



Contents lists available at ScienceDirect

Journal of the Mechanical Behavior of Biomedical Materials

journal homepage: <http://www.elsevier.com/locate/jmbbm>

Exploring the potential of transfer learning for metamodels of heterogeneous material deformation

E. Lejeune^{*}, B. Zhao

Department of Mechanical Engineering, Boston University, Boston, MA, 02215, USA

ARTICLE INFO

Keywords:

Soft tissue mechanics
Benchmark data
Machine learning

ABSTRACT

From the nano-scale to the macro-scale, biological tissue is spatially heterogeneous. Even when tissue behavior is well understood, the exact subject specific spatial distribution of material properties is often unknown. And, when developing computational models of biological tissue, it is usually prohibitively computationally expensive to simulate every plausible spatial distribution of material properties for each problem of interest. Therefore, one of the major challenges in developing accurate computational models of biological tissue is capturing the potential effects of this spatial heterogeneity. Recently, machine learning based metamodels have gained popularity as a computationally tractable way to overcome this problem because they can make predictions based on a limited number of direct simulation runs. These metamodels are promising, but they often still require a high number of direct simulations to achieve an acceptable performance. Here we show that transfer learning, a strategy where knowledge gained while solving one problem is transferred to solving a different but related problem, can help overcome this limitation. Critically, transfer learning can be used to leverage both low-fidelity simulation data and simulation data that is the outcome of solving a different but related mechanical problem. In this paper, we extend Mechanical MNIST, our open source benchmark dataset of heterogeneous material undergoing large deformation, to include a selection of low-fidelity simulation results that require $\approx 2 - 4$ orders of magnitude less CPU time to run. Then, we show that transferring the knowledge stored in metamodels trained on these low-fidelity simulation results can vastly improve the performance of metamodels used to predict the results of high-fidelity simulations. In the most dramatic examples, metamodels trained on 100 high fidelity simulations but pre-trained on 60,000 low-fidelity simulations achieves nearly the same test error as metamodels trained on 60,000 high-fidelity simulations (1 – 1.5% mean absolute percent error). In addition, we show that transfer learning is an effective method for leveraging data from different load cases, and for leveraging low-fidelity two-dimensional simulations to predict the outcomes of high-fidelity three-dimensional simulations. Looking forward, we anticipate that transfer learning will enable us to better capture the influence of tissue spatial heterogeneity on the mechanical behavior of biological materials across multiple different domains.

1. Introduction

Mechanical models of biological materials are useful for applications ranging from understanding the role of mechanics during development to predicting the outcomes of possible medical interventions (Ambrosi et al., 2011; Lejeune et al., 2016). And, the complex heterogeneous behavior of biological materials makes constructing relevant mechanical models a both challenging and interesting problem (Genet et al., 2015). Biological materials typically have a spatially heterogeneous micro- and nano-structure (Gu et al., 2017; Lejeune and Sacks, 2019). In addition, the material properties of biological tissue often vary on macroscopic

length scales (Bersi et al., 2020). Computational modeling is the ideal tool for directly investigating the effects of this heterogeneity in both predictive modeling and material characterization (Hugenberg et al., 2020; Yu et al., 2019). With computational modeling, it is possible to capture spatial heterogeneity and subsequently measure its effects (Gokhale et al., 2008). Critically, making meaningful predictions about the behavior of biological materials will often require multiple simulations to account for both subject-specific variations and model uncertainty (Lee et al., 2018; Tran et al., 2019). In other words, rather than representing a biological material through the outcome of a single simulation, examining the outcomes of a suite of simulations is often

^{*} Corresponding author.

E-mail address: elejeune@bu.edu (E. Lejeune).

<https://doi.org/10.1016/j.jmbbm.2020.104276>

Received 6 August 2020; Received in revised form 28 October 2020; Accepted 13 December 2020

Available online 31 December 2020

1751-6161/© 2020 Published by Elsevier Ltd.

more appropriate (Lejeune and Linder, 2019; Lejeune et al., 2019). However, particularly when simulating complex microstructure and/or whole organ geometries, this approach can be prohibitively computationally expensive (Yang et al., 2010). For example, it is usually prohibitively computationally expensive to simulate every plausible spatial distribution of material properties for a given problem of interest (Bessa et al., 2017; Gu et al., 2018; Liu et al., 2016).

