. Universal physics transformers: A framework for efficiently scaling neural operators. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 25152–25194. Curran Associates, Inc., 2024a. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/2cd36d327f33d47b372d4711edd08de0-Paper-Conference.pdf.
- Benedikt Alkin, Tobias Kronlachner, Samuele Papa, Stefan Pirker, Thomas Lichtenegger, and Johannes Brandstetter. Neuraldem real-time simulation of industrial particulate flows. *CoRR*, abs/2411.09678, 2024b. doi: 10.48550/ARXIV.2411.09678. URL https://doi.org/10.48550/arXiv.2411.09678.
- Martín Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *CoRR*, abs/1907.02893, 2019. URL http://arxiv.org/abs/1907.02893.
- R. Arularasan and R. Velraj. Modeling and simulation of a parallel plate heat sink using computational fluid dynamics. *The International Journal of Advanced Manufacturing Technology*, 51(1):415–419, 2010. ISSN 1433-3015. doi: 10.1007/s00170-008-1867-9. URL https://doi.org/10.1007/s00170-008-1867-9.
- Mahsa Baktashmotlagh, Mehrtash T Harandi, Brian C Lovell, and Mathieu Salzmann. Unsupervised domain adaptation by domain invariant projection. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition*, pp. 769–776, 2013.
- Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19, 2006.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Mach. Learn.*, 79(1-2):151–175, 2010. doi: 10.1007/S10994-009-5152-4. URL https://doi.org/10.1007/s10994-009-5152-4.
- Maurits Bleeker, Matthias Dorfer, Tobias Kronlachner, Reinhard Sonnleitner, Benedikt Alkin, and Johannes Brandstetter. Neuralcfd: Deep learning on high-fidelity automotive aerodynamics simulations. *CoRR*, abs/2502.09692, 2025. doi: 10.48550/ARXIV.2502.09692. URL https://doi.org/10.48550/arXiv.2502.09692.

- John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In Annie Zaenen and Antal van den Bosch (eds.), *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pp. 440–447, Prague, Czech Republic, June 2007. Association for Computational Linguistics. URL https://aclanthology.org/P07-1056/.
- Cristian Bodnar, Wessel P. Bruinsma, Ana Lucic, Megan Stanley, Anna Allen, Johannes Brandstetter, Patrick Garvan, Maik Riechert, Jonathan A. Weyn, Haiyu Dong, Jayesh K. Gupta, Kit Thambiratnam, Alexander T. Archibald, Chun-Chieh Wu, Elizabeth Heider, Max Welling, Richard E. Turner, and Paris Perdikaris. A foundation model for the earth system. *Nat.*, 641(8065):1180–1187, 2025. doi: 10.1038/S41586-025-09005-Y. URL https://doi.org/10.1038/s41586-025-09005-y.
- Florent Bonnet, Jocelyn Ahmed Mazari, Paola Cinnella, and Patrick Gallinari. AirfRANS: High fidelity computational fluid dynamics dataset for approximating reynolds-averaged navier—stokes solutions. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. URL https://arxiv.org/abs/2212.07564.
- Johannes Brandstetter. Envisioning better benchmarks for machine learning pde solvers. *Nature Machine Intelligence*, 7:2–3, January 2025. doi: 10.1038/s42256-024-00962-z. URL https://doi.org/10.1038/s42256-024-00962-z. Published online: 13 December 2024.
- Steven L. Brunton and J. Nathan Kutz. Machine learning for partial differential equations: Data-driven discovery, model reduction, and control. *Journal of Computational Dynamics*, 7(2):343–360, 2020. doi: 10.3934/jcd.2020015.
- Timothy A Burress, Steven L Campbell, Chester Coomer, Curtis William Ayers, Andrew A Wereszczak, Joseph Philip Cunningham, Laura D Marlino, Larry Eugene Seiber, and Hua-Tay Lin. Evaluation of the 2010 toyota prius hybrid synergy drive system. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States). Power..., 2011.
- David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. In *Advances in neural information processing systems*, volume 9, pp. 705–712, 1996.
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, 2017.
- Marius-Constantin Dinu, Markus Holzleitner, Maximilian Beck, Hoan Duc Nguyen, Andrea Huber, Hamid Eghbal-zadeh, Bernhard Alois Moser, Sergei V. Pereverzyev, Sepp Hochreiter, and Werner Zellinger. Addressing parameter choice issues in unsupervised domain adaptation by aggregation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.* OpenReview.net, 2023. URL https://openreview.net/forum?id=M95oDwJXayG.
- Alexander Dorninger, Simon Weitzhofer, Markus Schörgenhumer, Albert Sorgdrager, and Eike Janssen. Automated mechanical rotor design assessment based on 2d fea results. In 2021 11th International Electric Drives Production Conference (EDPC), pp. 1–8, 2021. doi: 10.1109/EDPC53547.2021.9684221.
- Laurent Dumas. Chapter 1 cfd-based optimization for automotive aerodynamics. 2007. URL https://api.semanticscholar.org/CorpusID:601424.
- Mohamed Elrefaie, Angela Dai, and Faez Ahmed. Drivaernet: A parametric car dataset for data-driven aerodynamic design and graph-based drag prediction. volume Volume 3A: 50th Design Automation Conference (DAC) of *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pp. V03AT03A019. Curran Associates, Inc., 08 2024a. URL https://doi.org/10.1115/DETC2024-143593.
- Mohamed Elrefaie, Florin Morar, Angela Dai, and Faez Ahmed. Drivaernet++: A large-scale multimodal car dataset with computational fluid dynamics simulations and deep learning benchmarks. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang

(eds.), Advances in Neural Information Processing Systems, volume 37, pp. 499-536. Curran Associates, Inc., 2024b. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/013cf29a9e68e4411d0593040a8a1eb3-Paper-Datasets_and_Benchmarks_Track.pdf.

- European Union. Artificial intelligence act (eu regulation 2024/1689), article 15: Accuracy, robustness and cybersecurity. http://data.europa.eu/eli/reg/2024/1689/oj, 2024. OJ L, 2024/1689, 12 July 2024.
- Lawrence C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 2nd edition, 2010.
- Luis Fernando Folle, Tiago Nunes Lima, Matheus Passos Sarmento Santos, Bruna Callegari, Bruno Caetano dos Santos Silva, Luiz Gustavo Souza Zamorano, and Rodrigo Santiago Coelho. A review on sheet metal forming behavior in high-strength steels and the use of numerical simulations. *Metals*, 14(12), 2024. ISSN 2075-4701. doi: 10.3390/met14121428. URL https://www.mdpi.com/2075-4701/14/12/1428.
- Nicola Rares Franco, Andrea Manzoni, and Paolo Zunino. Mesh-informed neural networks for operator learning in finite element spaces. *Journal of Scientific Computing*, 97, 2022. URL https://api.semanticscholar.org/CorpusID:258547333.
- Andreas Fürst, Florian Sestak, Artur P. Toshev, Benedikt Alkin, Nikolaus A. Adams, Andreas Mayr, Günter Klambauer, and Johannes Brandstetter. UPT++: Latent point set neural operators for modeling system state transitions. In *ICLR 2025 Workshop on Machine Learning Multiscale Processes*, 2025. URL https://openreview.net/forum?id=u3sEGcYyqw.
- L.M. Galantucci and L. Tricarico. Thermo-mechanical simulation of a rolling process with an fem approach. *Journal of Materials Processing Technology*, 92-93:494-501, 1999. ISSN 0924-0136. doi: https://doi.org/10.1016/S0924-0136(99)00242-3. URL https://www.sciencedirect.com/science/article/pii/S0924013699002423.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks, 2015.
- Josh Gardner, Zoran Popovic, and Ludwig Schmidt. Benchmarking distribution shift in tabular data with tableshift. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/a76a757ed479ale6a5f8134bea492f83-Abstract-Datasets_and Benchmarks.html.
- M.E. Gerlach, M. Zajonc, and B. Ponick. Mechanical stress and deformation in the rotors of high-speed pmsm and im. *Elektrotechnik & Informationstechnik*, 138(2):96–109, 2021. doi: 10.1007/s00502-021-00866-5. URL https://doi.org/10.1007/s00502-021-00866-5.
- Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. A kernel method for the two-sample-problem. In B. Schölkopf, J. Platt, and T. Hoffman (eds.), *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2006. URL https://proceedings.neurips.cc/paper_files/paper/2006/file/e9fb2eda3d9c55a0d89c98d6c54b5b3e-Paper.pdf.
- Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized pde modeling. *arXiv preprint arXiv:2209.15616*, 2022.
- N.K. Gupta. Steel Rolling: Principle, Process & Application. CRC Press, 2021. ISBN 9781032022161. URL https://books.google.at/books?id=wlEgzgEACAAJ.
- William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 1025–1035, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

- Maximilian Herde, Bogdan Raonic, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel de Bezenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for PDEs. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=JC1VKK3UXk.
 - Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, Olivier J. Hénaff, Matthew M. Botvinick, Andrew Zisserman, Oriol Vinyals, and João Carreira. Perceiver IO: A general architecture for structured inputs & outputs. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net, 2022. URL https://openreview.net/forum?id=filj7WpI-g.
 - Seo Yeon Jo, Seojun Hong, Heung Nam Han, and Myoung-Gyu Lee. Modeling and simulation of steel rolling with microstructure evolution: An overview. *steel research international*, 94(2):2200260, 2023. doi: https://doi.org/10.1002/srin.202200260. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/srin.202200260.
 - Fredrik D Johansson, David Sontag, and Rajesh Ranganath. Support and invertibility in domain-invariant representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 527–536. PMLR, 2019.
 - Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. *CoRR*, abs/2012.07421, 2020. URL https://arxiv.org/abs/2012.07421.
 - Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces. *CoRR*, abs/2108.08481, 2021. URL https://arxiv.org/abs/2108.08481.
 - Donald R Lesuer. Experimental investigations of material models for ti-6al-4v titanium and 2024-t3 aluminum. *DOT/FAA/R-00/25*, 2000.
 - Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *CoRR*, abs/2010.08895, 2020a. URL https://arxiv.org/abs/2010.08895.
 - Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. *CoRR*, abs/2003.03485, 2020b.
 - Zongyi Li, Nikola Borislavov Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Prakash Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and Anima Anandkumar. Geometry-informed neural operator for large-scale 3d PDEs. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=86dXbqT5Ua.
 - Phillip Lippe, Bas Veeling, Paris Perdikaris, Richard E. Turner, and Johannes Brandstetter. Pde-refiner: Achieving accurate long rollouts with neural PDE solvers. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/d529b943af3dba734f8a7d49efcb6d09-Abstract-Conference.html.
 - Q. Liu and H. Xue. Adversarial spectral kernel matching for unsupervised time series domain adaptation. Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), 30, 2021.
 - M. Long, Z. Cao, J. Wang, and M. I. Jordan. Conditional adversarial domain adaptation. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.

- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, March 2021. ISSN 2522-5839. doi: 10.1038/s42256-021-00302-5. URL http://dx.doi.org/10.1038/s42256-021-00302-5.
- Xufei Lu, Xin Lin, Michele Chiumenti, Miguel Cervera, JunJie Li, Liang Ma, Lei Wei, Yunlong Hu, and Weidong Huang. Finite element analysis and experimental validation of the thermomechanical behavior in laser solid forming of ti-6al-4v. *Additive Manufacturing*, 21:30–40, 2018.
- Yining Luo, Yingfa Chen, and Zhen Zhang. Cfdbench: A comprehensive benchmark for machine learning methods in fluid dynamics. *CoRR*, abs/2310.05963, 2023. doi: 10.48550/ARXIV.2310.05963. URL https://doi.org/10.48550/arXiv.2310.05963.
- P. Majumdar. Simulation of Thermal Systems, chapter 9, pp. 528-561. John Wiley Sons, Ltd, 2021. ISBN 9781118956922. doi: https://doi.org/10.1002/9781118956922.ch9. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118956922.ch9.
- Joaquim R.R.A. Martins. Aerodynamic design optimization: Challenges and perspectives. Computers Fluids, 239:105391, 2022. ISSN 0045-7930. doi: https://doi.org/10.1016/j.compfluid. 2022.105391. URL https://www.sciencedirect.com/science/article/pii/S0045793022000615.
- Florian Menter, M. Kuntz, and RB Langtry. Ten years of industrial experience with the sst turbulence model. *Heat and Mass Transfer*, 4, 01 2003.
- Amil Merchant, Simon L. Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk. Scaling deep learning for materials discovery. *Nat.*, 624(7990):80–85, 2023. doi: 10.1038/S41586-023-06735-9.
- Daniel Musekamp, Marimuthu Kalimuthu, David Holzmüller, Makoto Takamoto, and Mathias Niepert. Active learning for neural PDE solvers. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=x4ZmQaumRg.
- Kevin Musgrave, Serge J. Belongie, and Ser-Nam Lim. Unsupervised domain adaptation: A reality check. *CoRR*, abs/2111.15672, 2021. URL https://arxiv.org/abs/2111.15672.
- Ruben Ohana, Michael McCabe, Lucas Meyer, Rudy Morel, Fruzsina J. Agocs, Miguel Beneitez, Marsha Berger, Blakesley Burkhart, Stuart B. Dalziel, Drummond B. Fielding, Daniel Fortunato, Jared A. Goldberg, Keiya Hirashima, Yan-Fei Jiang, Rich R. Kerswell, Suryanarayana Maddu, Jonah Miller, Payel Mukhopadhyay, Stefan S. Nixon, Jeff Shen, Romain Watteaux, Bruno Régaldo-Saint Blancard, François Rozet, Liam H. Parker, Miles Cranmer, and Shirley Ho. The well: a large-scale collection of diverse physics simulations for machine learning. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), Advances in Neural Information Processing Systems, volume 37, pp. 44989–45037. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/4f9a5acd91ac76569f2fe291b1f4772b-Paper-Datasets_and_Benchmarks_Track.pdf.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010. doi: 10.1109/TKDE.2009.191.
- Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, and Animashree Anandkumar. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *CoRR*, abs/2202.11214, 2022. URL https://arxiv.org/abs/2202.11214.

- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4196–4206, 2023. URL https://arxiv.org/abs/2212.09748.
 - Xingchao Peng, Ben Usman, Kuniaki Saito, Neela Kaushik, Judy Hoffman, and Kate Saenko. Syn2real: A new benchmark forsynthetic-to-real visual domain adaptation. *CoRR*, abs/1806.09755, 2018.
 - Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.
 - Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pp. 3942–3951, 2018. URL https://arxiv.org/abs/1709.07871.
 - Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W. Battaglia. Learning meshbased simulation with graph networks. *CoRR*, abs/2010.03409, 2020. URL http://dblp.uni-trier.de/db/journals/corr/corr2010.html#abs-2010-03409.
 - Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, pp. 77–85. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.16. URL https://doi.org/10.1109/CVPR.2017.16.
 - Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. *Dataset Shift in Machine Learning*. The MIT Press, 2009. ISBN 0262170051.
 - Mohamed Ragab, Emadeldeen Eldele, Wee Ling Tan, Chuan-Sheng Foo, Zhenghua Chen, Min Wu, Chee Keong Kwoh, and Xiaoli Li. ADATIME: A benchmarking suite for domain adaptation on time series data. *CoRR*, abs/2203.08321, 2022. doi: 10.48550/ARXIV.2203.08321. URL https://doi.org/10.48550/arXiv.2203.08321.
 - Md Atiqur Rahman, S. M. Mozammil Hasnain, Prabhu Paramasivam, and Abinet Gosaye Ayanie. Advancing thermal management in electronics: a review of innovative heat sink designs and optimization techniques. *RSC Adv.*, 14:31291–31319, 2024. doi: 10.1039/D4RA05845C. URL http://dx.doi.org/10.1039/D4RA05845C.
 - Zheng Ren, Yongxin Yang, Bingbing Chen, Yaqing Li, Chengzhong Xu, Timothy M Hospedales, and Tao Wang. A survey of deep active learning. *ACM Computing Surveys (CSUR)*, 54(9):1–36, 2021.
 - T. Konstantin Rusch, Michael M. Bronstein, and Siddhartha Mishra. A survey on oversmoothing in graph neural networks. *arXiv preprint arXiv:2303.10993*, 2023.
 - Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
 - Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the International Conference on Machine Learning*, pp. 3076–3085. PMLR, 2017.
 - Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227–244, 2000. ISSN 0378-3758. doi: https://doi.org/10.1016/S0378-3758(00)00115-4. URL https://www.sciencedirect.com/science/article/pii/S0378375800001154.
 - Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. Covariate shift adaptation by importance weighted cross validation. *J. Mach. Learn. Res.*, 8:985–1005, 2007a. doi: 10.5555/1314498.1390324. URL https://dl.acm.org/doi/10.5555/1314498.1390324.

Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul von Bünau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In John C. Platt, Daphne Koller, Yoram Singer, and Sam T. Roweis (eds.), Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007, pp. 1433–1440. Curran Associates, Inc., 2007b. URL https://proceedings.neurips.cc/paper/2007/hash/be83ab3ecd0db773eb2dc1b0a17836a1-Abstract.html.

- Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation, 2016. URL https://arxiv.org/abs/1607.01719.
- Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 1596–1611. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/0a9747136d411fb83f0cf81820d44afb-Paper-Datasets_and_Benchmarks.pdf.
- A.Erman Tekkaya. State-of-the-art of simulation of sheet metal forming. *Journal of Materials Processing Technology*, 103(1):14–22, 2000. ISSN 0924-0136. doi: https://doi.org/10.1016/S0924-0136(00)00413-1. URL https://www.sciencedirect.com/science/article/pii/S0924013600004131.
- David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen van der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical Image Analysis*, 58: 101544, 2019. ISSN 1361-8415. doi: https://doi.org/10.1016/j.media.2019.101544. URL https://www.sciencedirect.com/science/article/pii/S1361841519300799.
- Artur Toshev, Harish Ramachandran, Jonas A. Erbesdobler, Gianluca Galletti, Johannes Brandstetter, and Nikolaus A. Adams. JAX-SPH: A differentiable smoothed particle hydrodynamics framework. In *ICLR 2024 Workshop on AI4DifferentialEquations In Science*, 2024. URL https://openreview.net/forum?id=8X5PXVmsHW.
- Artur P. Toshev, Gianluca Galletti, Fabian Fritz, Stefan Adami, and Nikolaus A. Adams. Lagrangebench: a lagrangian fluid mechanics benchmarking suite. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, 2023.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7167–7176, 2017.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. *CoRR*, abs/1706.07522, 2017. URL http://arxiv.org/abs/1706.07522.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023.
- Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast transformer solver for pdes on general geometries. In *International Conference on Machine Learning*, 2024.