One way to overcome this limitation is through constructing

metamodels (Hanakata et al., 2018). Metamodels, “models of models” often referred to as “surrogate models,” are used to predict model outputs – often referred to as quantities of interest (QoI) – from given model input parameters (Lejeune, 2020a; Queipo et al., 2005). For example, in the work that we present here, we focus on predicting the QoI change in strain energy, $\Delta\psi$, from a bitmap that defines a spatial distribution of material properties. This is illustrated in Fig. 1a. Metamodels are constructed (i.e. trained) from the outcomes of full-fidelity model runs

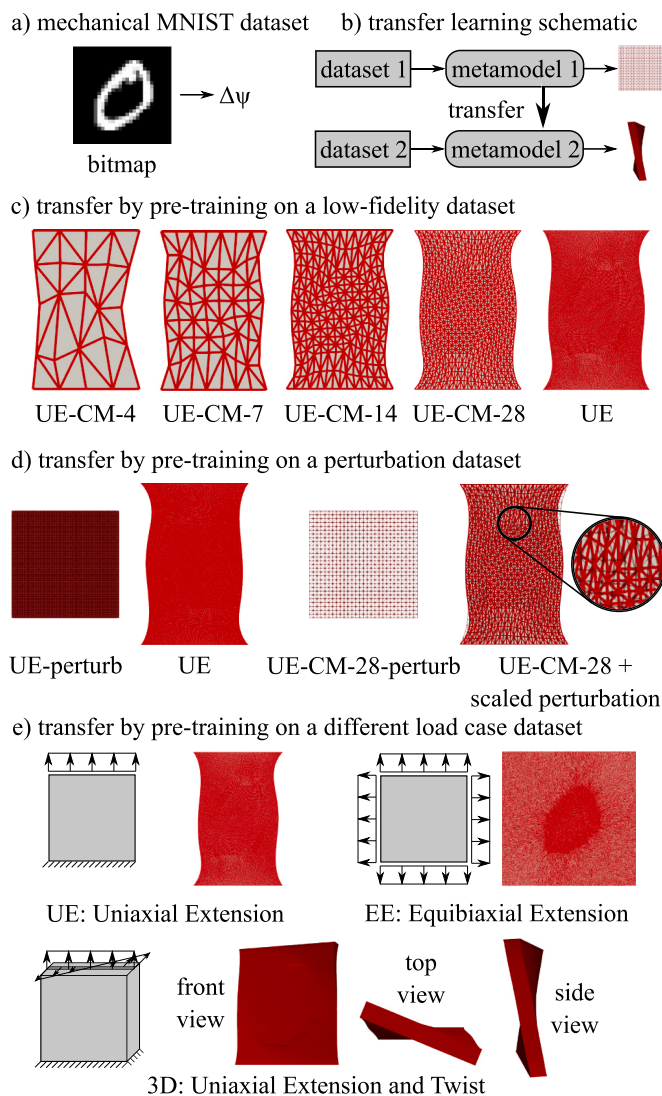


Fig. 1. a) The “Mechanical MNIST” dataset contains FEA simulation results, in this paper we consider the relationship between the input bitmap that prescribes material properties and the change in strain energy due to applied displacement; b) A schematic of a model trained on one dataset being transferred to make predictions on another dataset; c) Illustration of the levels of mesh refinement explored in our low fidelity models; d) An illustration of a perturbation compared to the final level of applied displacement, we note that the scaled result of the perturbation is an imperfect match to the final deformed configuration; e) Schematic illustration of the uniaxial extension (UE), equibiaxial extension (EE), and three-dimensional uniaxial extension and twist (3D) load cases.

(Wang and Sun, 2019). Typically, metamodels are constructed to be computationally cheap to execute (Vu et al., 2017; Yang and Perdikaris, 2019). For example, a trained metamodel will typically execute in fractions of a second while the equivalent finite element simulations typically take orders of magnitude longer to complete (Wang and Sun, 2018). If a metamodel is successfully trained, it will be able to make accurate predictions for input parameter combinations outside of those used to generate training simulations and thus enable a computationally cheap exploration of the input parameter space (Teichert and Garikipati, 2019). This means that metamodels are an invaluable tool for conducting design optimization, sensitivity analysis, uncertainty quantification, and some forms of multiscale modeling (Alber et al., 2019; Bessa et al., 2019; Peirlinck et al., 2019; SahliCostabal et al., 2018).