- Jianfei Yang, Hanjie Qian, Yuecong Xu, Kai Wang, and Lihua Xie. Can we evaluate domain adaptation models without target-domain labels? In *International Conference on Learning Representations*, 2024.
- Kaichao You, Ximei Wang, Mingsheng Long, and Michael I. Jordan. Towards accurate model selection in deep unsupervised domain adaptation. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, 9-15 June 2019, Long Beach, California, USA, volume 97 of Proceedings of Machine Learning Research, pp. 7124–7133. PMLR, 2019. URL http://proceedings.mlr.press/v97/you19a.html.
- Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. Central moment discrepancy (cmd) for domain-invariant representation learning, 2019. URL https://arxiv.org/abs/1702.08811.
- Werner Zellinger, Bernhard A Moser, and Susanne Saminger-Platz. On generalization in moment-based domain adaptation. *Annals of Mathematics and Artificial Intelligence*, 89(3):333–369, 2021a.
- Werner Zellinger, Natalia Shepeleva, Marius-Constantin Dinu, Hamid Eghbal-zadeh, Hoan Duc Nguyen, Bernhard Nessler, Sergei V. Pereverzyev, and Bernhard Alois Moser. The balancing principle for parameter choice in distance-regularized domain adaptation. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, pp. 20798–20811, 2021b. URL https://proceedings.neurips.cc/paper/2021/hash/ae0909a324fb2530e205e52d40266418-Abstract.html.
- Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, et al. A generative model for inorganic materials design. *Nature*, pp. 1–3, 2025.
- Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *Proceedings of the International Conference on Machine Learning*, pp. 7404–7413, 2019.
- Y. Zhu, F. Zhuang, J. Wang, G. Ke, J. Chen, J. Bian, H. Xiong, and Q. He. Deep subdomain adaptation network for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 32(4):1713–1722, 2021.

LLM USAGE DISCLOSURE

In general, LLM tools were used to refine writing in parts of the paper. DeepSeek-R1 and GPT-5 were additionally used to make visualizations prettier, speed up the development of plotting functions, and dump experimental results neatly into latex tabled tables. Beyond that, they were not used to a significant degree in other parts of the code, as neither Copilot nor Cursor are used by the main author. AI assistants were strictly editors and decorators, i.e. they were not involved in ideation, reordering ideas, or at any higher or lower conceptual level.

A ON NEURAL OPERATORS

One prominent approach in neural surrogate modeling for PDEs is operator learning (Kovachki et al., 2021; Li et al., 2020a; Lu et al., 2021; Alkin et al., 2024a; Li et al., 2020b). In this setting, an operator maps input functions, such as boundary or initial conditions, to the corresponding solution of the PDE. During training, neural operators typically learn from input-output pairs of discretized functions (Kovachki et al., 2021; Li et al., 2020a; Lu et al., 2021; Alkin et al., 2024a). While some methods expect regular, grid based inputs (Li et al., 2020a), others can be applied to any kind of data structure (Alkin et al., 2024a; Li et al., 2020b). One notable property is *discretization invariance*, which, along with the ability to handle irregular data, enables generalization across different resolutions and mesh geometries. This is a highly desirable property for industrial simulations (Pfaff et al., 2020; Alkin et al., 2024a; Fürst et al., 2025; Li et al., 2023; Franco et al., 2022), where non-uniform meshes are the standard due to the computational and modeling advantages. In this work, we focus on domain adaptation rather than benchmarking discretization invariance, and include neural surrogates that may not satisfy this property, such as (Hamilton et al., 2017). Such models have been leveraged in several large scale industrial contexts, including CFD for automotive (Bleeker et al., 2025) or Discrete Element Method (DEM) simulations for industrial processes (Alkin et al., 2024b).

B DETAILED RESULTS

Complementing the summary in Table 2 of the main paper, the following sections present detailed results for each dataset. For every dataset, we present a complete empirical evaluation of our benchmark that compares the performance for all combinations of models, UDA algorithms and model selection strategies across all output fields.

While these quantitative metrics offer a high level summary of model performance, industry practitioners often need a more fine grained picture to assess the neural surrogate's capabilities under distribution shifts. To address this, we include additional analyses and visualizations alongside the quantitative results. First, we provide error distribution histograms to better illustrate the difficulty of the domain shift occurring in each dataset. Additionally, we present fringe and scatter plots comparing model predictions with the respective ground truth numerical solutions.

B.1 HOT ROLLING

Table 3 presents the complete benchmarking results for the *hot rolling* dataset.

To gain more insights, we conduct additional analyses on the best performing model, selected based on having the lowest average normalized target domain error across all fields. Figure 7 shows the error distribution of this model and clearly highlights the substantial distribution shift between the source and target domain of the *hot rolling* dataset. Errors in the target domain are noticeably larger, almost up to an order of magnitude higher than those observed in the source domain.

To further illustrate the model's performance, we analyze two representative samples, one from the source and one from the target domain. Since the most critical field for downstream applications is PEEQ, we restrict the following analysis on this scalar field only.

Table 4 presents a summary of the absolute PEEQ prediction errors for the selected source and target samples. Additionally, Figure 8 and Figure 9 visualize the ground truth, predictions, and absolute errors for these samples using fringe plots.

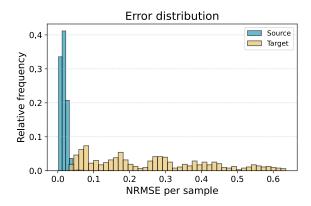


Figure 7: Distribution of NRMSE (averaged across all fields) for the test sets of the source (blue) and target (yellow) domains in the *hot rolling* dataset. Bar height indicates the relative frequency of samples within each bin.

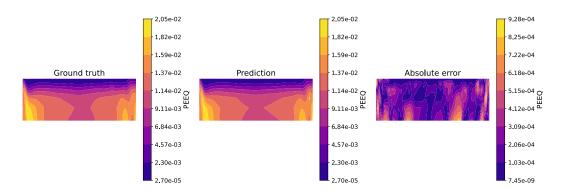


Figure 8: Fringe plot of the *hot rolling* dataset (representative source sample). Shown is the ground truth (left) and predicted (middle) PEEQ, as well as the absolute error (right).

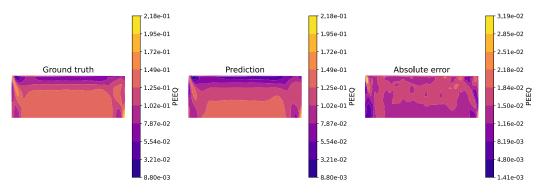


Figure 9: Fringe plot of the *hot rolling* dataset (representative target sample). Shown is the ground truth (left) and predicted (middle) PEEQ, as well as the absolute error (right).

Table 3: RMSE (mean \pm std over 4 seeds) on the *hot rolling* dataset at *medium* difficulty. Values are target domain errors (lower is better). Bold marks the overall best model + UDA algorithm + model selection combination. For each architecture, the unregularized baseline row is shaded beige, whereas the best UDA + selection within that architecture is underlined and shaded green. Asterisks denote unstable runs ($\geq 10 \times$ the column median).

Model	DA	Model	All Fields Nor	malized Avg (-)	Rel Custo	n Error (-)	Deform	nation (mm)	Logarithmic	Strain (×10 ⁻²)	Equivalent Plas	tic Strain (×10 ⁻²)	Mises S	itress (MPa)	Stre	ss (MPa)
.mouci	Algorithm	Selection	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT
			$0.016(\pm0.000)$	$0.365(\pm0.130)$	$0.234(\pm0.081)$	$0.192(\pm 0.000)$	$0.525(\pm0.023)$	$5.715(\pm 1.567)$	$0.018(\pm0.000)$	$0.997(\pm 0.377)$	$0.033(\pm0.000)$	$2.113(\pm 0.789)$	$1.972(\pm0.024)$	19.790(±7.186)	$1.234(\pm 0.010)$	11.421(±3.891)
	DANN DANN DANN DANN	DEV IWV SB TB	0.014(±0.000) 0.014(±0.000) 0.014(±0.000) 0.014(±0.000)	$1.175(\pm 0.053)$ $0.289(\pm 0.147)$ $0.692(\pm 0.511)$ $0.230(\pm 0.041)$	$0.047(\pm0.003)$ $0.042(\pm0.002)$ $0.043(\pm0.004)$ $0.042(\pm0.001)$	$0.726(\pm 0.051)$ $0.183(\pm 0.045)$ $0.466(\pm 0.318)$ $0.170(\pm 0.026)$	$0.577(\pm0.061)$ $0.561(\pm0.032)$ $0.573(\pm0.043)$ $0.604(\pm0.010)$	$17.363(\pm0.803)$ $5.359(\pm1.848)$ $11.090(\pm7.161)$ $4.640(\pm0.593)$	$0.019(\pm 0.001)$ $0.018(\pm 0.000)$ $0.018(\pm 0.000)$ $0.018(\pm 0.001)$	$3.452(\pm 0.176)$ $0.792(\pm 0.186)$ $2.120(\pm 1.506)$ $0.740(\pm 0.134)$	$0.035(\pm0.001)$ $0.033(\pm0.001)$ $0.034(\pm0.001)$ $0.034(\pm0.001)$	7.290(±0.405) 1.622(±0.306) 4.510(±3.201) 1.549(±0.275)	$2.056(\pm 0.050)$ $1.992(\pm 0.037)$ $1.991(\pm 0.045)$ $2.017(\pm 0.047)$	111.626(±7.317) 24.471(±22.423) 60.352(±51.358) 14.867(±3.085)	$1.264(\pm 0.033)$ $1.246(\pm 0.025)$ $1.237(\pm 0.027)$ $1.248(\pm 0.028)$	$59.263(\pm 5.594)$ $13.737(\pm 11.828)$ $31.612(\pm 25.882)$ $8.665(\pm 1.635)$
GraphSAGE	CMD CMD CMD CMD	DEV IWV SB TB	0.015(±0.001) 0.014(±0.000) 0.014(±0.000) 0.014(±0.000)	$\begin{array}{c} 1.447(\pm 0.202) \\ \textbf{0.228}(\pm \textbf{0.021}) \\ \hline 0.786(\pm 0.535) \\ 0.221(\pm 0.007) \end{array}$	0.047(±0.005) 0.043(±0.003) 0.041(±0.001) 0.044(±0.001)	$0.799(\pm 0.069)$ $0.170(\pm 0.012)$ $0.526(\pm 0.327)$ $0.165(\pm 0.003)$	$0.617(\pm0.040)$ $0.577(\pm0.023)$ $0.571(\pm0.040)$ $0.583(\pm0.033)$	18.383(±2.116) 4.622(±0.283) 12.160(±7.174) 4.607(±0.261)	$0.020(\pm 0.001)$ $0.018(\pm 0.000)$ $0.018(\pm 0.000)$ $0.018(\pm 0.000)$	3.781(±0.544) 0.742(±0.071) 2.403(±1.541) 0.711(±0.014)	$0.037(\pm0.003)$ $0.033(\pm0.001)$ $0.033(\pm0.000)$ $0.033(\pm0.000)$	7.764(±1.210) 1.563(±0.153) 5.068(±3.248) 1.507(±0.045)	$2.169(\pm 0.151)$ $1.980(\pm 0.040)$ $1.974(\pm 0.034)$ $1.992(\pm 0.021)$	136.324(±23.104) 14.494(±1.375) 68.509(±55.017) 14.288(±1.040)	1.324(±0.062) 1.237(±0.032) 1.228(±0.024) 1.245(±0.019)	95.502(±20.973) 8.386(±0.819) 36.834(±28.921) 8.275(±0.632)
	Deep Coral Deep Coral Deep Coral Deep Coral	IWV SB	$\begin{array}{c} 0.014(\pm 0.000) \\ 0.014(\pm 0.000) \\ 0.014(\pm 0.000) \\ 0.014(\pm 0.000) \end{array}$	$0.668(\pm 0.351)$ $0.282(\pm 0.056)$ $0.511(\pm 0.420)$ $0.212(\pm 0.012)$	$0.041(\pm 0.002)$ $0.041(\pm 0.000)$ $0.041(\pm 0.000)$ $0.040(\pm 0.002)$	$\begin{array}{c} 0.458(\pm 0.213) \\ 0.210(\pm 0.033) \\ 0.358(\pm 0.260) \\ 0.158(\pm 0.014) \end{array}$	$\begin{array}{c} 0.519(\pm 0.066) \\ 0.569(\pm 0.040) \\ 0.548(\pm 0.031) \\ 0.590(\pm 0.045) \end{array}$	$10.566(\pm 4.767)$ $5.416(\pm 0.356)$ $8.679(\pm 5.518)$ $4.547(\pm 0.361)$	$0.018(\pm 0.000)$ $0.018(\pm 0.000)$ $0.018(\pm 0.000)$ $0.018(\pm 0.000)$	$2.042(\pm 1.017)$ $0.874(\pm 0.103)$ $1.597(\pm 1.222)$ $0.679(\pm 0.050)$	$0.033(\pm0.000)$ $0.033(\pm0.000)$ $0.033(\pm0.000)$ $0.033(\pm0.001)$	$4.291(\pm 2.145)$ $1.841(\pm 0.199)$ $3.385(\pm 2.597)$ $1.427(\pm 0.095)$	$1.992(\pm 0.014)$ $1.977(\pm 0.017)$ $1.970(\pm 0.010)$ $1.992(\pm 0.028)$	56.628(±37.354) 20.781(±8.267) 41.196(±43.492) 13.829(±0.609)	$1.241(\pm 0.008)$ $1.231(\pm 0.014)$ $1.227(\pm 0.015)$ $1.237(\pm 0.018)$	$30.864(\pm 18.487)$ $11.823(\pm 4.643)$ $22.041(\pm 21.683)$ $8.097(\pm 0.242)$
			$0.023(\pm0.001)$	$0.469(\pm0.055)$	$0.270(\pm0.046)$	$0.685(\pm0.000)$	$2.240(\pm0.001)$	$11.474(\pm 0.290)$	$0.026(\pm0.001)$	$1.225(\pm0.165)$	$0.051(\pm0.002)$	$2.519(\pm 0.385)$	$2.860(\pm0.138)$	$27.611(\pm 5.693)$	$1.674(\pm 0.071)$	$16.226(\pm 3.967)$
	DANN DANN DANN DANN	DEV IWV SB TB	$0.019(\pm 0.001)$ $0.020(\pm 0.001)$ $0.019(\pm 0.001)$ $0.020(\pm 0.001)$	1.137(±0.054) 0.974(±0.419) 0.951(±0.347) 0.336(±0.054)	$0.061(\pm0.004)$ $0.065(\pm0.004)$ $0.063(\pm0.002)$ $0.060(\pm0.004)$	$0.733(\pm 0.032)$ $0.638(\pm 0.268)$ $0.633(\pm 0.213)$ $0.240(\pm 0.047)$	$2.236(\pm0.001)$ $2.243(\pm0.008)$ $2.239(\pm0.004)$ $2.239(\pm0.002)$	18.268(±0.381) 16.949(±3.600) 16.497(±3.351) 11.137(±0.329)	$0.027(\pm 0.001)$ $0.028(\pm 0.002)$ $0.027(\pm 0.001)$ $0.027(\pm 0.001)$	$3.419(\pm 0.150)$ $2.953(\pm 1.232)$ $2.906(\pm 1.015)$ $1.092(\pm 0.206)$	$0.052(\pm0.004)$ $0.054(\pm0.003)$ $0.052(\pm0.003)$ $0.052(\pm0.002)$	7.266(±0.319) 6.263(±2.630) 6.165(±2.173) 2.312(±0.421)	$2.909(\pm 0.165)$ $2.949(\pm 0.177)$ $2.886(\pm 0.135)$ $2.988(\pm 0.138)$	106.917(±6.089) 89.454(±42.987) 85.396(±36.953) 22.461(±4.074)	$1.688(\pm 0.091)$ $1.727(\pm 0.102)$ $1.679(\pm 0.085)$ $1.733(\pm 0.060)$	54.542(±3.799) 45.685(±21.204) 43.890(±17.807) 12.876(±2.085)
PointNet	CMD CMD CMD CMD	DEV IWV SB TB	0.020(±0.001) 0.020(±0.001) 0.019(±0.001) 0.019(±0.001)	$1.164(\pm 0.107)$ $1.232(\pm 0.036)$ $0.387(\pm 0.059)$ $0.353(\pm 0.078)$	$0.068(\pm0.009)$ $0.072(\pm0.007)$ $0.060(\pm0.004)$ $0.058(\pm0.003)$	$0.790(\pm 0.066)$ $0.836(\pm 0.039)$ $0.261(\pm 0.084)$ $0.249(\pm 0.077)$	$2.240(\pm0.002)$ $2.240(\pm0.002)$ $2.241(\pm0.002)$ $2.240(\pm0.002)$	$18.411(\pm 1.132)$ $19.099(\pm 0.332)$ $11.327(\pm 0.574)$ $11.231(\pm 0.508)$	$0.028(\pm 0.001)$ $0.028(\pm 0.001)$ $0.027(\pm 0.001)$ $0.026(\pm 0.000)$	$3.591(\pm 0.305)$ $3.775(\pm 0.107)$ $1.201(\pm 0.297)$ $1.147(\pm 0.284)$	$0.054(\pm0.003)$ $0.055(\pm0.002)$ $0.051(\pm0.001)$ $0.051(\pm0.001)$	7.643(\pm 0.591) 8.015(\pm 0.237) 2.511(\pm 0.699) 2.402(\pm 0.634)	2.996(±0.178) 3.053(±0.149) 2.852(±0.079) 2.843(±0.083)	102.004(±14.284) 110.674(±4.420) 27.922(±6.676) 23.881(±4.994)	1.749(±0.081) 1.787(±0.065) 1.675(±0.053) 1.684(±0.060)	52.048(±6.396) 55.953(±2.196) 16.105(±3.828) 13.680(±2.721)
	Deep Coral Deep Coral Deep Coral Deep Coral	SB	$0.020(\pm 0.001)$ $0.020(\pm 0.001)$ $0.019(\pm 0.000)$ $0.019(\pm 0.000)$	$1.036(\pm0.102)$ $1.048(\pm0.167)$ $0.977(\pm0.158)$ $0.346(\pm0.078)$	$0.067(\pm0.010)$ $0.067(\pm0.004)$ $0.060(\pm0.004)$ $0.058(\pm0.003)$	$0.678(\pm 0.087)$ $0.661(\pm 0.117)$ $0.651(\pm 0.079)$ $0.236(\pm 0.061)$	$2.241(\pm 0.002)$ $2.238(\pm 0.003)$ $2.243(\pm 0.002)$ $2.239(\pm 0.003)$	$17.119(\pm 1.256)$ $17.395(\pm 1.880)$ $16.764(\pm 1.497)$ $11.099(\pm 0.287)$	$0.028(\pm 0.001)$ $0.028(\pm 0.001)$ $0.027(\pm 0.001)$ $0.027(\pm 0.001)$	$3.071(\pm 0.374)$ $3.077(\pm 0.508)$ $2.947(\pm 0.409)$ $1.100(\pm 0.270)$	$0.055(\pm0.003)$ $0.055(\pm0.002)$ $0.052(\pm0.001)$ $0.051(\pm0.001)$	6.501(±0.887) 6.461(±1.120) 6.257(±0.856) 2.304(±0.618)	3.003(±0.133) 2.984(±0.108) 2.866(±0.081) 2.857(±0.089)	96.974(±10.676) 100.276(±20.956) 88.933(±21.502) 24.024(±6.005)	1.768(±0.079) 1.775(±0.050) 1.677(±0.043) 1.693(±0.037)	$50.257(\pm 3.638)$ $52.079(\pm 8.798)$ $45.919(\pm 10.860)$ $13.911(\pm 3.132)$
			$0.028(\pm0.001)$	*	*	4.614(±0.000)	$0.580(\pm0.035)$	*	$0.036(\pm0.001)$	*	$0.070(\pm0.002)$	*	$3.541(\pm0.088)$	*	$2.056(\pm0.038)$	*
	DANN DANN DANN DANN	DEV IWV SB TB	$0.024(\pm 0.001)$ $0.023(\pm 0.001)$ $0.023(\pm 0.000)$ $0.023(\pm 0.001)$	1.316(±0.047) * * 1.248(±0.044)	$0.084(\pm 0.011)$ $0.081(\pm 0.005)$ $0.078(\pm 0.004)$ $0.081(\pm 0.013)$	0.819(±0.030) * 0.772(±0.028)	$0.567(\pm0.029)$ $0.558(\pm0.035)$ $0.557(\pm0.026)$ $0.581(\pm0.052)$	19.257(±0.628) * 18.345(±0.512)	$0.037(\pm0.002)$ $0.036(\pm0.002)$ $0.035(\pm0.000)$ $0.036(\pm0.002)$	3.826(±0.120) * * 3.634(±0.123)	$0.074(\pm0.005)$ $0.071(\pm0.004)$ $0.069(\pm0.001)$ $0.072(\pm0.004)$	8.116(±0.248) * 7.702(±0.270)	$3.543(\pm 0.135)$ $3.499(\pm 0.119)$ $3.423(\pm 0.040)$ $3.483(\pm 0.132)$	135.376(±7.612) * * 126.737(±6.261)	$2.054(\pm0.071)$ $2.036(\pm0.071)$ $1.991(\pm0.021)$ $2.032(\pm0.072)$	67.099(±2.973) * 63.425(±2.241)
Transolver	CMD CMD CMD CMD	DEV IWV SB TB	0.024(±0.001) 0.024(±0.000) 0.023(±0.001) 0.024(±0.001)	1.122(±0.311) 2.699(±3.476) 0.899(±0.359) 0.567(±0.137)	$0.085(\pm0.008)$ $0.089(\pm0.005)$ $0.079(\pm0.004)$ $0.082(\pm0.004)$	$0.716(\pm 0.151)$ $1.404(\pm 1.603)$ $0.594(\pm 0.228)$ $0.405(\pm 0.100)$	$0.616(\pm0.032)$ $0.618(\pm0.032)$ $0.589(\pm0.026)$ $0.615(\pm0.046)$	16.889(±3.999) 79.583(±130.154) 13.867(±5.331) 9.350(±2.011)	$0.037(\pm0.001)$ $0.038(\pm0.001)$ $0.035(\pm0.001)$ $0.038(\pm0.002)$	3.344(±0.794) 5.006(±4.391) 2.742(±1.053) 1.798(±0.445)	$0.073(\pm0.002)$ $0.074(\pm0.001)$ $0.070(\pm0.003)$ $0.074(\pm0.003)$	7.127(±1.619) 13.120(±14.213) 5.798(±2.234) 3.834(±1.010)	$3.580(\pm 0.093)$ $3.655(\pm 0.025)$ $3.525(\pm 0.129)$ $3.601(\pm 0.152)$	107.295(±44.374) 359.583(±531.879) 77.712(±38.132) 41.985(±11.585)	2.064(±0.046) 2.098(±0.026) 2.034(±0.051) 2.085(±0.091)	54.102(±20.096) 255.183(±416.162) 41.261(±18.879) 23.186(±5.851)
	Deep Coral	DEV IWV SB TB	0.023(±0.000) 0.023(±0.001) 0.023(±0.001) 0.024(±0.000)	3.195(±4.937) * 3.614(±4.685) 0.656(±0.188)	$0.076(\pm0.003)$ $0.085(\pm0.013)$ $0.076(\pm0.005)$ $0.078(\pm0.005)$	1.514(±2.070) 4.755(±4.841) 1.904(±1.890) 0.427(±0.118)	$0.589(\pm0.010)$ $0.606(\pm0.052)$ $0.582(\pm0.017)$ $0.589(\pm0.021)$	90.875(±159.599) * 94.482(±157.213) 10.201(±2.894)	$0.035(\pm0.001)$ $0.036(\pm0.003)$ $0.034(\pm0.002)$ $0.037(\pm0.001)$	9.083(±13.802) * 10.348(±13.030) 1.985(±0.589)	$0.069(\pm0.002)$ $0.071(\pm0.006)$ $0.068(\pm0.004)$ $0.073(\pm0.003)$	22.402(±35.350) * 25.760(±33.377) 4.247(±1.285)	$3.420(\pm0.087)$ $3.511(\pm0.124)$ $3.406(\pm0.075)$ $3.523(\pm0.057)$	248.180(±377.121) * 269.209(±363.762) 53.791(±18.383)	1.992(±0.038) 2.036(±0.068) 1.988(±0.048) 2.045(±0.048)	173.301(±281.052) * 199.236(±266.513) 29.346(±8.583)

Table 4: Absolute error of PEEQ predictions for representative samples from the source and target domain of the *hot rolling* dataset. Lowest value per metric is bold.

Metric	Source	Target
Mean	2.07e-04	1.46e-02
Std	1.87e-04	2.73e-03
Median	1.66e-04	1.49e-02
Q_{01}	7.45e-09	6.82e-03
Q_{25}	6.38e-05	1.37e-02
Q_{75}	2.99e-04	1.58e-02
Q_{99}	7.61e-04	2.14e-02

B.2 SHEET METAL FORMING

In contrast to the substantial shift observed in the hot rolling dataset, the distribution shift in the *sheet metal forming* dataset is moderate. Table 5 presents the detailed performance across all models, algorithms, and selections for this dataset.

To further illustrate model behavior under distribution shift, we examine the best performing model, selected by lowest normalized average target domain error. The error distribution (Figure 10) shows a moderate distribution shift between the source and target domain with some outliers in the target domain.

To better understand the model's predictive behavior in this setting, we analyze best- and worst-case examples in each domain, again focusing on the critical PEEQ field. Table 6 provides a statistical summary of the absolute PEEQ prediction errors across the selected cases. Fringe plots in Figures 11 to 14 provide a visual understanding of model accuracy. These visualizations emphasize that while the best prediction in the target domain remains reasonably accurate, others (e.g., the worst case sample) exhibit notable discrepancies in the localized regions around the bends that we are most interested in.

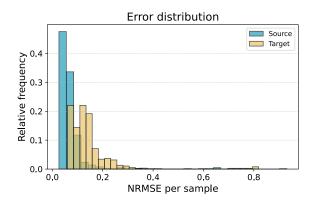


Figure 10: Distribution of NRMSE (averaged across all fields) for the test sets of the source (blue) and target (yellow) domains in the *sheet metal forming* dataset. Bar height indicates the relative frequency of samples within each bin.

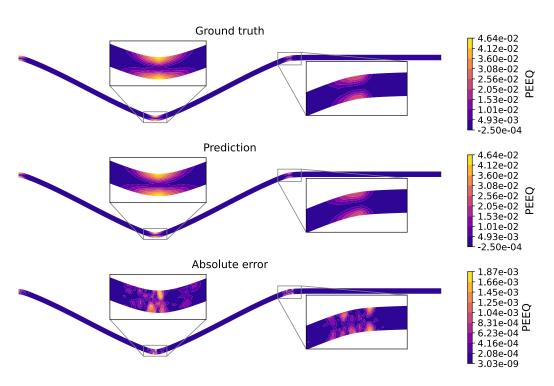


Figure 11: Fringe plot of the *sheet metal forming* dataset (best source sample). Shown is the ground truth (top) and predicted (middle) PEEQ, aswell as the absolute error (bottom).

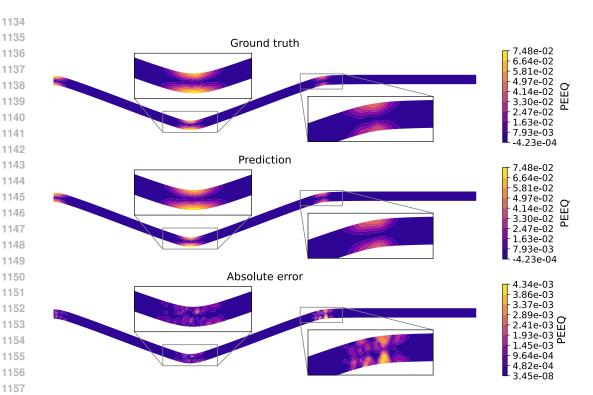


Figure 12: Fringe plot of the sheet metal forming dataset (worst source sample). Shown is the ground truth (top) and predicted (middle) PEEQ, as well as the absolute error (bottom).

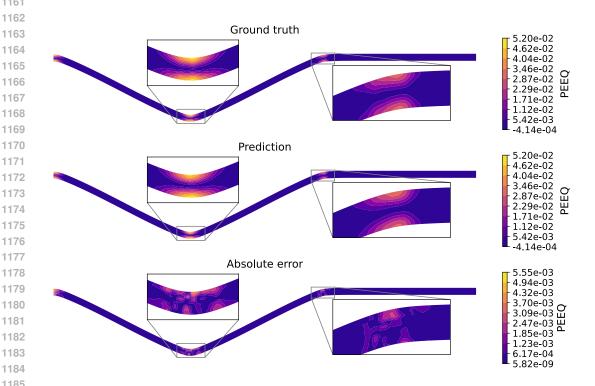


Figure 13: Fringe plot of the *sheet metal forming* dataset (best target sample). Shown is the ground truth (top) and predicted (middle) PEEQ, as well as the absolute error (bottom).

Table 5: RMSE (mean \pm std over 4 seeds) on the *sheet metal forming* dataset at *medium* difficulty. Values are target domain errors (lower is better). Bold marks the overall best model + UDA algorithm + model selection combination. For each architecture, the unregularized baseline row is shaded beige, whereas the best UDA + selection within that architecture is underlined and shaded green.