For metamodels that predict the mechanical behavior of heterogeneous material, convolutional neural networks (CNNs) are a popular choice (Schwarzer et al., 2019; Yang et al., 2019; Zhang and Gu, 2020). This is because CNNs are designed to capture spatial relationships (Yang et al., 2020). Recently, CNNs have been used to capture the mechanical behavior of heterogeneous materials with applications such as fracture prediction (Hsu et al., 2020), design optimization (Yu et al., 2019), and material homogenization (Rao and Liu, 2020). However, CNNs are notorious for requiring a large amount of training data (D'souza et al., 2020). This means that generating sufficient training data will still be prohibitively computationally expensive in many cases. Transfer learning, where knowledge gained while solving one problem is transferred to a different but related problem, may help us overcome this issue (Pan and Yang, 2009). Specifically, if easy to acquire data (either computationally cheap or generated for a previous study) can be leveraged, far fewer harder to acquire or new simulations will be necessary. A schematic of transfer learning is illustrated in Fig. 1b. In the context of simulating biological tissue with the finite element method (Costabal et al., 2019; Javili et al., 2015), examples of easier to acquire simulation data include simulations conducted with linear elements where quadratic elements would be more appropriate, simulations conducted on an un-refined mesh, simulations with less complexity, simulations in a lower dimensional space, and simulations of an identical tissue run previously to solve a different problem. We note that recently multi-fidelity modeling, a specific form of transfer learning where low fidelity simulation data is leveraged to make predictions about high fidelity simulation data, has gained popularity for constructing metamodels of computationally expensive simulations (Bonfiglio et al., 2018). Here we choose a few relevant examples of simulation data from our Mechanical MNIST benchmark dataset (details provided in Section 2.1) to investigate the efficacy of metamodeling via transfer learning for problems related to heterogeneous tissue undergoing large deformation. Examples of the “different but related” datasets

that we investigate in this paper are shown in Fig. 1c–e.

In this paper, we begin in Section 2 by describing our dataset, our metamodeling strategy, and our chosen method for transfer learning. Then, in Section 3, we describe and discuss the results of our study. Critically, we find that transfer learning is an effective way to leverage data from low-fidelity simulations, perturbation simulations, and simulations with different loading conditions. Finally, we conclude in Section 4. In the Additional Information Section, we provide information on how to access all of the data and the code required to reproduce our results.

2. Methods

Here we begin in Section 2.1 with a brief description of the Mechanical MNIST dataset and the additions to the dataset that were required to conduct the research presented in this paper (Lejeune, 2020b). Then, in Section 2.2, we briefly describe our metamodel selection and in Section 2.3 we describe our methods for implementing transfer learning. We end in Section 2.4 with a brief note on alternative approaches to transfer learning and multi-fidelity modeling. Details for accessing all of the relevant data and software are provided in the Additional Information Section.

2.1. Mechanical MNIST dataset

The Mechanical MNIST dataset (Lejeune, 2020b) is a benchmark dataset with data from finite element simulations inspired by the MNIST dataset of handwritten digits popular within the computer vision research community (LeCun et al., 1998). Briefly, the dataset is constructed by using the MNIST bitmaps, illustrated in Fig. 1, to define the material properties of a heterogeneous block of Neo-Hookean material. The white pixels correspond to Young's Modulus $E = 100$, the black pixels correspond to $E = 1$, and E for the grayscale pixels is computed by linearly interpolating. The two-dimensional block of material is then subject to an applied displacement, where boundary conditions for the Uniaxial Extension (UE), Equibiaxial Extension (EE) and three-dimensional extension and twist (3D) load cases are shown schematically in Fig. 1e. The total change in strain energy $\Delta\psi$ is computed for each level of applied displacement. As illustrated in Fig. 1a, we consider a single QoI, $\Delta\psi$, and train our metamodels to predict this QoI when given an input bitmap. Identical to the original MNIST dataset, each subset of Mechanical MNIST contains 60,000 training data points and 10,000 test data points. Further details are presented in the manuscript accompanying the original dataset (Lejeune, 2020b) and in Appendix B. Despite the “toy problem” nature of the dataset, it is specifically focused on heterogeneous materials undergoing large deformation and is thus