Model	DA		All Fields Nor	malized Avg (+)	Rel Custo	m Error (-)	Deforma	ition (mm)	Logarithmic S	Strain (×10 ⁻²)	Equivalent Plas	tic Strain (×10 ⁻²)	Mises S	tress (MPa)	Stre	ss (MPa)
	Algorithm	Selection	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT
		-	$0.070(\pm 0.002)$	$0.376(\pm 0.028)$	$0.238(\pm 0.013)$	$2.496(\pm0.088)$	$1.411(\pm 0.070)$	$1.939(\pm 0.530)$	$0.024(\pm0.001)$	$0.156(\pm0.014)$	$0.043(\pm0.001)$	$0.272(\pm 0.026)$	$11.022(\pm 0.324)$	46.097(±4.911)	$5.548(\pm0.198)$	$31.225(\pm 1.554)$
	DANN DANN DANN DANN	DEV IWV SB TB	$0.056(\pm 0.004)$ $0.057(\pm 0.003)$ $0.055(\pm 0.002)$ $0.057(\pm 0.003)$	* 0.329(±0.027) 1.139(±0.411) 0.323(±0.025)	$0.237(\pm 0.006)$ $0.236(\pm 0.011)$ $0.236(\pm 0.005)$ $0.239(\pm 0.013)$	$10.849(\pm 9.217)$ $5.599(\pm 4.103)$ $5.603(\pm 1.282)$ $3.580(\pm 0.384)$	$1.347(\pm 0.045)$ $1.406(\pm 0.071)$ $1.404(\pm 0.035)$ $1.416(\pm 0.055)$	$16.199(\pm 21.097)$ $2.095(\pm 0.188)$ $7.810(\pm 6.066)$ $2.021(\pm 0.156)$	$0.023(\pm 0.001)$ $0.023(\pm 0.001)$ $0.022(\pm 0.001)$ $0.023(\pm 0.001)$	$0.965(\pm 1.238)$ $0.158(\pm 0.010)$ $0.467(\pm 0.147)$ $0.156(\pm 0.010)$	$0.042(\pm 0.003)$ $0.042(\pm 0.003)$ $0.040(\pm 0.001)$ $0.042(\pm 0.003)$	* 0.269(±0.011) 0.921(±0.452) 0.265(±0.004)	$10.597(\pm 0.564)$ $10.758(\pm 0.277)$ $10.732(\pm 0.406)$ $10.728(\pm 0.218)$	406.576(±403.134) 52.401(±7.908) 186.098(±37.057) 49.234(±5.606)	$5.334(\pm 0.299)$ $5.387(\pm 0.134)$ $5.372(\pm 0.168)$ $5.405(\pm 0.167)$	177.376(±187.164) 34.644(±4.404) 94.370(±16.943) 33.375(±4.786)
GraphSAGE	CMD CMD CMD CMD	DEV IWV SB TB	$\begin{array}{c} 0.055(\pm 0.002) \\ 0.055(\pm 0.001) \\ 0.055(\pm 0.001) \\ 0.057(\pm 0.001) \end{array}$	$0.857(\pm 0.475)$ $0.407(\pm 0.124)$ $0.569(\pm 0.306)$ $0.289(\pm 0.036)$	$\begin{array}{c} 0.236(\pm 0.002) \\ 0.235(\pm 0.008) \\ 0.230(\pm 0.008) \\ 0.234(\pm 0.008) \end{array}$	$7.008(\pm 3.495)$ $4.179(\pm 1.820)$ $5.015(\pm 3.239)$ $3.005(\pm 0.236)$	$1.355(\pm 0.058)$ $1.326(\pm 0.031)$ $1.433(\pm 0.024)$ $1.345(\pm 0.059)$	$6.409(\pm 3.878)$ $2.455(\pm 1.014)$ $4.708(\pm 4.280)$ $2.028(\pm 0.798)$	$0.022(\pm 0.001)$ $0.022(\pm 0.001)$ $0.022(\pm 0.000)$ $0.023(\pm 0.000)$	$\begin{array}{c} 0.380(\pm 0.177) \\ 0.201(\pm 0.057) \\ 0.290(\pm 0.160) \\ 0.139(\pm 0.017) \end{array}$	$0.041(\pm 0.002)$ $0.041(\pm 0.001)$ $0.040(\pm 0.000)$ $0.042(\pm 0.001)$	$0.645(\pm 0.290)$ $0.358(\pm 0.112)$ $0.497(\pm 0.273)$ $0.243(\pm 0.028)$	$10.590(\pm 0.343)$ $10.730(\pm 0.065)$ $10.550(\pm 0.163)$ $10.828(\pm 0.169)$	145.233(±79.730) 61.685(±20.293) 99.069(±58.189) 43.746(±5.836)	$5.287(\pm0.140)$ $5.354(\pm0.115)$ $5.299(\pm0.071)$ $5.437(\pm0.110)$	87.123(±42.179) 37.004(±7.364) 55.134(±35.744) 29.606(±3.478)
	Deep Coral Deep Coral Deep Coral Deep Coral	DEV IWV SB TB	$\begin{array}{c} 0.054(\pm 0.002) \\ 0.055(\pm 0.002) \\ 0.056(\pm 0.002) \\ 0.056(\pm 0.003) \end{array}$	$0.422(\pm 0.092)$ $0.353(\pm 0.075)$ $0.364(\pm 0.105)$ $0.287(\pm 0.011)$	$\begin{array}{c} 0.224(\pm 0.005) \\ 0.230(\pm 0.015) \\ 0.240(\pm 0.007) \\ 0.239(\pm 0.015) \end{array}$	$\begin{array}{c} \star \\ 4.170(\pm 0.972) \\ 3.685(\pm 0.837) \\ 2.905(\pm 0.494) \end{array}$	$1.364(\pm 0.031)$ $1.389(\pm 0.055)$ $1.392(\pm 0.071)$ $1.395(\pm 0.068)$	$3.601(\pm 1.971)$ $2.449(\pm 1.115)$ $2.386(\pm 0.735)$ $1.825(\pm 0.369)$	$\begin{array}{c} 0.022(\pm0.001) \\ 0.022(\pm0.001) \\ 0.022(\pm0.001) \\ 0.023(\pm0.001) \end{array}$	$\begin{array}{c} 0.191(\pm 0.023) \\ 0.170(\pm 0.032) \\ 0.177(\pm 0.055) \\ 0.137(\pm 0.007) \end{array}$	$0.040(\pm 0.001)$ $0.041(\pm 0.002)$ $0.041(\pm 0.001)$ $0.041(\pm 0.002)$	$0.337(\pm0.056)$ $0.304(\pm0.078)$ $0.310(\pm0.090)$ $0.242(\pm0.008)$	$\begin{array}{c} 10.325(\pm0.414) \\ 10.585(\pm0.289) \\ 10.744(\pm0.189) \\ 10.781(\pm0.333) \end{array}$	65.531(±28.549) 48.326(±6.560) 52.764(±11.554) 44.161(±3.225)	$5.199(\pm 0.178)$ $5.320(\pm 0.174)$ $5.368(\pm 0.092)$ $5.398(\pm 0.179)$	43.920(±13.392) 34.667(±5.465) 35.332(±8.182) 29.228(±1.451)
			$0.077(\pm0.011)$	$0.226(\pm 0.047)$	$0.238(\pm 0.067)$	$0.881(\pm 0.135)$	$2.012(\pm 0.149)$	$2.556(\pm0.948)$	$0.024(\pm0.004)$	$0.087(\pm0.022)$	$0.045(\pm0.007)$	$0.160(\pm0.039)$	$11.357(\pm 2.106)$	$31.435(\pm 6.317)$	$8.067(\pm0.634)$	$16.525(\pm 3.262)$
	DANN DANN DANN DANN	DEV IWV SB TB	$0.067(\pm 0.006)$ $0.071(\pm 0.006)$ $0.067(\pm 0.005)$ $0.076(\pm 0.004)$	$1.184(\pm 1.941)$ $0.269(\pm 0.177)$ $0.359(\pm 0.153)$ $0.166(\pm 0.008)$	$0.329(\pm 0.013)$ $0.330(\pm 0.011)$ $0.327(\pm 0.012)$ $0.336(\pm 0.008)$	5.501(±8.113) 1.009(±0.370) 1.721(±0.875) 1.015(±0.194)	$2.250(\pm0.036)$ $2.254(\pm0.020)$ $2.250(\pm0.022)$ $2.270(\pm0.037)$	5.938(±7.076) 4.586(±5.092) 5.573(±4.577) 2.089(±0.144)	$0.025(\pm0.002)$ $0.026(\pm0.002)$ $0.025(\pm0.002)$ $0.028(\pm0.001)$	$0.705(\pm 1.197)$ $0.135(\pm 0.082)$ $0.181(\pm 0.076)$ $0.084(\pm 0.010)$	$0.046(\pm0.004)$ $0.049(\pm0.004)$ $0.047(\pm0.004)$ $0.053(\pm0.002)$	1.522(±2.652) 0.243(±0.150) 0.328(±0.138) 0.149(±0.016)	$12.112(\pm 1.331)$ $12.935(\pm 1.330)$ $12.090(\pm 1.186)$ $14.069(\pm 1.203)$	125.666(±180.946) 49.068(±41.241) 63.622(±34.926) 24.299(±2.097)	$8.601(\pm0.260)$ $8.758(\pm0.267)$ $8.522(\pm0.221)$ $9.041(\pm0.253)$	100.301(±164.456) 22.930(±17.436) 30.676(±14.620) 13.427(±0.788)
PointNet	CMD CMD CMD CMD	DEV IWV SB TB	$0.086(\pm 0.039)$ $0.072(\pm 0.003)$ $0.060(\pm 0.006)$ $0.069(\pm 0.006)$	$0.272(\pm 0.127)$ $0.244(\pm 0.146)$ $0.252(\pm 0.066)$ $0.173(\pm 0.013)$	$0.319(\pm 0.018)$ $0.314(\pm 0.011)$ $0.253(\pm 0.050)$ $0.269(\pm 0.046)$	$1.501(\pm 0.749)$ $1.587(\pm 0.689)$ $1.185(\pm 0.297)$ $1.186(\pm 0.308)$	$2.431(\pm 0.364)$ $2.253(\pm 0.039)$ $1.988(\pm 0.069)$ $2.099(\pm 0.124)$	2.887(±0.885) 2.660(±0.994) 3.698(±1.484) 2.114(±0.141)	$0.037(\pm0.025)$ $0.027(\pm0.001)$ $0.022(\pm0.002)$ $0.026(\pm0.003)$	$0.133(\pm 0.059)$ $0.117(\pm 0.070)$ $0.124(\pm 0.029)$ $0.089(\pm 0.011)$	$0.069(\pm0.046)$ $0.050(\pm0.002)$ $0.042(\pm0.005)$ $0.049(\pm0.005)$	$0.239(\pm 0.105)$ $0.213(\pm 0.125)$ $0.221(\pm 0.049)$ $0.158(\pm 0.019)$	$13.564(\pm 3.638)$ $13.159(\pm 0.877)$ $10.166(\pm 1.459)$ $12.260(\pm 0.750)$	51.155(±35.944) 47.263(±38.465) 38.406(±13.599) 25.184(±1.660)	$9.306(\pm 1.488)$ $8.843(\pm 0.213)$ $7.737(\pm 0.316)$ $8.365(\pm 0.388)$	22.677(±11.247) 20.956(±12.370) 20.153(±5.512) 13.693(±0.839)
	Deep Coral Deep Coral Deep Coral	DEV IWV SB TB	$0.067(\pm 0.008)$ $0.065(\pm 0.007)$ $0.060(\pm 0.009)$ $0.069(\pm 0.014)$	$0.228(\pm 0.065)$ $0.193(\pm 0.024)$ $0.182(\pm 0.021)$ $0.158(\pm 0.006)$	$0.290(\pm 0.075)$ $0.289(\pm 0.075)$ $0.242(\pm 0.079)$ $0.283(\pm 0.080)$	$1.626(\pm 0.438)$ $1.565(\pm 0.325)$ $1.154(\pm 0.368)$ $0.914(\pm 0.144)$	$2.201(\pm 0.189)$ $2.177(\pm 0.177)$ $2.042(\pm 0.185)$ $2.129(\pm 0.184)$	$2.613(\pm 0.839)$ $2.323(\pm 0.412)$ $2.555(\pm 0.422)$ $2.004(\pm 0.051)$	$0.025(\pm0.003)$ $0.024(\pm0.003)$ $0.022(\pm0.004)$ $0.026(\pm0.006)$	0.119(±0.040) 0.096(±0.007) 0.084(±0.011) 0.078(±0.005)	$0.046(\pm0.006)$ $0.045(\pm0.006)$ $0.042(\pm0.008)$ $0.049(\pm0.011)$	$0.213(\pm 0.067)$ $0.171(\pm 0.016)$ $0.150(\pm 0.023)$ $0.140(\pm 0.009)$	$12.087(\pm 1.995)$ $11.672(\pm 1.692)$ $10.156(\pm 2.001)$ $12.320(\pm 3.129)$	36.983(±12.354) 32.046(±6.501) 31.345(±5.362) 22.942(±1.429)	8.439(±0.665) 8.366(±0.596) 7.837(±0.674) 8.432(±0.932)	18.516(±5.099) 16.155(±2.945) 16.017(±2.153) 12.967(±0.350)
			$0.070(\pm 0.002)$	$0.168(\pm 0.029)$	$0.216(\pm 0.011)$	$0.609(\pm 0.136)$	$1.168(\pm0.012)$	$1.189(\pm 0.293)$	$0.022(\pm0.001)$	$0.070(\pm0.015)$	$0.041(\pm0.001)$	$0.126(\pm 0.029)$	$12.862(\pm0.461)$	23.014(±4.849)	$6.033(\pm0.161)$	$10.852(\pm 1.952)$
	DANN DANN DANN DANN	DEV IWV SB TB	$\begin{array}{c} 0.057(\pm 0.002) \\ 0.056(\pm 0.002) \\ 0.056(\pm 0.002) \\ 0.058(\pm 0.002) \end{array}$	$0.206(\pm 0.051)$ $0.170(\pm 0.018)$ $0.172(\pm 0.016)$ $0.133(\pm 0.016)$	$\begin{array}{c} 0.212(\pm 0.007) \\ 0.210(\pm 0.008) \\ 0.210(\pm 0.008) \\ 0.209(\pm 0.009) \end{array}$	$0.984(\pm 0.326)$ $1.011(\pm 0.239)$ $1.000(\pm 0.236)$ $0.592(\pm 0.106)$	$1.211(\pm 0.062)$ $1.171(\pm 0.057)$ $1.207(\pm 0.062)$ $1.249(\pm 0.054)$	$2.625(\pm 1.493)$ $1.595(\pm 0.395)$ $1.679(\pm 0.366)$ $1.205(\pm 0.276)$	$\begin{array}{c} 0.021(\pm 0.001) \\ 0.021(\pm 0.001) \\ 0.021(\pm 0.001) \\ 0.021(\pm 0.001) \\ 0.022(\pm 0.001) \end{array}$	$\begin{array}{c} 0.103(\pm 0.022) \\ 0.084(\pm 0.007) \\ 0.085(\pm 0.006) \\ 0.064(\pm 0.013) \end{array}$	$0.040(\pm 0.001)$ $0.040(\pm 0.002)$ $0.040(\pm 0.002)$ $0.041(\pm 0.001)$	$0.187(\pm 0.038)$ $0.156(\pm 0.014)$ $0.157(\pm 0.012)$ $0.117(\pm 0.025)$	$\begin{array}{c} 12.275(\pm 0.537) \\ 12.205(\pm 0.524) \\ 12.074(\pm 0.284) \\ 12.560(\pm 0.653) \end{array}$	36.777(±15.101) 28.101(±5.563) 28.661(±5.284) 21.245(±1.910)	$5.787(\pm0.203)$ $5.737(\pm0.227)$ $5.709(\pm0.179)$ $5.924(\pm0.299)$	$17.571(\pm 6.801)$ $13.630(\pm 2.672)$ $13.862(\pm 2.528)$ $10.337(\pm 0.834)$
Transolver	CMD CMD CMD CMD	DEV IWV SB TB	$0.058(\pm 0.002)$ $0.056(\pm 0.001)$ $0.056(\pm 0.002)$ $0.062(\pm 0.001)$	$0.275(\pm 0.130)$ $0.202(\pm 0.098)$ $0.235(\pm 0.097)$ $0.131(\pm 0.008)$	$0.208(\pm 0.006)$ $0.204(\pm 0.002)$ $0.207(\pm 0.004)$ $0.212(\pm 0.011)$	3.204(±4.010) 2.864(±4.222) 3.117(±4.077) 0.794(±0.130)	$1.226(\pm 0.058)$ $1.208(\pm 0.033)$ $1.214(\pm 0.063)$ $1.263(\pm 0.042)$	3.832(±3.244) 2.084(±1.690) 2.739(±1.545) 1.023(±0.223)	$0.022(\pm0.001)$ $0.021(\pm0.000)$ $0.021(\pm0.001)$ $0.023(\pm0.000)$	$0.137(\pm0.063)$ $0.105(\pm0.055)$ $0.122(\pm0.053)$ $0.065(\pm0.005)$	$0.041(\pm 0.001)$ $0.040(\pm 0.001)$ $0.040(\pm 0.001)$ $0.044(\pm 0.001)$	$0.248(\pm 0.111)$ $0.187(\pm 0.093)$ $0.215(\pm 0.090)$ $0.117(\pm 0.007)$	$12.806(\pm 0.937)$ $12.374(\pm 0.580)$ $12.145(\pm 0.515)$ $13.505(\pm 0.428)$	49.516(±31.330) 30.122(±14.430) 35.915(±14.900) 20.285(±1.747)	$5.985(\pm 0.366)$ $5.815(\pm 0.209)$ $5.731(\pm 0.224)$ $6.326(\pm 0.169)$	24.985(±14.508) 16.390(±8.969) 19.679(±9.245) 9.821(±0.838)
	Deep Coral Deep Coral Deep Coral Deep Coral	IWV SB TB	$0.058(\pm 0.001)$ $0.057(\pm 0.001)$ $0.057(\pm 0.001)$ $0.059(\pm 0.001)$	$\begin{array}{c} {\bf 0.155(\pm 0.015)} \\ \hline 0.255(\pm 0.206) \\ 0.263(\pm 0.201) \\ 0.138(\pm 0.014) \end{array}$	$\begin{array}{c} 0.209(\pm0.008) \\ 0.210(\pm0.008) \\ 0.209(\pm0.009) \\ 0.205(\pm0.008) \end{array}$	$0.803(\pm0.187)$ $1.370(\pm0.967)$ $1.692(\pm1.028)$ $0.802(\pm0.260)$	$1.220(\pm 0.043)$ $1.219(\pm 0.022)$ $1.199(\pm 0.019)$ $1.227(\pm 0.016)$	1.356(±0.309) 2.945(±3.136) 3.277(±2.944) 0.957(±0.036)	0.022(±0.001) 0.021(±0.000) 0.021(±0.001) 0.022(±0.000)	0.079(±0.007) 0.130(±0.109) 0.135(±0.106) 0.068(±0.012)	0.042(±0.001) 0.041(±0.000) 0.040(±0.001) 0.042(±0.001)	$0.143(\pm 0.011)$ $0.235(\pm 0.195)$ $0.244(\pm 0.189)$ $0.124(\pm 0.023)$	$\begin{array}{c} 12.849(\pm 0.230) \\ 12.630(\pm 0.275) \\ 12.509(\pm 0.180) \\ 12.970(\pm 0.502) \end{array}$	24.395(±2.894) 43.648(±38.095) 44.318(±37.691) 22.062(±2.213)	$6.023(\pm0.091)$ $5.953(\pm0.136)$ $5.878(\pm0.082)$ $6.080(\pm0.207)$	12.115(±1.501) 22.342(±20.099) 22.645(±19.921) 10.846(±0.704)

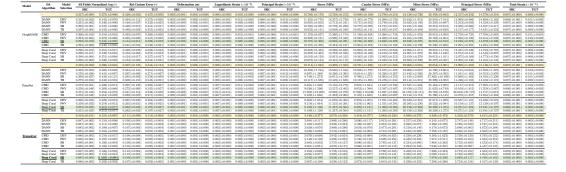
Table 6: Absolute error of PEEQ predictions for the best and worst samples from the source and target domain of the *sheet metal forming* dataset. Lowest value per metric is bold.

Metric	Sou	ırce	Target			
11200220	Best	Worst	Best	Worst		
Mean	5.47e-05	1.68e-04	1.72e-04	1.86e-03		
Std	1.35e-04	3.80e-04	4.52e-04	5.97e-03		
Median	1.96e-05	4.35e-05	5.58e-05	2.85e-04		
Q_{01}	2.43e-07	6.52e-07	1.03e-06	6.66e-06		
Q_{25}	7.49e-06	1.84e-05	2.60e-05	1.53e-04		
Q_{75}	4.05e-05	1.05e-04	9.10e-05	4.67e-04		
Q99	7.37e-04	1.96e-03	2.56e-03	3.40e-02		

B.3 ELECTRIC MOTOR DESIGN

Table 7 presents the complete benchmarking results for the *electric motor design* dataset. For this dataset the relative degradation in model performance in the target domain is in general smaller than in the previous two presented above.

Table 7: RMSE (mean \pm std over 4 seeds) on the *electric motor design* dataset at *medium* difficulty. Values are target domain errors (lower is better). Bold marks the overall best model + UDA algorithm + model selection combination. For each architecture, the unregularized baseline row is shaded beige, whereas the best UDA + selection within that architecture is underlined and shaded green.



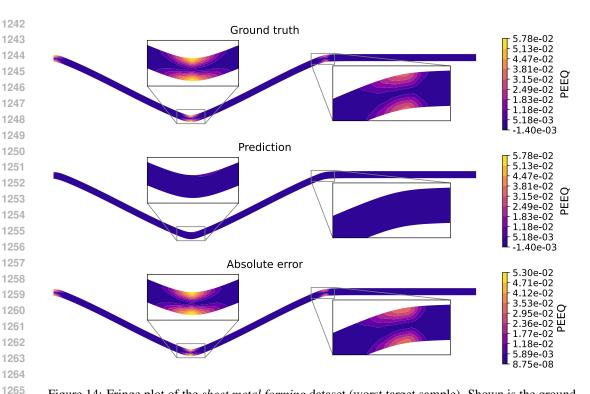


Figure 14: Fringe plot of the sheet metal forming dataset (worst target sample). Shown is the ground truth (top) and predicted (middle) PEEQ, as well as the absolute error (bottom).

To assess the effect of the domain shift on prediction accuracy in the *electric motor design* dataset further, Figure 15 shows the distribution of NRMSEs for the best performing model, selected by lowest average error in the target domain, in the source and target domain.

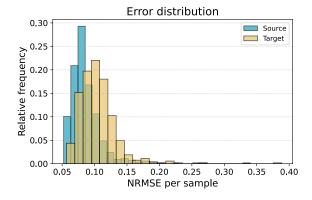


Figure 15: Distribution of NRMSE (averaged across all fields) for the test sets of the source (blue) and target (yellow) domains in the *electric motor design* dataset. Bar height indicates the relative frequency of samples within each bin.

In this task, the Mises stress is used as a scalar summary of the multi-axial stress state and is particularly interesting for downstream analysis and optimization. We therefore focus our closer inspection on this field.

Table 8 presents a comparison of absolute Mises stress errors for the best and worst samples from both the source and target test sets. The corresponding fringe plots are shown in Figures 16 to 19, comparing the ground truth and predicted fields alongside their absolute errors. They show that the best samples are predicted very well, whereas the worst sample of the source domain visually appears

slightly worse than the one of the target domain. On average, however, it is still predicted more accurately than the worst sample of the target domain, as shown in Table 8.

Table 8: Absolute error (MPa) of Mises stress predictions for the best and worst samples from the source and target domain of the *electric motor design* dataset. Lowest value per metric is bold.

Metric	So	urce	Target				
1,100110	Best	Worst	Best	Worst			
Mean	2.00	20.50	2.67	23.21			
Std	2.73	39.50	3.13	23.32			
Median	1.26	13.63	1.68	13.09			
Q_{01}	0.02	0.18	0.03	0.11			
Q_{25}	0.60	5.08	0.75	2.87			
Q ₇₅	2.24	24.05	3.50	41.97			
Q_{99}	13.43	140.17	15.05	78.28			

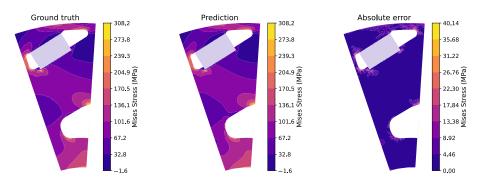


Figure 16: Fringe plot of the *electric motor design* dataset (best source sample). Shown is the ground truth (left) and predicted (middle) Mises stress, as well as the absolute error (right).

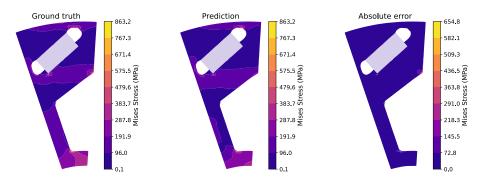


Figure 17: Fringe plot of the *electric motor design* dataset (worst source sample). Shown is the ground truth (left) and predicted (middle) Mises stress, as well as the absolute error (right).

B.4 HEATSINK DESIGN

Table 9 presents the complete benchmarking results for the *heatsink design* dataset.

We again further investigate model performance under distribution shift by examining predictions from the best performing model, selected by lowest average error in the target domain. Figure 20 presents the respective distribution of prediction errors in the source and target domain, clearly indicating the negative effects of the distribution shift on model performance.

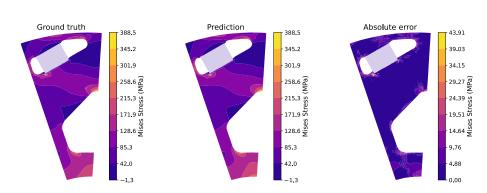


Figure 18: Fringe plot of the *electric motor design* dataset (best target sample). Shown is the ground truth (left) and predicted (middle) Mises stress, as well as the absolute error (right).