Table 1

Summary of the different datasets, all within Mechanical MNIST, that we investigate in this paper. The perturbation displacement is 0.001 units compared to the side length of 28 units. Two additional datasets, UE-CM-7-quad and UE-CM-4-quad, are described in Appendix A. All data is freely available with access details provided in the Additional Information Section.

Name	Description
UE	uniaxial extension, full fidelity dataset (fully refined mesh, quadratic triangular elements, applied displacement is 1/2 of a side length)
EE	equibiaxial extension, full fidelity dataset
3D	uniaxial extension and out of plane twist, full fidelity three dimensional dataset (fully refined mesh, quadratic tetrahedral elements, applied displacement is 1/7 of a side length, twist is $\pi/8$ radians, block thickness is 1/7 of a side length)
UE-CM-28	uniaxial extension, $28 \times 28 \times 2$ linear triangular elements
UE-CM-14	uniaxial extension, $14 \times 14 \times 2$ linear triangular elements
UE-CM-7	uniaxial extension, $7 \times 7 \times 2$ linear triangular elements
UE-CM-4	uniaxial extension, $4 \times 4 \times 2$ linear triangular elements
UE-perturb	uniaxial extension, applied displacement is a perturbation
UE-CM-28-perturb	uniaxial extension, $28 \times 28 \times 2$ linear triangular elements, applied displacement is a perturbation