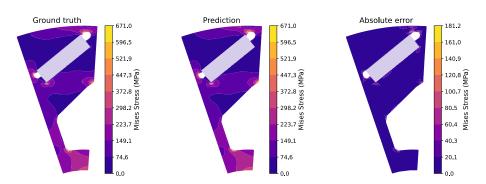


Figure 19: Fringe plot of the *electric motor design* dataset (worst target sample). Shown is the ground truth (left) and predicted (middle) Mises stress, as well as the absolute error (right).

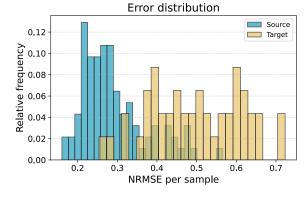


Figure 20: Distribution of NRMSE (averaged across all fields) for the test sets of the source (blue) and target (yellow) domains in the *heatsink design* dataset. Bar height indicates the relative frequency of samples within each bin.

Table 9: RMSE (mean \pm std over 4 seeds) on the *heatsink* dataset at *medium* difficulty. Values are target domain errors (lower is better). Bold marks the overall best model + UDA algorithm + model selection combination. For each architecture, the unregularized baseline row is shaded beige, whereas the best UDA + selection within that architecture is underlined and shaded green.

Model	DA	Model	All Fields Nor	malized Avg (-)	Tempera	ature (K)	Velocit	ty (m/s)	Rel Custo	m Error (-)	Pressur	re (kPa)
Model	Algorithm	Selection	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT
	-	-	$0.525(\pm0.026)$	$0.568(\pm0.030)$	15.581(±1.535)	21.126(±2.365)	$0.054(\pm0.002)$	$0.044(\pm0.000)$	$0.036(\pm0.006)$	$0.050(\pm0.009)$	$385.602(\pm 34.377)$	1879.390(±239.203)
	DANN DANN DANN DANN	DEV IWV SB TB	$\begin{array}{c} 0.315(\pm 0.120) \\ 0.317(\pm 0.118) \\ 0.228(\pm 0.016) \\ 0.304(\pm 0.036) \end{array}$	$0.458(\pm 0.039)$ $0.467(\pm 0.043)$ $0.494(\pm 0.026)$ $0.397(\pm 0.019)$	10.562(±5.446) 10.707(±5.293) 6.668(±1.013) 10.964(±1.411)	$\begin{array}{c} 20.006(\pm 2.840) \\ 19.269(\pm 3.565) \\ 20.129(\pm 2.380) \\ 15.719(\pm 1.387) \end{array}$	$\begin{array}{c} 0.040(\pm 0.011) \\ 0.040(\pm 0.011) \\ 0.031(\pm 0.002) \\ 0.041(\pm 0.005) \end{array}$	$\begin{array}{c} 0.049(\pm0.007) \\ 0.049(\pm0.007) \\ 0.055(\pm0.002) \\ 0.043(\pm0.002) \end{array}$	$\begin{array}{c} 0.022(\pm0.016) \\ 0.022(\pm0.016) \\ 0.010(\pm0.002) \\ 0.021(\pm0.004) \end{array}$	$\begin{array}{c} 0.042(\pm 0.010) \\ 0.040(\pm 0.013) \\ 0.041(\pm 0.008) \\ 0.029(\pm 0.006) \end{array}$	786.388(±1051.364) 782.535(±1054.138) 207.063(±13.650) 331.207(±140.774)	2037.844(±331.813) 2358.560(±646.551) 2102.712(±614.979) 1907.682(±232.159)
PointNet	CMD CMD CMD CMD	DEV IWV SB TB	$\begin{array}{c} 0.424(\pm 0.003) \\ 0.239(\pm 0.008) \\ 0.238(\pm 0.007) \\ 0.302(\pm 0.086) \end{array}$	$0.443(\pm 0.003)$ $0.480(\pm 0.020)$ $0.477(\pm 0.023)$ $0.442(\pm 0.018)$	$16.292(\pm 0.119)$ $7.577(\pm 0.479)$ $7.483(\pm 0.406)$ $10.801(\pm 4.087)$	$\begin{array}{c} 20.676(\pm0.141) \\ 18.524(\pm1.213) \\ 18.666(\pm1.057) \\ 17.800(\pm2.256) \end{array}$	$\begin{array}{c} 0.042(\pm 0.001) \\ 0.033(\pm 0.001) \\ 0.033(\pm 0.001) \\ 0.037(\pm 0.004) \end{array}$	$\begin{array}{c} 0.042(\pm 0.001) \\ 0.051(\pm 0.002) \\ 0.051(\pm 0.002) \\ 0.046(\pm 0.004) \end{array}$	$\begin{array}{c} 0.042(\pm 0.000) \\ 0.012(\pm 0.001) \\ 0.012(\pm 0.001) \\ 0.022(\pm 0.014) \end{array}$	$\begin{array}{c} 0.048(\pm 0.001) \\ 0.035(\pm 0.003) \\ 0.035(\pm 0.003) \\ 0.036(\pm 0.009) \end{array}$	$\begin{array}{c} 2388.245(\pm 17.831) \\ 192.803(\pm 4.658) \\ 196.144(\pm 5.022) \\ 756.939(\pm 1076.860) \end{array}$	$2464.906(\pm 43.671)$ $2455.172(\pm 118.483)$ $2375.674(\pm 157.350)$ $2288.916(\pm 108.201)$
	Deep Coral Deep Coral Deep Coral Deep Coral	DEV IWV SB TB	$0.276(\pm 0.072)$ $0.276(\pm 0.072)$ $0.270(\pm 0.061)$ $0.343(\pm 0.063)$	$0.396(\pm 0.049)$ $0.396(\pm 0.049)$ $0.394(\pm 0.048)$ $0.384(\pm 0.042)$	9.382(±3.681) 9.382(±3.681) 9.071(±3.069) 12.763(±3.067)	18.009(±2.336) 18.009(±2.336) 17.428(±1.939) 18.517(±2.502)	$0.038(\pm0.010)$ $0.038(\pm0.010)$ $0.037(\pm0.009)$ $0.047(\pm0.009)$	$0.044(\pm0.006)$ $0.044(\pm0.006)$ $0.044(\pm0.006)$ $0.042(\pm0.004)$	$\begin{array}{c} 0.017(\pm 0.011) \\ 0.017(\pm 0.011) \\ 0.016(\pm 0.008) \\ 0.027(\pm 0.009) \end{array}$	$\begin{array}{c} 0.033(\pm 0.012) \\ 0.033(\pm 0.012) \\ 0.031(\pm 0.010) \\ 0.040(\pm 0.012) \end{array}$	240.309(±86.936) 240.309(±86.936) 224.078(±54.512) 324.404(±102.861)	$1035.599(\pm 570.783)$ $1035.599(\pm 570.783)$ $1037.166(\pm 573.817)$ $1438.782(\pm 427.051)$
	-	-	$0.245(\pm0.004)$	$0.446(\pm 0.018)$	$4.331(\pm0.057)$	$9.720(\pm0.273)$	$0.024(\pm0.000)$	$0.038(\pm0.001)$	$0.007(\pm0.000)$	$0.010(\pm0.001)$	$307.475(\pm 33.199)$	$1687.115(\pm 262.894)$
	DANN DANN DANN DANN	DEV IWV SB TB	$0.183(\pm 0.001)$ $0.214(\pm 0.056)$ $0.183(\pm 0.002)$ $0.242(\pm 0.067)$	$0.436(\pm 0.011)$ $0.425(\pm 0.023)$ $0.422(\pm 0.015)$ $0.406(\pm 0.024)$	$4.322(\pm 0.082)$ $5.911(\pm 2.935)$ $4.342(\pm 0.076)$ $7.479(\pm 3.602)$	$\begin{array}{c} 15.263(\pm 0.933) \\ 14.712(\pm 1.180) \\ 13.943(\pm 1.120) \\ 14.101(\pm 0.337) \end{array}$	$\begin{array}{c} 0.024(\pm 0.000) \\ 0.028(\pm 0.007) \\ 0.024(\pm 0.000) \\ 0.032(\pm 0.009) \end{array}$	$0.049(\pm 0.001)$ $0.047(\pm 0.004)$ $0.047(\pm 0.002)$ $0.045(\pm 0.003)$	$0.007(\pm 0.000)$ $0.011(\pm 0.006)$ $0.007(\pm 0.000)$ $0.014(\pm 0.008)$	$\begin{array}{c} 0.023(\pm 0.003) \\ 0.023(\pm 0.002) \\ 0.021(\pm 0.002) \\ 0.023(\pm 0.003) \end{array}$	$302.724(\pm 26.912)$ $374.431(\pm 150.069)$ $295.253(\pm 24.111)$ $436.461(\pm 137.106)$	$2031.634(\pm 48.462)$ $2166.222(\pm 193.875)$ $2101.032(\pm 65.772)$ $2066.973(\pm 155.445)$
Transolver	CMD CMD CMD CMD	DEV IWV SB TB	$0.312(\pm 0.088)$ $0.181(\pm 0.001)$ $0.180(\pm 0.002)$ $0.180(\pm 0.002)$	$0.437(\pm 0.040)$ $0.375(\pm 0.032)$ $0.413(\pm 0.044)$ $0.367(\pm 0.026)$	$11.200(\pm 4.591)$ $4.272(\pm 0.141)$ $4.265(\pm 0.099)$ $4.238(\pm 0.137)$	$17.013(\pm 3.502)$ $11.934(\pm 2.484)$ $13.661(\pm 2.706)$ $11.225(\pm 2.088)$	$0.029(\pm 0.004)$ $0.024(\pm 0.000)$ $0.024(\pm 0.000)$ $0.024(\pm 0.000)$	$0.043(\pm 0.002)$ $0.042(\pm 0.004)$ $0.046(\pm 0.005)$ $0.041(\pm 0.003)$	$0.030(\pm 0.015)$ $0.007(\pm 0.000)$ $0.007(\pm 0.000)$ $0.007(\pm 0.000)$	$0.038(\pm 0.014)$ $0.015(\pm 0.006)$ $0.020(\pm 0.007)$ $0.014(\pm 0.005)$	$1886.652(\pm 1080.114)$ $274.411(\pm 9.762)$ $291.801(\pm 20.389)$ $260.403(\pm 22.204)$	2701.731(±449.442) 1816.960(±32.439) 2008.434(±140.227) 1893.827(±183.300)
	Deep Coral Deep Coral Deep Coral Deep Coral	DEV IWV SB TB	$0.181(\pm 0.001)$ $0.181(\pm 0.001)$ $0.181(\pm 0.002)$ $0.181(\pm 0.001)$	$\begin{array}{c} \textbf{0.310}(\pm 0.019) \\ \hline 0.326(\pm 0.029) \\ 0.334(\pm 0.024) \\ 0.310(\pm 0.019) \end{array}$	$4.257(\pm 0.074)$ $4.269(\pm 0.077)$ $4.276(\pm 0.039)$ $4.257(\pm 0.074)$	8.717(±0.824) 9.493(±1.143) 9.624(±1.234) 8.717(±0.824)	$0.024(\pm0.000)$ $0.024(\pm0.000)$ $0.024(\pm0.000)$ $0.024(\pm0.000)$	$0.037(\pm0.002)$ $0.037(\pm0.002)$ $0.038(\pm0.002)$ $0.037(\pm0.002)$	$0.007(\pm0.000)$ $0.007(\pm0.000)$ $0.007(\pm0.000)$ $0.007(\pm0.000)$	$0.009(\pm0.002)$ $0.010(\pm0.001)$ $0.010(\pm0.002)$ $0.009(\pm0.002)$	290.470(±21.975) 293.812(±6.818) 290.325(±21.486) 290.470(±21.975)	1334.241(±173.205) 1545.184(±365.222) 1620.377(±262.907) 1334.241(±173.205)
	-	-	$0.244(\pm 0.002)$	$0.441(\pm 0.024)$	$4.316(\pm0.028)$	$13.033(\pm 1.059)$	$0.025(\pm0.000)$	$0.040(\pm0.002)$	$0.006(\pm0.000)$	$0.013(\pm 0.004)$	231.715(±14.315)	816.180(±48.994)
	DANN DANN DANN DANN	DEV IWV SB TB	$0.188(\pm 0.011)$ $0.222(\pm 0.053)$ $0.184(\pm 0.002)$ $0.273(\pm 0.092)$	$0.446(\pm 0.026)$ $0.443(\pm 0.070)$ $0.480(\pm 0.018)$ $0.398(\pm 0.038)$	4.651(±0.781) 6.731(±3.132) 4.285(±0.072) 9.411(±4.841)	15.580(±0.609) 15.179(±1.591) 15.689(±0.806) 15.644(±3.334)	$0.026(\pm0.002)$ $0.030(\pm0.007)$ $0.025(\pm0.000)$ $0.037(\pm0.012)$	$0.050(\pm0.003)$ $0.048(\pm0.006)$ $0.051(\pm0.001)$ $0.043(\pm0.004)$	$0.007(\pm0.002)$ $0.011(\pm0.007)$ $0.006(\pm0.000)$ $0.019(\pm0.013)$	$0.024(\pm0.003)$ $0.024(\pm0.006)$ $0.023(\pm0.003)$ $0.027(\pm0.011)$	222.968(±13.097) 247.259(±32.883) 243.565(±23.766) 285.494(±72.700)	2164.527(±302.094) 2380.114(±727.167) 2729.454(±516.782) 1872.493(±365.646)
UPT	CMD CMD CMD CMD	DEV IWV SB TB	$0.210(\pm 0.055)$ $0.182(\pm 0.000)$ $0.179(\pm 0.001)$ $0.182(\pm 0.000)$	$0.406(\pm 0.046)$ $0.363(\pm 0.015)$ $0.444(\pm 0.010)$ $0.363(\pm 0.015)$	$5.994(\pm 3.353)$ $4.297(\pm 0.038)$ $4.135(\pm 0.026)$ $4.297(\pm 0.038)$	$14.289(\pm 2.054)$ $12.908(\pm 0.487)$ $16.130(\pm 0.627)$ $12.908(\pm 0.487)$	$0.028(\pm0.007)$ $0.025(\pm0.000)$ $0.024(\pm0.000)$ $0.025(\pm0.000)$	$0.046(\pm0.005)$ $0.043(\pm0.001)$ $0.050(\pm0.001)$ $0.043(\pm0.001)$	$0.010(\pm0.008)$ $0.006(\pm0.000)$ $0.006(\pm0.000)$ $0.006(\pm0.000)$	$0.021(\pm 0.006)$ $0.017(\pm 0.001)$ $0.025(\pm 0.002)$ $0.017(\pm 0.001)$	$236.012(\pm 21.869)$ $221.370(\pm 9.452)$ $230.955(\pm 7.947)$ $221.370(\pm 9.452)$	1873.691(±394.368) 1365.372(±256.867) 1918.684(±52.452) 1365.372(±256.867)
	Deep Coral Deep Coral Deep Coral	DEV IWV SB TB	$0.183(\pm 0.001)$ $0.183(\pm 0.001)$ $0.182(\pm 0.000)$ $0.182(\pm 0.000)$	$0.345(\pm 0.013)$ $0.339(\pm 0.020)$ $0.328(\pm 0.012)$ $0.321(\pm 0.008)$	$4.318(\pm 0.067)$ $4.329(\pm 0.056)$ $4.305(\pm 0.045)$ $4.347(\pm 0.039)$	$13.290(\pm 0.655)$ $13.190(\pm 0.784)$ $12.610(\pm 1.328)$ $12.637(\pm 0.949)$	$0.025(\pm0.000)$ $0.025(\pm0.000)$ $0.025(\pm0.000)$ $0.025(\pm0.000)$	$0.041(\pm0.001)$ $0.041(\pm0.002)$ $0.040(\pm0.002)$ $0.039(\pm0.001)$	$0.006(\pm0.000)$ $0.006(\pm0.000)$ $0.006(\pm0.000)$ $0.006(\pm0.000)$	$0.010(\pm0.002)$ $0.010(\pm0.002)$ $0.014(\pm0.008)$ $0.013(\pm0.007)$	221.345(±8.107) 222.325(±7.415) 218.254(±10.676) 218.167(±11.586)	809.556(±99.407) 769.346(±63.531) 865.504(±170.301) 791.890(±121.633)

In this task, the temperature field is the most critical for downstream analysis and optimization, which is why we focus our detailed analysis on it.

Table 10 compares the absolute temperature prediction errors for the best and worst samples from both the source and target test sets. The corresponding scatter plots are shown in Figures 21 to 24, comparing the ground truth and predicted temperature fields, alongside their absolute errors.

While the best source domain prediction is quite accurate, with low average and percentile errors (Table 10, Figure 21), the 99th percentile of the worst source domain prediction reaches up to 29K. Given a total temperature range of 100K, this represents a relative error of nearly 30%. The worst target domain prediction is even less accurate, showing substantial visual and quantitative deviations from the ground truth (Table 10, Figure 24).

Table 10: Absolute error (K) of temperature predictions for the best and worst samples in the source and target domain of the *heatsink design* dataset. Lowest value per metric is bold.

Metric	Sou	ırce	Target			
1,101110	Best	Worst	Best	Worst		
Mean	1.84e+00	5.79e+00	2.23e+00	1.42e+01		
Std	1.94e+00	5.90e+00	2.85e+00	1.46e+01		
Median	1.25e+00	4.06e+00	1.31e+00	8.84e+00		
Q_{01}	2.17e-02	7.51e-02	2.41e-02	1.62e-01		
Q_{25}	5.49e-01	1.92e+00	5.95e-01	4.49e+00		
Q_{75}	2.44e+00	7.48e+00	2.68e+00	1.87e+01		
Q ₉₉	9.26e+00	2.88e+01	1.49e+01	6.61e+01		

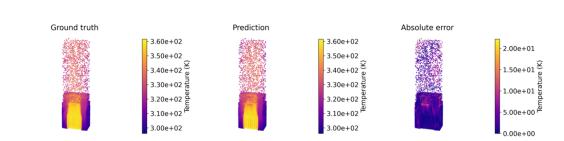


Figure 21: Sliced scatter plot of the *heatsink design* dataset (best source sample). Shown is the ground truth (left) and predicted (middle) temperature field, as well as the absolute error (right).