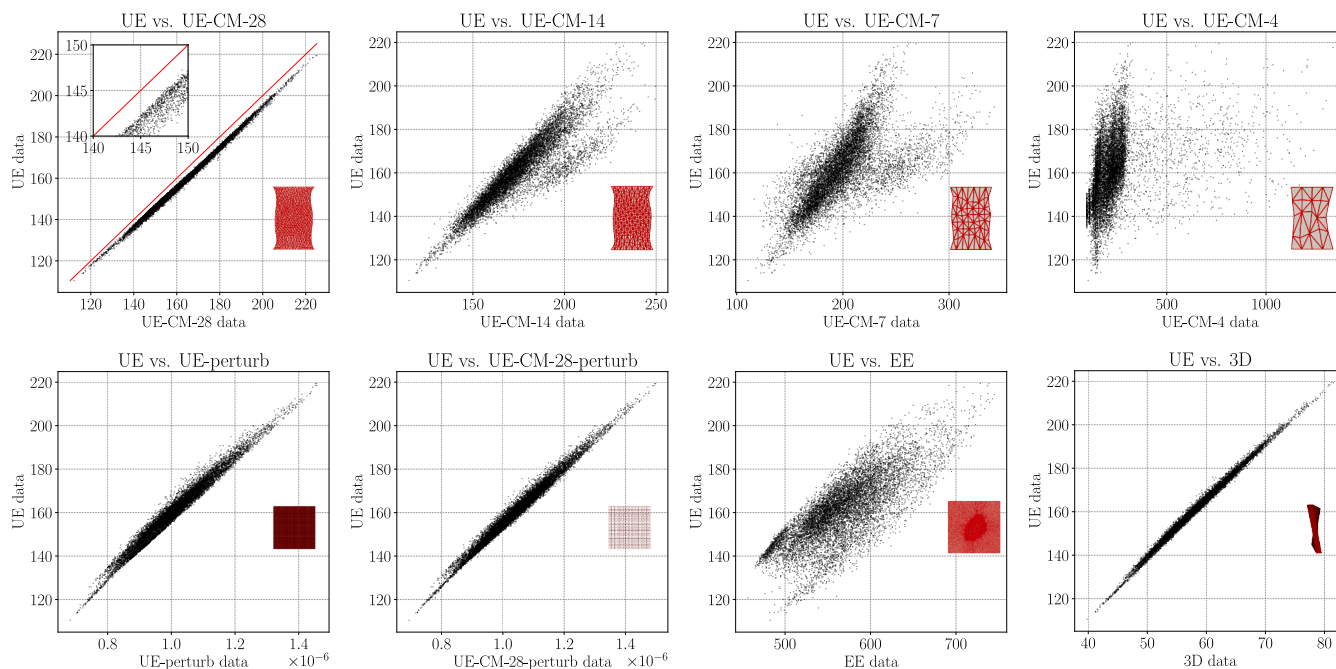


Fig. 2. All plots illustrate the correlation between the standard Uniaxial Extension (UE) dataset and another dataset: UE-CM-28, UE-CM-14, UE-CM-7, UE-CM-4, UE-perturb, UE-CM-28-perturb, EE, and 3D. All axis represent the QoI, change in strain energy $\Delta\psi$ at the end of the simulation, for each dataset. Each point corresponds to an input bitmap in the test set (10,000 points total per plot).

relevant to our broader understanding of the behavior of heterogeneous soft tissue.

In this paper, we use nine different datasets that all fall under the broader Mechanical MNIST umbrella. We provide a summary in [Table 1](#). For this paper specifically, we created the four coarse mesh (“CM”) datasets. Other than changing the mesh size and switching from quadratic to linear elements, these simulations are identical to the simulations in the “UE” high-fidelity dataset. However, because the MNIST input bitmaps are 28×28 , the grids coarser than 28×28 are unable to perfectly resolve the spatial pattern. In addition we generated the three-dimensional dataset “3D” specifically for this paper. All other datasets had already been generated and published [43, 44]. The perturbation datasets “UE-perturb” and “UE-CM-28-perturb” are contained within “UE” and “UE-CM-28” respectively. We also detail two additional datasets “UE-CM-7-quad” and “UE-CM-4-quad” in [Appendix A](#). The open source finite element library FEniCS was used to run all simulations [45, 46].

2.2. Metamodel

For each dataset shown in [Table 1](#), we train a metamodel to predict $\Delta\psi$ from the input bitmap. In all cases, we used the same metamodel architecture. Specifically, we use a convolutional neural network (CNN) ([LeCun et al., 2015](#)) with a convolutional layer with 20 feature maps, followed by a max pooling layer, followed by a convolutional layer with 50 feature maps, followed by a max pooling layer, followed by a fully connected layer with 50 nodes. The output layer is a single node that predicts $\Delta\psi$. We use rectified linear unit (ReLU) activation functions ([Ramachandran et al., 2017](#)). We construct and train the CNNs with open source library PyTorch ([Paszke et al., 2017](#)). Details for accessing the code used to implement our metamodel are given in the Additional Information Section. Metamodel performance is discussed in [Section 3.1](#). We note that our investigation of potential metamodels was not exhaustive. This model architecture is a slight modification of our

previously published baseline model that was chosen based on model architectures that achieves good performance on the original MNIST dataset ([Lejeune, 2020b](#)). To our knowledge, the performance that we achieve in this paper represents the new best performance on the Mechanical MNIST datasets. However, we anticipate that this performance will be improved upon in the future.




Fig. 3. This plot shows the test performance of each metamodel (see [Section 2.2](#)) with respect to the number of training samples for each dataset. We note that the metamodel trained on the UE dataset has 3.96% mean absolute percent error when trained with 100 samples, and 1.00% error when trained with 60,000 samples.