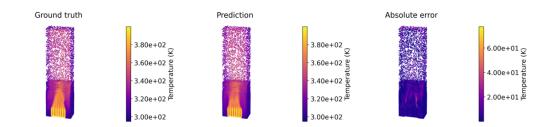


Figure 22: Sliced scatter plot of the *heatsink design* dataset (worst source sample). Shown is the ground truth (left) and predicted (middle) temperature field, as well as the absolute error (right).

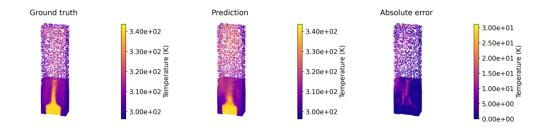


Figure 23: Sliced scatter plot of the *heatsink design* dataset (best target sample). Shown is the ground truth (left) and predicted (middle) temperature field, as well as the absolute error (right).

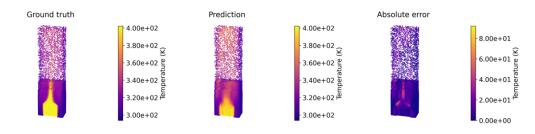


Figure 24: Sliced scatter plot of the *heatsink design* dataset (worst target sample). Shown is the ground truth (left) and predicted (middle) temperature field, as well as the absolute error (right).

C DISTRIBUTION SHIFTS

Table 11 provides an overview of the parameter ranges chosen to define source and target domains for different task difficulties across all datasets. To gain more insights into the parameter importance besides the domain experts' opinion, we visualize the latent space of the conditioning network for all presented datasets in Figures 25 to 28.

Table 11: Defined distribution shifts (source and target domains) of each dataset and each difficulty.

Dataset	Parameter	Difficulty	Source range (no. samples)	Target range (no. samples)
Rolling	Reduction $r(-)$	easy medium hard	[0.01, 0.13) (4000) [0.01, 0.115) (3500) [0.01, 0.10) (3000)	[0.13, 0.15] (750) [0.115, 0.15] (1250) [0.10, 0.15] (1750)
Forming	Thickness t (mm)	easy medium hard	[2, 4.8) (3060) [2, 4.3) (2550) [2, 4.1) (2295)	[4.8, 5] (255) [4.3, 5] (765) [4.1, 5] (1020)
Electric Motor	Rotor slot diameter 3 $d_{r3}\ (mm)$	easy medium hard	[100, 122)(2693) [99, 120) (2143) [99, 118) (1728)	[122, 126](504) [120, 126] (1054) [118, 126] (1469)
Heatsink	# fins	easy medium hard	[5, 13) (404) [5, 12) (365) [5, 11) (342)	[13, 14] (56) [12, 15] (95) [11, 15] (118)

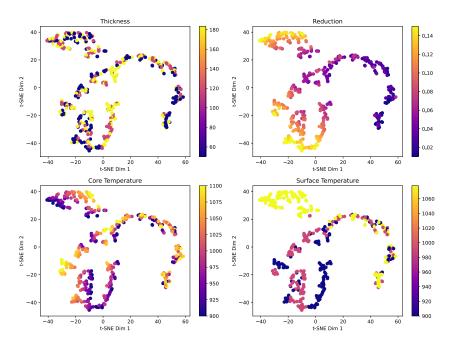


Figure 25: T-SNE visualization of the conditioning vectors for the *hot rolling* dataset. Point color indicates the magnitude of the respective parameter. While the slab thickness t appears to be uniformly distributed, the remaining three exhibit distinct clustering patterns. Taking into account domain knowledge from industry experts, we defined the reduction parameter r as the basis for constructing distribution shifts.

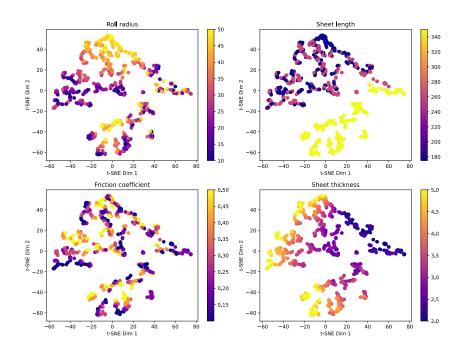


Figure 26: T-SNE visualization of the conditioning vectors for the *sheet metal forming* dataset. Point color indicates the magnitude of the respective parameter. The sheet length l shows the most distinct groupings, but with only three discrete values, it is unsuitable for defining domain splits. The friction coefficient μ appears uniformly distributed across the embedding. In contrast, sheet thickness t and roll radius t show clustering behavior, making them more appropriate candidates for inducing distribution shifts. We choose t as the domain defining parameter.

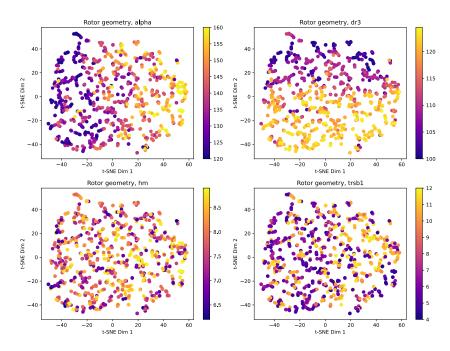


Figure 27: T-SNE visualization of the conditioning vectors for the *electric motor design* dataset. Point color indicates the magnitude of the respective parameter. For clarity, we only show selected parameters. The only parameter for which exhibits see some structure in the latent space is d_{r3} , we therefore choose this to be our domain defining parameter in accordance with domain experts.

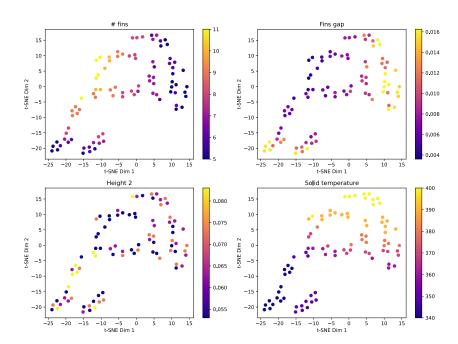


Figure 28: T-SNE visualization of the conditioning vectors for the *heatsink design* dataset. Point color indicates the magnitude of the respective parameter. Height 2 is distributed equally across the representation, but the other parameters show concrete grouping behavior. We therefore choose the number of fins as the domain defining parameter.

D MODEL ARCHITECTURES

This section provides explanations of all model architectures used in our benchmark. All models are implemented in PyTorch and are adapted to our conditional regression task. All models have in common, that they take node coordinates as inputs and embed them using a sinusoidal positional encoding. Additionally, all models are conditioned on the input parameters of the respective simulation sample, which are encoded through a conditioning network described below.

Conditioning Network. The conditioning module used for all neural surrogate architectures embeds the simulation input parameters into a latent vector used for conditioning. The network consists of a sinusoidal encoding followed by a simple MLP. The dimension of the latent encoding is 8 throughout all experiments.

PointNet. Our PointNet implementation is adapted from (Qi et al., 2017) for node-level regression. Input node coordinates are first encoded using sinusoidal embeddings and passed through an encoder MLP. The resulting representations are aggregated globally using max pooling over nodes to obtain a global feature vector. To propagate this global feature, it is concatenated back to each point's feature vector. This fused representation is then fed into a final MLP, which produces the output fields. The conditioning is performed by concatenating the conditioning vector to the global feature before propagating it to the nodes features. We use a PointNet base dimension of 16 for the small model and 32 for the larger model.

GraphSAGE. We adapt GraphSAGE (Hamilton et al., 2017) to the conditional mesh regression setting. Again, input node coordinates are embedded using a sinusoidal encoding and passed through an MLP encoder. The main body of the model consists of multiple GraphSAGE message passing layers with mean aggregation. We support two conditioning modes, namely concatenating the latent conditioning vector to the node features, or applying FiLM style modulation (Perez et al., 2018) to the node features before each message passing layer. We always use FiLM modulation in the presented results. After message passing, the node representations are passed through a final MLP

decoder to produce the output fields. The base dimension of the model is kept at 128 and we employ 4 GraphSAGE layers.

Transolver. The Transolver model follows the originally introduced architecture (Wu et al., 2024). Similar to the other models, node coordinates first are embedded using a sinusoidal encoding and passed through an MLP encoder to produce initial features. Through learned assignement, each node then gets mapped to a slice, and inter- as well as intra-slice attention is performed. Afterwards, fields are decoded using an MLP readout. The architecture supports two conditioning modes: concatenation, where the conditioning vector is concatenated to the input node features before projection, or modulation through DiT layers across the network. For our experiments, DiT is used. We choose a latent dimension of 128, a slice base of 32 and we apply four attention blocks for the small model. For the larger model, we scale to 256, 128 and 8 layers respectively.

UPT. Our UPT implementation builds on the architecture proposed in (Alkin et al., 2024a). First, a fixed number of supernodes are uniformly sampled from the input nodes. Node coordinates are embedded using a sinusoidal encoding followed by an MLP. The supernodes aggregate features from nearby nodes using one-directional message passing and serve as tokens for subsequent transformer processing. They are then processed by stack of DiT blocks, which condition the network on the simulation input parameters. For prediction, we employ a DiT Perceiver (Jaegle et al., 2022) decoder that performs cross-attention between the latent representation and a set of query positions. This allows the model to generate field predictions at arbitrary spatial locations, which is a desirable property for inference. We sample 4096 supernodes and use a base dimension of 192. We use 8 DiT blocks for processing and 4 DiT Perceiver blocks for decoding.

E EXPERIMENTS

This section provides a detailed overview of the performed experiments for this benchmark. First, we explain the benchmarking setup used to generate the benchmarking results in detail in Appendix E.1 and the evaluation procedure in Appendix E.2. Furthermore, we provide information about training times for the presented methods in Appendix E.3.

E.1 EXPERIMENTAL SETUP

Dataset Splits. We split each dataset into source and target domains as outlined in Section 3.5 and Appendix C. Within source domains, we use a 50%/25%/25% split for training, validation, and testing, respectively. For target domains, where labels are unavailable during training in our UDA setup, we use a 50%/50% split for training and test sets. The large validation and test sets are motivated the industrial relevance of our benchmark, where reliable performance estimation on unseen data is a crucial factor.

Training Pipeline. For training, we use a dataset wide per field z-score normalization strategy, with statistics computed on the source domain training set. We use a batch size of 16 and the AdamW optimizer (Loshchilov & Hutter, 2019) with a weight decay of 1e-5 and a cosine learning rate schedule, starting from 1e-3. Gradients are clipped to a maximum norm of 1. For the large scale *heatsink design* dataset, we enable Automatic Mixed Precision (AMP) to reduce memory consumption and training time. Additionally, we use Exponential Moving Average (EMA) updates with a decay factor of 0.95 to stabilize training.

Performance metrics are evaluated every 10 epochs, and we train all models for a maximum of 3000 epochs with early stopping after 500 epochs of no improvement on the source domain validation loss.

Domain Adaptation Specifics. To enable UDA algorithms, we jointly sample mini batches from the source and target domains at each training step and pass them thorugh the model. Since target labels are not available, we compute supervised losses only on the source domain outputs. In addition, we compute DA losses on the latent representations of source and target domains in order to encourage domain invariance.

Since a crucial factor in the performance of UDA algorithms is the choice of the domain adaptation loss weight λ , we perform extensive sweeps over this hyperparameter and select models using the unsupervised model selection strategies described in Section 4.3.

For the three smaller datasets, we sweep λ logarithmically over $\lambda \in \{10^{-1}, 10^{-2}, \dots, 10^{-9}\}$, while for the large scale *Heatsink design* dataset, we sweep a smaller range, namely $\lambda \in \{10^2, 10^{-1}, \dots, 10^{-2}\}$, motivated by the balancing principle (Zellinger et al., 2021b).

Table 12 provides an overview of the number of trained models for benchmarking performance of all models and all UDA algorithms on the *medium* difficulty domain shifts across all datasets.

Table 12: Overview of the benchmarking setup and number of trained models across all datasets.

Dataset	Models	UDA algorithms	λ values	# seeds	# models trained
Rolling	PointNet, GraphSAGE, Transolver	Deep Coral, CMD, DANN w/o UDA	$\{10^{-1}; 10^{-9}\}$	4 4	324 12
Forming	PointNet, GraphSAGE, Transolver	Deep Coral, CMD, DANN w/o UDA	$\{10^{-1}; 10^{-9}\}$	4 4	324 12
Motor	PointNet, GraphSAGE, Transolver	Deep Coral, CMD, DANN w/o UDA	$\{10^{-1}; 10^{-9}\}$	4 4	324 12
Heatsink	PointNet, Transover, UPT	Deep Coral, CMD, DANN w/o UDA	$\{10^2; 10^{-2}\}$	4 4	180 12
Sum					1,200

Additional Details. For the three smaller datasets, we use smaller networks, while for the large scale *heatsink design* dataset, we train larger model configurations to accommodate the increased data complexity. An overview of model sizes along with average training times per dataset is provided in Table 13. We also refer to the accompanying code repository for a complete listing of all model hyperparameters, where we provide all baseline configuration files and detailed step by step instructions for reproducibility of our results.

Another important detail is that, during training on the *heatsink design* dataset, we randomly subsample 16,000 nodes from the mesh in each training step to ensure computational tractability. However, all reported performance metrics are computed on the full resolution of the data without any subsampling.

E.2 EVALUATION METRICS

We report the RMSE for each predicted output field. For field *i*, the RMSE is defined as:

$$\text{RMSE}_{i}^{\text{field}} = \frac{1}{M} \sum_{m=1}^{M} \sqrt{\frac{1}{N_{m}} \sum_{n=1}^{N_{m}} \left(y_{m,n}^{(i)} - f(x)_{m,n}^{(i)}\right)^{2}},$$

where M is the number of test samples (graphs), N_m the number of nodes in graph m, $y_{m,n}^{(i)}$ the ground truth value of field i at node n of graph m, and $f(x)_{m,n}^{(i)}$ the respective model prediction.

For aggregated evaluation, we define the total Normalized RMSE (NRMSE) as:

$$NRMSE = \frac{1}{K} \sum_{i=1}^{K} NRMSE_{i}^{field},$$

where K is the number of predicted fields. For this metric, all individual field errors are computed on normalized fields before aggregation.

In addition to the error on the fields, we report the mean Euclidean error of the predicted node displacement. This is computed based on the predicted coordinates $\hat{\mathbf{c}}_{m,n} \in \mathbb{R}^d$ and the ground truth coordinates $\mathbf{c}_{m,n} \in \mathbb{R}^d$, where $d \in \{2,3\}$ is the spatial dimensionality, as follows:

$$\text{RMSE}^{\text{deformation}} = \frac{1}{M} \sum_{m=1}^{M} \sqrt{\frac{1}{N_m} \sum_{n=1}^{N_m} \left\| \mathbf{c}_{m,n} - \hat{\mathbf{c}}_{m,n} \right\|_2}.$$

E.3 COMPUTATIONAL RESOURCES AND TIMINGS

While generating the results reported on the *medium* difficulty level of our benchmark, we measured average training times per dataset and model architecture. While the total compute budget is difficult to estimate due to distributed training runs across various hardware setups, we report standardized average training times for 2000 epochs in Table 13, measured on a single NVIDIA H100 GPU using batch size of 16.

Table 13: Average training times (averaged for 2000 epochs) and parameter counts for each model on the *medium* difficulty benchmark tasks. Times are measured on a H100 GPU using a batch size of 16.

Dataset	Model	# parameters	Avg. training time (h)
	PointNet	0.3M	1.2
Rolling	GraphSAGE	0.2M	3
	Transolver	0.57M	2.1
	PointNet	0.3M	2.8
Forming	GraphSAGE	0.2M	8
	Transolver	0.57M	4.4
	PointNet	0.3M	5.6
Motor	GraphSAGE	0.2M	11.5
	Transolver	0.57M	6.5
	PointNet	1.08M	4.9
Heatsink	Transolver	4.07M	5.3
	UPT	5.77M	5.5

F DATASET DETAILS

F.1 HOT ROLLING

The hot rolling dataset represents a hot rolling process in which a metal slab undergoes plastic deformation to form a sheet metal product. The model considers a plane-strain representation of a heated steel slab segment with a core temperature $T_{\rm core}$ and a surface temperature $T_{\rm Surf}$, initially at thickness t, passing through a simplified roll stand with a nominal roll gap g (see Figure 2a). This roll gap effectively matches the exit thickness of the workpiece. Given the material properties, the initial temperature distribution over the slab thickness and the specified pass reduction, the model aims to capture the evolution of the thermo-mechanical state of the workpiece as it traverses the roll gap.

To reduce computational complexity, the analysis is confined to the vertical midplane along the rolling direction based on a plane-strain assumption. This is well justified by the high width-to-thickness ratio characteristic of the workpiece. Additionally, vertical symmetry is also exploited. Consequently, only the upper half of the workpiece and the upper work roll are modeled.

The workpiece is discretized using plane-strain, reduced-integration, quadrilateral elements. Mesh generation is fully automated, with the element size calibrated according to findings from a mesh convergence study. In terms of mechanical behavior, the workpiece is modeled as elasto-plastic with isotropic hardening, employing tabulated flow curves representative for a titanium alloy (Lesuer, 2000; Lu et al., 2018). The elastic modulus and flow stress are temperature dependent, with the latter also influenced by the plastic strain rate. In contrast, material density and Poisson's ratio are assumed to remain constant. The work roll with a diameter of 1000 mm is idealized as an analytically defined rigid body.

In addition to the mechanical behavior, the elements also feature a temperature degree of freedom that captures thermal phenomena, which are in turn fully coupled with the mechanical field. Heat conduction within the workpiece is governed by temperature dependent thermal conductivity and specific heat capacity. Heat transfer at the interface between the workpiece and the roll is modeled as proportional to the temperature difference between the contacting surfaces, using a heat transfer coefficient of $5 \text{ mW/mm}^2\text{K}$. The model also accounts for internal heat generation due to plastic

deformation, based on the standard assumption that 90% of plastic work is converted into heat. Additionally, all frictional energy is assumed to be fully transformed into heat and evenly divided between the workpiece and the roll. However, since the analysis focuses on the workpiece, only the portion of this heat entering the workpiece is considered.

The FE simulation is performed with the *Abaqus* explicit solver using a relatively high mass scaling factor of 100. This mass scaling proved to be a suitable choice for maintaining both computational efficiency and solution accuracy. The pre-processing, evaluation and post-processing of the simulations was automated in Python. A full factorial design of experiments was conducted by varying the parameters outlined in Table 14. Simulation outputs from Abaqus (.odb files) were converted to a more suitable .h5 format in post-processing, enabling seamless integration into the SIMSHIFT framework. All simulations were run on a Gigabyte Aorus 15P KD consumer laptop equipped with an Intel Core i7-11800H CPU (8 cores, 16 threads, 2.30–4.60 GHz), 16 GB DDR4 RAM at 3200 MHz and a 1 TB NVMe SSD. The single-core CPU time for one simulation was 25 seconds on average, depending on the mesh size and convergence speed.

Table 14: Input parameter ranges for the *hot rolling* simulations. Samples are generated by equally spacing each parameter within the specified range using the indicated number of steps, resulting in $5 \times 19 \times 10 \times 5 = 4750$ total samples.

Parameter	Description	Min	Max	Steps
t (mm)	Initial slab thickness.	50.0	183.3	5
reduction $(-)$	Reduction of initial slab thickness.	0.01	0.15	19
$T_{\rm core} (^{\circ}C)$	Core slab temperature.	900.0	1000.0	10
T_{surf} (${}^{\circ}C$)	Surface slab temperature.	900.0	1077.77	5

F.2 SHEET METAL FORMING

For the *sheet metal forming* dataset, a w-shaped bending process was selected due to its complex contact interactions and the highly nonlinear progression of bending forces. For this purpose, a parameterized 2D FE model of the process was developed using the commercial FEM software *Abaqus* and its implicit solver, with the simulation pipeline implemented in Python. The initial configuration of the finite element model is shown in Figure 29 and described below.

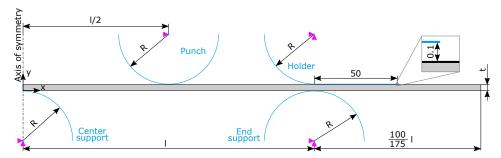


Figure 29: Bending process abstraction, initial configuration.

Due to geometric and loading symmetry, only the right half of the sheet with a thickness t was modeled. The die and punch were idealized as rigid circular segments with a shared radius r. Additionally, a rigid blank holder comprising an arc and a straight segment was positioned $0.1~\mathrm{mm}$ above the sheet to maintain contact and restrain vertical motion. The required sheet length was determined by the support span l, enabling material flow toward the center in response to the downward motion of the punch.

The sheet was discretized using bilinear, plane-strain quadrilateral elements with reduced integration and hourglass control (Abaqus element type CPE4R). A prior mesh convergence study indicated that accurate simulation results require a minimum of 10 element rows across the sheet thickness. The element size was fixed at $0.125 \times 0.1~\mathrm{mm}$ to ensure a uniform aspect ratio, constraining the sheet thickness to $t>1~\mathrm{mm}$.

The sheet material was modeled as elastoplastic with von Mises plasticity and linear isotropic hardening. The following properties were assigned: Young's modulus of 210 GPa, Poisson's ratio of 0.3, yield stress of 410 MPa, and hardening modulus of 2268 MPa.

For all contact interfaces, a normal contact formulation with surface-to-surface discretization, penalty enforcement, and finite-sliding tracking was employed. Tangential contact was modeled via a Coulomb friction law with a coefficient μ .

The supports and blank holder were fixed by constraining horizontal and vertical translations as well as in-plane rotations. These constraints were applied at the centroid of each arc segment, representing the reference point for the respective rigid body. The punch was similarly constrained against horizontal movement and rotation but retained vertical mobility. The deformed configuration following a vertical displacement U of the punch is illustrated in Figure 30.

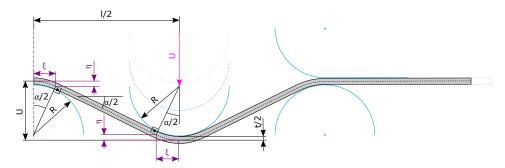


Figure 30: Bending process abstraction, deformed configuration.

A full factorial design of experiments was conducted by varying the parameters outlined in Table 15. As for the hot rolling simulations, outputs from Abaqus (.odb files) were converted to .h5 format in post-processing, to integrate them into the SIMSHIFT framework. All simulations were run on a Gigabyte Aorus 15P KD consumer laptop equipped with an Intel Core i7-11800H CPU (8 cores, 16 threads, 2.30–4.60 GHz), 16 GB DDR4 RAM at 3200 MHz and a 1 TB NVMe SSD. The single-core CPU time for one simulation run was 300 seconds on average, depending on mesh size and convergence speed.

Table 15: Input parameter ranges for the *sheet metal forming* simulations. Samples are generated by equally spacing each parameter within the specified range using the indicated number of steps, resulting in $17 \times 13 \times 3 \times 5 = 3315$ total samples.

Parameter	Description	Min	Max	Steps
r (mm)	Roll radius.	10.0	50.0	17
t(mm)	Sheet thickness.	2.0	5.0	13
l(mm)	Sheet length.	175.0	350.0	3
$\mu(-)$	Friction coefficient between the sheet and the rolls.	0.1	0.5	5

F.3 ELECTRIC MOTOR DESIGN

The *electric motor design* dataset includes a structural FE simulation of a rotor within electric machinery, subjected to mechanical loading at burst speed. The rotor topology is modeled after the motor architecture of the 2010 Toyota Prius (Burress et al., 2011), an industry-recognized benchmark frequently used for validation and comparison in academic and industrial research. The Prius rotor topology is based on a V-shaped magnet configuration as shown in Figure 31.

Structural rotor simulations are essential in multi-physics design optimization, where motor performance is evaluated across multiple domains including electromagnetic, thermal, acoustic, and structural. Using a design optimization framework, stator and rotor design are iteratively refined to identify Pareto-optimal solutions based on objectives such as efficiency, torque, weight, and speed. In this process, the structural FE model predicts stress and deformation due to loading ensuring the rotor's structural integrity.

The set up and execution of the structural simulations for this dataset are automated and implemented in the open source design optimization framework *SyMSpace*³. The FE simulation of the rotor is performed using a mixed 2D plane stress and plane strain formulation with triangular elements. To enhance computational efficiency, geometric symmetry is exploited and only a 1/16 sector of the full rotor is modeled. The mechanical simulation is static and evaluates the rotor under centrifugal loading, incorporating press-fit conditions between the rotor core and shaft, as well as contact interactions between the rotor core and embedded magnets.

An elastic material behavior is employed for all components, including the rotor core, shaft, and magnets. Material properties are summarized in Table 16. Based on the parametrized CAD model of the rotor topology, the geometry is automatically meshed using *Netgen*⁴. The design optimization tool also automatically identifies nodes for boundary conditions and contact surfaces and applies the corresponding constraints and interactions required for the simulation. The implicit FE solver *HOTINT* is used to compute the quasi-static response of the system, providing local stress and strain fields across the rotor topology.

Table 16: Material parameters for the structural *electric motor design* simulations.

	Rotor Core	Rotor Shaft	Permanent Magnet
Material	NO27-14 Y420HP	42CrMo4	BMN-40SH
Density (kg/dm ³)	7.6	7.72	7.55
Possions ratio (-)	0.29	0.3	0.24
Young's Modulus (kN/mm ²)	185.0	210.0	175.0
Tensile Strength (kN/mm ²)	550.0	850.0	250.0

To generate the electric motor dataset, a comprehensive motor optimization study was conducted using *SyMSpace*, based on design specifications of the 2010 Toyota Prius. The optimization aimed to minimize multiple performance metrics, including motor mass, material costs, rotor torque ripple, motor losses, coil temperature, stator terminal current, and elastic rotor deformation. A genetic algorithm was employed to explore the design space and identify Pareto-optimal solutions. In the process, 3,196 motor configurations were evaluated by varying, among other factors, the rotor's topological parameters within the bounds specified in Table 17. The outputs of the structural simulations were generated in .vtk format and then stored in .h5 files, allowing direct integration into the SIMSHIFT framework. Each structural simulation required approximately 4 to 5 minutes of single-core CPU time on a Intel Core i9-14900KS processor (24 Cores, 3200 MHz), depending on convergence speed of the contact algorithm.

³https://symspace.lcm.at/

⁴https://ngsolve.org/

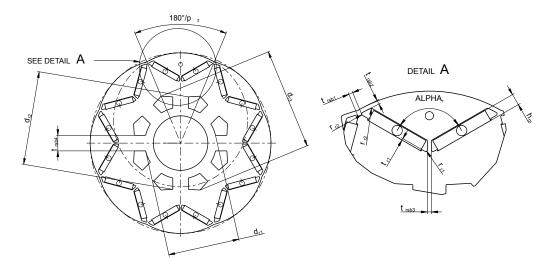


Figure 31: Technical drawing of the electrical motor. Sampling ranges for the shown parameters can be found in Table 17.

Table 17: Input parameters for the *electric motor design* simulations. Since the design space was explored by a genetic algorithm, the parameters are not uniformly sampled as in the previous simulation scenarios. In total, 3196 simulations were performed.

Parameter	Description	Min	Max
$d_{si} (mm)$	Stator inner diameter.	150.0	180.0
$h_m (mm)$	Magnet height.	6.0	9.0
α_r (\circ)	Angle between magnets.	120.0	160.0
$t_{r1} \ (mm)$	Magnet step.	1.0	5.0
$r_{r1} (mm)$	Rotor slot fillet radius 1.	0.5	2.5
$r_{r2} \ (mm)$	Rotor slot fillet radius 2.	0.5	3.5
$r_{r3} \ (mm)$	Rotor slot fillet radius 3.	0.5	5.0
$r_{r4} (mm)$	Rotor slot fillet radius 4.	0.5	3.0
$t_{rsb1} \ (mm)$	Thickness saturation bar 1.	4.0	12.0
$t_{rsb2} \ (mm)$	Thickness saturation bar 2.	1.0	3.0
$t_{rsb3} \ (mm)$	Thickness saturation bar 3.	1.2	4.0
$t_{rsb4} (mm)$	Thickness saturation bar 4.	5.0	12.0
$d_{r1} (mm)$	Rotor slot diameter 1.	60.0	80.0
$d_{r2} (mm)$	Rotor slot diameter 2.	80.0	120.0
d_{r3} (mm)	Rotor slot diameter 3.	100.0	125.0

F.4 HEATSINK DESIGN

The *heatsink design* dataset consists of heatsink geometries similar to the example shown in Figure 32, placed centrally at the bottom of a surrounding box-shaped domain filled with air. The dimensions of the surrounding enclosure are $0.14 \text{ m} \times 0.14 \text{ m} \times 0.5 \text{ m}$ (length × width × height).

The geometric configuration of each heatsink is defined by several parameters, which were varied within specified bounds for the design study. These parameters and their corresponding value ranges are summarized in Table 18. A total of 460 simulation cases were generated, with non-uniform sampling across the parameter space.

The dataset was generated using CFD simulations based on the Reynolds-Averaged Navier-Stokes (RANS) equations coupled with the energy equation. All simulations were conducted in the open-source CFD suite *OpenFOAM 9*.

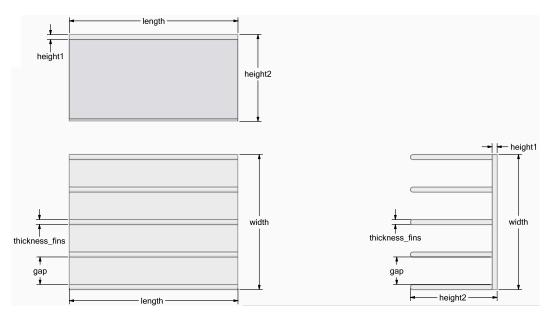


Figure 32: Technical drawing of the solid body in the *heatsink design* dataset. Some of the shown parameters are varied for data generation (see Table 18).

Table 18: Geometric and physical parameters of the *heatsink design* simulations. The variable parameters were not uniformly sampled. In total, 460 simulations were performed.

Parameter	Description	fixed Value	Min	Max
length (m)	Heatsink length	0.1	-	-
width (m)	Heatsink width	0.08	-	-
height1 (m)	Baseplate height	0.003	-	-
T(amb) (K)	Ambient Temperature	300	-	-
fins $(-)$	Number of fins	-	5	14
gap (m)	Gap between fins	-	0.0023	0.01625
thickness_fins (m)	Thickness of fins	-	0.003	0.004
height2 (m)	Heatsink height	=	0.053	0.083
T (solid) (K)	Temperature of the solid fins	-	340	400

The computational domain was discretized using a finite volume method with second-order spatial discretization schemes. A structured hexahedral background mesh was generated with the blockMesh utility in OpenFOAM, followed by mesh refinement using snappyHexMesh to accurately resolve the heatsink structure defined in STL format.

To simulate buoyancy driven natural convection, the buoyantSimpleFoam solver was employed. This solver is designed for steady state, compressible, buoyant flows, using the SIMPLE algorithm for pressure-momentum coupling, extended with under relaxation techniques to enhance numerical stability and robust convergence.

Boundary conditions were applied as follows:

- Walls of the surrounding: no-slip velocity condition with fixed ambient temperature as defined in Table 18.
- Walls of the heatsink: no-slip velocity condition with solid temperature within the range specified for parameter T (solid) in Table 18.

Given the turbulent nature of the flow, the RANS equations were closed using the SST $k-\omega$ turbulence model (Menter et al., 2003). Near-wall regions were modeled using a y^+ -insensitive near-wall

treatment, allowing accurate resolution of boundary layers without the need for excessively fine meshes.

A mesh convergence study was conducted to ensure numerical accuracy. Depending on mesh resolution, each simulation required approximately 11 to 18 hours of single-core CPU time on an Intel Core i9-14900KS processor (24 cores, 3.2 GHz).

G ABLATION STUDIES

In the following sections, we present ablations on the SIMSHIFT framework.

G.1 GEOMETRIC ENCODING

The design concept of SIMSHIFT is to allow plug-in integration of any UDA algorithm and model architecture, as long as the model can be conditioned in some way (see Figure 1). However, explicitly conditioning models on scalar geometric parameters is not the only option: for instance, domain-specific information may be encoded implicitly in the mesh itself. To investigate this, we provide an ablation in which the model encodes the mesh directly and is not explicitly conditioned on the scalar parameters. Specifically, we replace the feed-forward conditioning network with a geometric PointNet based encoder to embed the input mesh into a global latent vector, on which UDA is then performed.

We report results of this setup on the *electric motor design* dataset. The setup follows the benchmarking procedure described in Section 4 and Appendix E.1: for each UDA algorithm, we train across 9 different regularizer strengths and 4 random seeds.

Table 19: RMSE (mean \pm std over 4 seeds) on the *electric motor design* dataset when using a PointNet geometry encoder. Values are target domain errors (lower is better). Bold marks the overall best model + UDA algorithm + model selection combination. For each architecture, the unregularized baseline row is shaded beige, whereas the best UDA + selection within that architecture is underlined and shaded green.



Table 19 shows that UDA algorithms can boost target performance compared to the unregularized baseline model. However compared to our chosen benchmarking design in Table 7, both the performance of the unregularized baseline as well as the one of the best performing UDA method is worse, which supports our choice of explicitly conditioning on scalar parameters in the main benchmark.

G.2 TWO-DIMENSIONAL SHIFTS

Defining shifts based on one parameter allows for controlled experiments, also given that the parameters were picked based on preliminary experiments (see Appendix C) and consultation with domain experts. In real-world scenarios, however, distribution shifts often affect multiple parameters simultaneously rather than only a single one. It is therefore important to investigate the performance of the benchmarked UDA algorithms under multidimensional parameter shifts. As a step in this direction, we provide an ablation on the *electric motor design* dataset for a two-dimensional parameter shift.

To be concise, we jointly shift the rotor slot diameter d_{r3} (parameter shift in the main benchmark) and the angle between the magnets α_r . Table 20 shows the corresponding two-dimensional distribution shift between the source and the target domain.

We train all models with each UDA algorithm following the procedure in Section 4.5 and Appendix E.1, and report the results in Table 21.

Table 20: Parameter ranges for the two-dimensional distribution shift on the *electric motor design* dataset.

Parameter	Source range	Target range
Rotor slot diameter 3 d_{r3} (mm)	[100, 120)	[120, 126]
Angle between magnets α_r $(^{\circ})$	[119, 153)	[153, 170]

Table 21: RMSE (mean \pm std over 4 seeds) on the *electric motor design* dataset at with a two-dimensional distribution shift in parameter space. Values are target domain errors (lower is better). Bold marks the overall best model + UDA algorithm + model selection combination. For each architecture, the unregularized baseline row is shaded beige, whereas the best UDA + selection within that architecture is underlined and shaded green.

Model	DA	Model	All Fields Non	malized Avg (-)	Deform	ation (m)	Logarithmic Strain (×10 ⁻²)		Principal St	rain (×10 ⁻²)	Stress	(MPa)	Cauchy S	tress (MPa)	Mises Str	ress (MPa)	Principal S	tress (MPa)	Total Strain (×10 ⁻²)	
	Algorithm	Selection	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT	SRC	TGT
		-	$0.324(\pm0.025)$	0.379(±0.017)	$0.001(\pm0.000)$	$0.000(\pm0.000)$	$0.008(\pm0.001)$	$0.010(\pm0.000)$	$0.008(\pm0.001)$	0.009(±0.000)	11.027(±0.778)	12.953(±0.502)	11.045(±0.775)	12.984(±0.507)	24.438(±1.730)	29.664(±0.853)	12.510(±0.814)	14.816(±0.481)	$0.007(\pm0.001)$	$0.009(\pm0.000)$
	DANN DANN DANN DANN	DEV IWV SB TB	0.284(±0.013) 0.284(±0.012) 0.284(±0.008) 0.299(±0.013)	$0.382(\pm 0.067)$ $0.342(\pm 0.010)$ $0.353(\pm 0.019)$ $0.343(\pm 0.008)$	$0.001(\pm 0.000)$ $0.001(\pm 0.000)$ $0.001(\pm 0.000)$ $0.001(\pm 0.000)$	$0.000(\pm0.000)$ $0.000(\pm0.000)$ $0.000(\pm0.000)$ $0.000(\pm0.000)$	$\begin{array}{c} 0.008(\pm 0.000) \\ 0.008(\pm 0.000) \\ 0.008(\pm 0.000) \\ 0.008(\pm 0.000) \end{array}$	$0.011(\pm 0.002)$ $0.009(\pm 0.000)$ $0.010(\pm 0.001)$ $0.009(\pm 0.000)$	$\begin{array}{c} 0.007(\pm 0.000) \\ 0.007(\pm 0.000) \\ 0.007(\pm 0.000) \\ 0.008(\pm 0.000) \end{array}$	0.010(±0.002) 0.009(±0.000) 0.010(±0.001) 0.009(±0.000)	10.412(±0.403) 10.439(±0.388) 10.402(±0.294) 10.957(±0.457)	$\begin{array}{c} 14.196(\pm 2.641) \\ 12.620(\pm 0.368) \\ 13.056(\pm 0.748) \\ 12.640(\pm 0.246) \end{array}$	10.431(±0.404) 10.457(±0.390) 10.421(±0.295) 10.977(±0.459)	14.224(±2.646) 12.647(±0.370) 13.084(±0.749) 12.666(±0.246)	22.782(±0.904) 22.994(±0.840) 22.890(±0.806) 24.296(±1.127)	32.960(±6.648) 28.910(±1.073) 30.202(±1.998) 28.930(±0.399)	11.857(±0.406) 11.918(±0.354) 11.882(±0.288) 12.430(±0.514)	16.314(±3.055) 14.474(±0.420) 14.993(±0.825) 14.460(±0.225)	0.007(±0.000) 0.007(±0.000) 0.007(±0.000) 0.007(±0.000)	$0.010(\pm 0.002)$ $0.009(\pm 0.000)$ $0.009(\pm 0.001)$ $0.009(\pm 0.000)$
GraphSAGE	CMD CMD CMD CMD	DEV IWV SB TB	$\begin{array}{c} 0.306(\pm 0.003) \\ 0.287(\pm 0.023) \\ 0.278(\pm 0.021) \\ 0.290(\pm 0.027) \end{array}$	$\begin{array}{c} 0.359(\pm 0.011) \\ 0.347(\pm 0.023) \\ 0.340(\pm 0.018) \\ 0.343(\pm 0.017) \end{array}$	$\begin{array}{c} 0.001(\pm 0.000) \\ 0.001(\pm 0.000) \\ 0.001(\pm 0.000) \\ 0.001(\pm 0.000) \end{array}$	$\begin{array}{c} 0.000(\pm 0.000) \\ 0.000(\pm 0.000) \\ 0.000(\pm 0.000) \\ 0.000(\pm 0.000) \end{array}$	$\begin{array}{c} 0.008(\pm 0.000) \\ 0.008(\pm 0.001) \\ 0.008(\pm 0.001) \\ 0.008(\pm 0.001) \end{array}$	$\begin{array}{c} 0.010(\pm 0.000) \\ 0.010(\pm 0.001) \\ 0.009(\pm 0.000) \\ 0.009(\pm 0.000) \end{array}$	$\begin{array}{c} 0.008(\pm 0.000) \\ 0.007(\pm 0.001) \\ 0.007(\pm 0.001) \\ 0.007(\pm 0.001) \end{array}$	$\begin{array}{c} 0.010(\pm0.000) \\ 0.009(\pm0.001) \\ 0.009(\pm0.000) \\ 0.009(\pm0.000) \end{array}$	$\begin{array}{c} 11.195(\pm0.100) \\ 10.515(\pm0.823) \\ 10.176(\pm0.756) \\ 10.609(\pm0.933) \end{array}$	$\begin{array}{c} 13.203(\pm 0.465) \\ 12.812(\pm 0.849) \\ 12.534(\pm 0.642) \\ 12.608(\pm 0.581) \end{array}$	$\begin{array}{c} 11.220(\pm0.104) \\ 10.536(\pm0.827) \\ 10.192(\pm0.754) \\ 10.627(\pm0.933) \end{array}$	$\begin{array}{c} 13.230(\pm0.468) \\ 12.837(\pm0.850) \\ 12.560(\pm0.644) \\ 12.634(\pm0.582) \end{array}$	24.742(±0.244) 23.067(±2.069) 22.292(±1.968) 23.292(±2.338)	30.241(±1.550) 29.412(±2.331) 28.655(±1.563) 28.678(±1.334)	12.652(±0.144) 11.924(±0.882) 11.580(±0.827) 12.061(±1.045)	$\begin{array}{c} 15.065(\pm 0.552) \\ 14.653(\pm 0.953) \\ 14.341(\pm 0.683) \\ 14.410(\pm 0.634) \end{array}$	$\begin{array}{c} 0.008(\pm 0.000) \\ 0.007(\pm 0.001) \\ 0.007(\pm 0.001) \\ 0.007(\pm 0.001) \\ 0.007(\pm 0.001) \end{array}$	$0.009(\pm 0.000)$ $0.009(\pm 0.001)$ $0.009(\pm 0.000)$ $0.009(\pm 0.000)$
	Deep Coral Deep Coral Deep Coral Deep Coral	DEV IWV SB TB	0.303(±0.006) 0.276(±0.002) 0.278(±0.002) 0.281(±0.012)	$0.358(\pm 0.002)$ $0.334(\pm 0.007)$ $0.339(\pm 0.007)$ $0.336(\pm 0.010)$	0.001(±0.000) 0.001(±0.000) 0.001(±0.000) 0.001(±0.000)	0.000(±0.000) 0.000(±0.000) 0.000(±0.000) 0.000(±0.000)	$0.008(\pm 0.000)$ $0.008(\pm 0.000)$ $0.008(\pm 0.000)$ $0.008(\pm 0.000)$	0.010(±0.000) 0.009(±0.000) 0.009(±0.000) 0.009(±0.000)	0.008(±0.000) 0.007(±0.000) 0.007(±0.000) 0.007(±0.000)	0.010(±0.000) 0.009(±0.000) 0.009(±0.000) 0.009(±0.000)	11.090(±0.161) 10.142(±0.089) 10.223(±0.101) 10.294(±0.385)	13.184(±0.057) 12.322(±0.261) 12.518(±0.253) 12.367(±0.350)	11.110(±0.160) 10.161(±0.089) 10.242(±0.101) 10.313(±0.385)	13.210(±0.057) 12.348(±0.262) 12.544(±0.253) 12.393(±0.351)	24.492(±0.274) 22.259(±0.186) 22.490(±0.309) 22.589(±0.841)	30 206(±0.372) 28 271(±0.656) 28 .784(±0.715) 28 .284(±0.679)	12.531(±0.176) 11.540(±0.087) 11.644(±0.129) 11.697(±0.391)	15.040(±0.128) 14.144(±0.288) 14.374(±0.295) 14.174(±0.346)	0.007(±0.000) 0.007(±0.000) 0.007(±0.000) 0.007(±0.000)	0.009(±0.000) 0.008(±0.000) 0.009(±0.000) 0.008(±0.000)
			$0.404(\pm0.061)$	$0.441(\pm 0.064)$	$0.001(\pm0.000)$	$0.001(\pm0.000)$	$0.010(\pm0.001)$	$0.011(\pm0.002)$	$0.010(\pm0.002)$	$0.011(\pm 0.002)$	13.536(±2.116)	14.947(±2.259)	13.530(±2.091)	14.957(±2.241)	29.759(±4.853)	33.860(±5.135)	15.433(±2.479)	17.310(±2.587)	$0.009(\pm0.001)$	$0.010(\pm 0.001)$
	DANN DANN DANN DANN	DEV IWV SB TB	0.333(±0.079) 0.344(±0.080) 0.320(±0.066) 0.320(±0.066)	$0.383(\pm 0.077)$ $0.379(\pm 0.074)$ $0.355(\pm 0.060)$ $0.355(\pm 0.060)$	$0.001(\pm 0.000)$ $0.001(\pm 0.000)$ $0.001(\pm 0.000)$ $0.001(\pm 0.000)$	0.000(±0.000) 0.000(±0.000) 0.000(±0.000) 0.000(±0.000)	$0.009(\pm0.002)$ $0.009(\pm0.002)$ $0.009(\pm0.002)$ $0.009(\pm0.002)$	0.010(±0.002) 0.010(±0.002) 0.010(±0.002) 0.010(±0.002)	$0.009(\pm0.002)$ $0.009(\pm0.002)$ $0.008(\pm0.002)$ $0.008(\pm0.002)$	0.010(±0.002) 0.010(±0.002) 0.009(±0.002) 0.009(±0.002)	12.052(±2.867) 12.476(±2.926) 11.565(±2.387) 11.565(±2.387)	14.014(±2.877) 13.850(±2.763) 12.905(±2.214) 12.905(±2.214)	12.056(±2.846) 12.469(±2.909) 11.568(±2.361) 11.568(±2.361)	14.030(±2.866) 13.860(±2.750) 12.921(±2.199) 12.921(±2.199)	26.183(±6.669) 27.020(±6.762) 25.087(±5.559) 25.087(±5.559)	31.715(±6.844) 30.977(±6.339) 28.844(±5.027) 28.844(±5.027)	13.702(±3.348) 14.155(±3.392) 13.108(±2.779) 13.108(±2.779)	16.135(±3.416) 15.937(±3.286) 14.832(±2.609) 14.832(±2.609)	0.008(±0.002) 0.008(±0.002) 0.008(±0.002) 0.008(±0.002)	0.010(±0.002) 0.009(±0.002) 0.009(±0.001) 0.009(±0.001)
PointNet	CMD CMD CMD CMD	DEV IWV SB TB	0.338(±0.064) 0.368(±0.031) 0.355(±0.037) 0.334(±0.059)	$0.398(\pm 0.042)$ $0.414(\pm 0.009)$ $0.403(\pm 0.032)$ $0.380(\pm 0.040)$	$0.001(\pm0.000)$ $0.001(\pm0.000)$ $0.001(\pm0.000)$ $0.001(\pm0.000)$	0.000(±0.000) 0.001(±0.000) 0.000(±0.000) 0.000(±0.000)	$0.009(\pm 0.002)$ $0.010(\pm 0.001)$ $0.010(\pm 0.001)$ $0.009(\pm 0.002)$	0.011(±0.001) 0.011(±0.000) 0.011(±0.001) 0.010(±0.001)	0.009(±0.002) 0.009(±0.001) 0.009(±0.001) 0.009(±0.002)	0.011(±0.001) 0.011(±0.000) 0.011(±0.001) 0.010(±0.001)	12.173(±2.395) 13.275(±1.104) 12.821(±1.387) 11.996(±2.224)	14.649(±1.516) 15.175(±0.343) 14.758(±1.300) 13.845(±1.431)	12.178(±2.381) 13.277(±1.085) 12.828(±1.367) 12.008(±2.217)	14.665(±1.508) 15.183(±0.328) 14.773(±1.293) 13.868(±1.418)	26.658(±5.253) 28.952(±2.713) 27.878(±3.425) 26.384(±4.897)	33.615(±3.357) 34.480(±1.329) 33.618(±3.351) 31.387(±3.242)	13.839(±2.688) 15.004(±1.304) 14.537(±1.634) 13.665(±2.474)	16.815(±1.806) 17.411(±0.451) 16.989(±1.460) 15.919(±1.717)	0.008(±0.002) 0.009(±0.001) 0.009(±0.001) 0.008(±0.001)	0.010(±0.001) 0.010(±0.000) 0.010(±0.001) 0.009(±0.001)
	Deep Coral Deep Coral Deep Coral Deep Coral	IWV SB	0.350(±0.039) 0.351(±0.042) 0.358(±0.042) 0.341(±0.045)	0.383(±0.038) 0.384(±0.040) 0.393(±0.042) 0.371(±0.045)	0.001(±0.000) 0.001(±0.000) 0.001(±0.000) 0.001(±0.000)	0.000(±0.000) 0.000(±0.000) 0.000(±0.000) 0.000(±0.000)	$0.009(\pm0.001)$ $0.009(\pm0.001)$ $0.010(\pm0.001)$ $0.009(\pm0.001)$	0.010(±0.001) 0.010(±0.001) 0.011(±0.001) 0.010(±0.001)	0.009(±0.001) 0.009(±0.001) 0.009(±0.001) 0.009(±0.001)	0.010(±0.001) 0.010(±0.001) 0.010(±0.001) 0.010(±0.001)	12.608(±1.434) 12.664(±1.539) 12.915(±1.522) 12.259(±1.691)	13.923(±1.438) 13.988(±1.552) 14.274(±1.566) 13.481(±1.683)	12.604(±1.396) 12.664(±1.506) 12.920(±1.505) 12.262(±1.653)	13.944(±1.431) 13.997(±1.523) 14.296(±1.557) 13.505(±1.673)	27.257(±3.653) 27.397(±3.913) 28.150(±3.736) 26.681(±3.884)	31.231(±3.564) 31.369(±3.801) 32.076(±3.728) 30.213(±3.909)	14.263(±1.750) 14.340(±1.891) 14.713(±1.880) 13.887(±1.948)	16.034(±1.720) 16.108(±1.847) 16.505(±1.872) 15.502(±1.950)	0.008(±0.001) 0.009(±0.001) 0.009(±0.001) 0.008(±0.001)	0.009(±0.001) 0.009(±0.001) 0.010(±0.001) 0.009(±0.001)
		-	$0.134(\pm0.009)$	0.146(±0.010)	$0.001(\pm0.000)$	$0.000(\pm0.000)$	$0.003(\pm0.000)$	$0.003(\pm0.000)$	0.003(±0.000)	$0.003(\pm0.000)$	4.411(±0.261)	$4.836(\pm0.289)$	4.419(±0.262)	4.845(±0.290)	9.486(±0.610)	10.368(±0.615)	5.252(±0.329)	$5.820(\pm 0.350)$	$0.003(\pm0.000)$	$0.003(\pm0.000)$
	DANN DANN DANN DANN	DEV IWV SB TB	0.111(±0.008) 0.107(±0.008) 0.099(±0.004) 0.102(±0.004)	0.121(±0.009) 0.115(±0.009) 0.110(±0.005) 0.111(±0.006)	0.001(±0.000) 0.001(±0.000) 0.001(±0.000) 0.001(±0.000)	0.000(±0.000) 0.000(±0.000) 0.000(±0.000) 0.000(±0.000)	$0.003(\pm0.000)$ $0.003(\pm0.000)$ $0.003(\pm0.000)$ $0.003(\pm0.000)$	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	0.003(±0.000) 0.003(±0.000) 0.002(±0.000) 0.003(±0.000)	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	3.920(±0.260) 3.817(±0.310) 3.549(±0.157) 3.644(±0.124)	4.347(±0.333) 4.114(±0.336) 3.945(±0.186) 3.968(±0.196)	3.927(±0.260) 3.824(±0.311) 3.555(±0.157) 3.650(±0.124)	4.355(±0.333) 4.121(±0.336) 3.952(±0.186) 3.975(±0.197)	8.341(±0.634) 8.048(±0.705) 7.368(±0.389) 7.540(±0.360)	9.329(±0.814) 8.699(±0.742) 8.361(±0.440) 8.386(±0.453)	4.688(±0.317) 4.569(±0.328) 4.228(±0.188) 4.332(±0.204)	5.236(±0.393) 4.982(±0.389) 4.772(±0.215) 4.802(±0.232)	0.003(±0.000) 0.003(±0.000) 0.002(±0.000) 0.002(±0.000)	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)
Transolver	CMD CMD CMD CMD	DEV IWV SB TB	0.107(±0.008) 0.102(±0.006) 0.104(±0.005) 0.101(±0.005)	$0.119(\pm 0.009)$ $0.112(\pm 0.007)$ $0.115(\pm 0.008)$ $0.110(\pm 0.008)$	$0.001(\pm0.000)$ $0.001(\pm0.000)$ $0.001(\pm0.000)$ $0.001(\pm0.000)$	0.000(±0.000) 0.000(±0.000) 0.000(±0.000) 0.000(±0.000)	$0.003(\pm0.000)$ $0.003(\pm0.000)$ $0.003(\pm0.000)$ $0.003(\pm0.000)$	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	3.796(±0.229) 3.631(±0.218) 3.701(±0.211) 3.570(±0.169)	4.251(±0.316) 3.990(±0.264) 4.100(±0.282) 3.922(±0.274)	3.803(±0.229) 3.638(±0.219) 3.707(±0.211) 3.577(±0.169)	4.259(±0.317) 3.998(±0.264) 4.107(±0.283) 3.929(±0.275)	8.091(±0.555) 7.714(±0.389) 7.820(±0.377) 7.658(±0.316)	9.037(±0.762) 8.407(±0.616) 8.685(±0.662) 8.303(±0.630)	4.591(±0.277) 4.402(±0.203) 4.463(±0.187) 4.344(±0.200)	5.141(±0.363) 4.838(±0.300) 4.973(±0.325) 4.761(±0.334)	0.002(±0.000) 0.002(±0.000) 0.002(±0.000) 0.002(±0.000)	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)
	Deep Coral Deep Coral Deep Coral Deep Coral	IWV SB	0.107(±0.004) 0.105(±0.004) 0.106(±0.004) 0.105(±0.003)	$0.116(\pm 0.005)$ $0.115(\pm 0.003)$ $0.116(\pm 0.003)$ $0.113(\pm 0.003)$	0.001(±0.000) 0.001(±0.000) 0.001(±0.000) 0.001(±0.000)	$0.000(\pm0.000)$ $0.000(\pm0.000)$ $0.000(\pm0.000)$ $0.000(\pm0.000)$	$\begin{array}{c} 0.003(\pm 0.000) \\ 0.003(\pm 0.000) \\ 0.003(\pm 0.000) \\ 0.003(\pm 0.000) \end{array}$	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	$\begin{array}{c} 0.003(\pm0.000) \\ 0.003(\pm0.000) \\ 0.003(\pm0.000) \\ 0.003(\pm0.000) \end{array}$	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)	3.794(±0.126) 3.721(±0.142) 3.764(±0.104) 3.701(±0.089)	4.121(±0.180) 4.113(±0.087) 4.157(±0.104) 4.016(±0.109)	3.800(±0.127) 3.727(±0.142) 3.770(±0.104) 3.707(±0.089)	4.128(±0.180) 4.119(±0.088) 4.164(±0.104) 4.023(±0.100)	8.042(±0.300) 7.854(±0.253) 7.949(±0.270) 7.933(±0.293)	8.790(±0.414) 8.748(±0.210) 8.832(±0.205) 8.514(±0.229)	4.571(±0.142) 4.482(±0.153) 4.527(±0.145) 4.508(±0.136)	5.033(±0.197) 4.994(±0.131) 5.040(±0.139) 4.891(±0.114)	0.002(±0.000) 0.002(±0.000) 0.002(±0.000) 0.002(±0.000)	0.003(±0.000) 0.003(±0.000) 0.003(±0.000) 0.003(±0.000)

Comparing these results with the original one-dimensional shift (Table 7), two observations stand out: (i) For two out of three architectures, both the unregularized baseline and the best UDA algorithm and model selection combination exhibit higher errors in the average field NRMSE than in the one-dimensional shift setting, confirming that the two-dimensional shift is a more challenging task. (ii) The relative improvements over the unregularized baselines are larger, indicating that UDA training provides greater benefits under this more challenging distribution shift.

These findings highlight the potential of UDA to handle increasingly complex distribution shifts, underscoring its practical relevance for real-world applications.