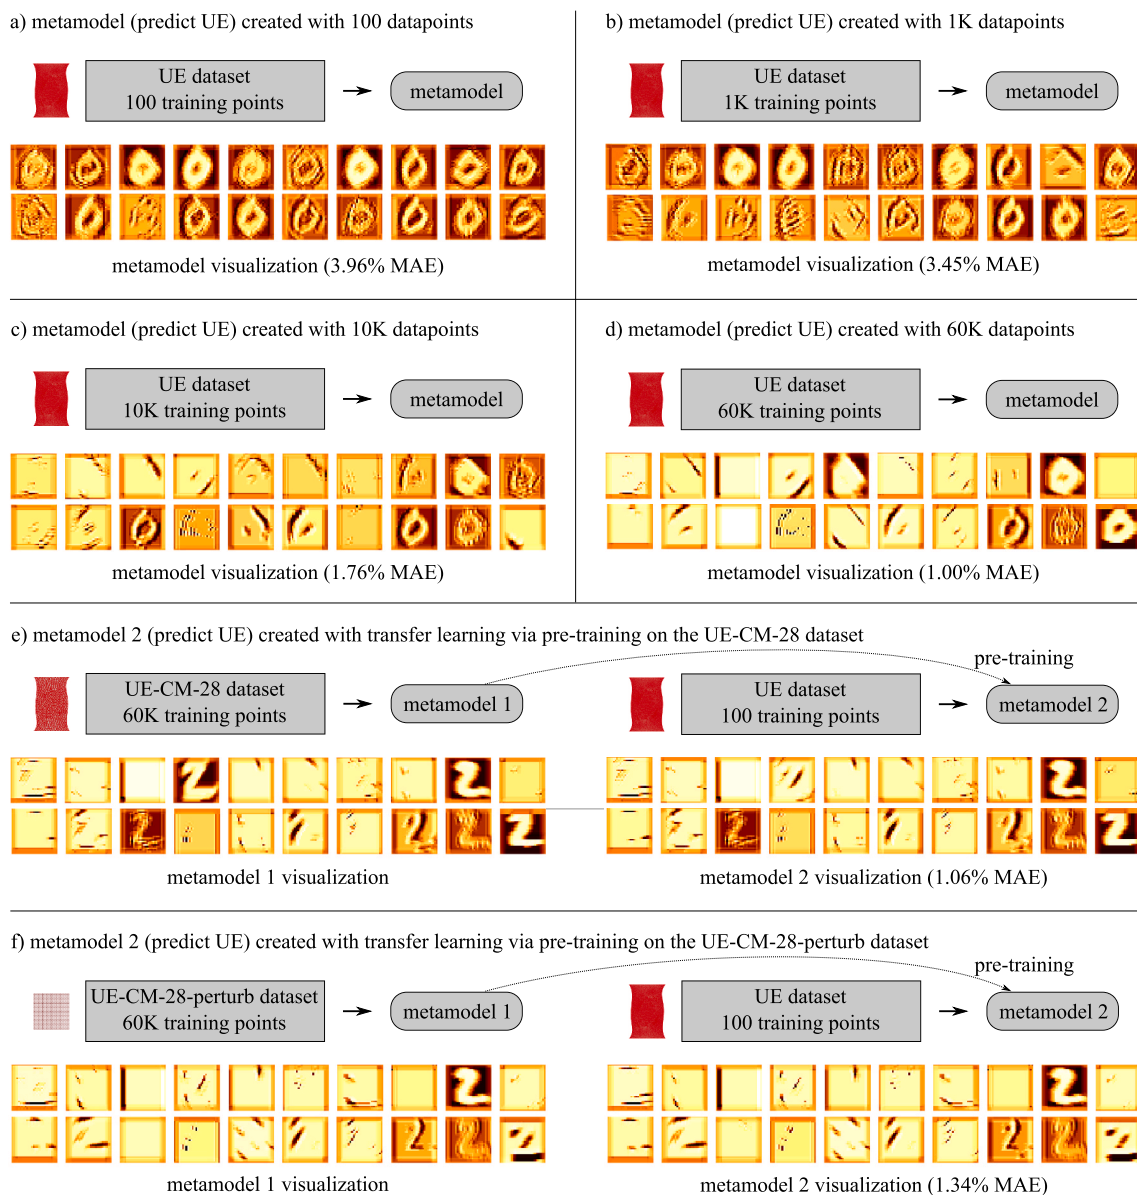


Fig. 4. Illustration of first layer activations for an example input bitmap for different numbers of training points from the UE dataset: a) 100 training points; b) 1,000 training points; c) 10,000 training points; d) 60,000 training points. Illustration of first layer activations for an example input bitmap for pre-trained metamodels: e) a metamodel pre-trained on the UE-CM-28 dataset is then transferred to make predictions on the UE dataset; f) a metamodel pre-trained on the UE CM-28-perturb dataset is then transferred to make predictions on the UE dataset. The “metamodel visualization” shows the first layer activations visualized with guided back-propagation (Springenberg et al., 2014; Ozbulak, 2019). Mean absolute percent error (MAE) is reported.

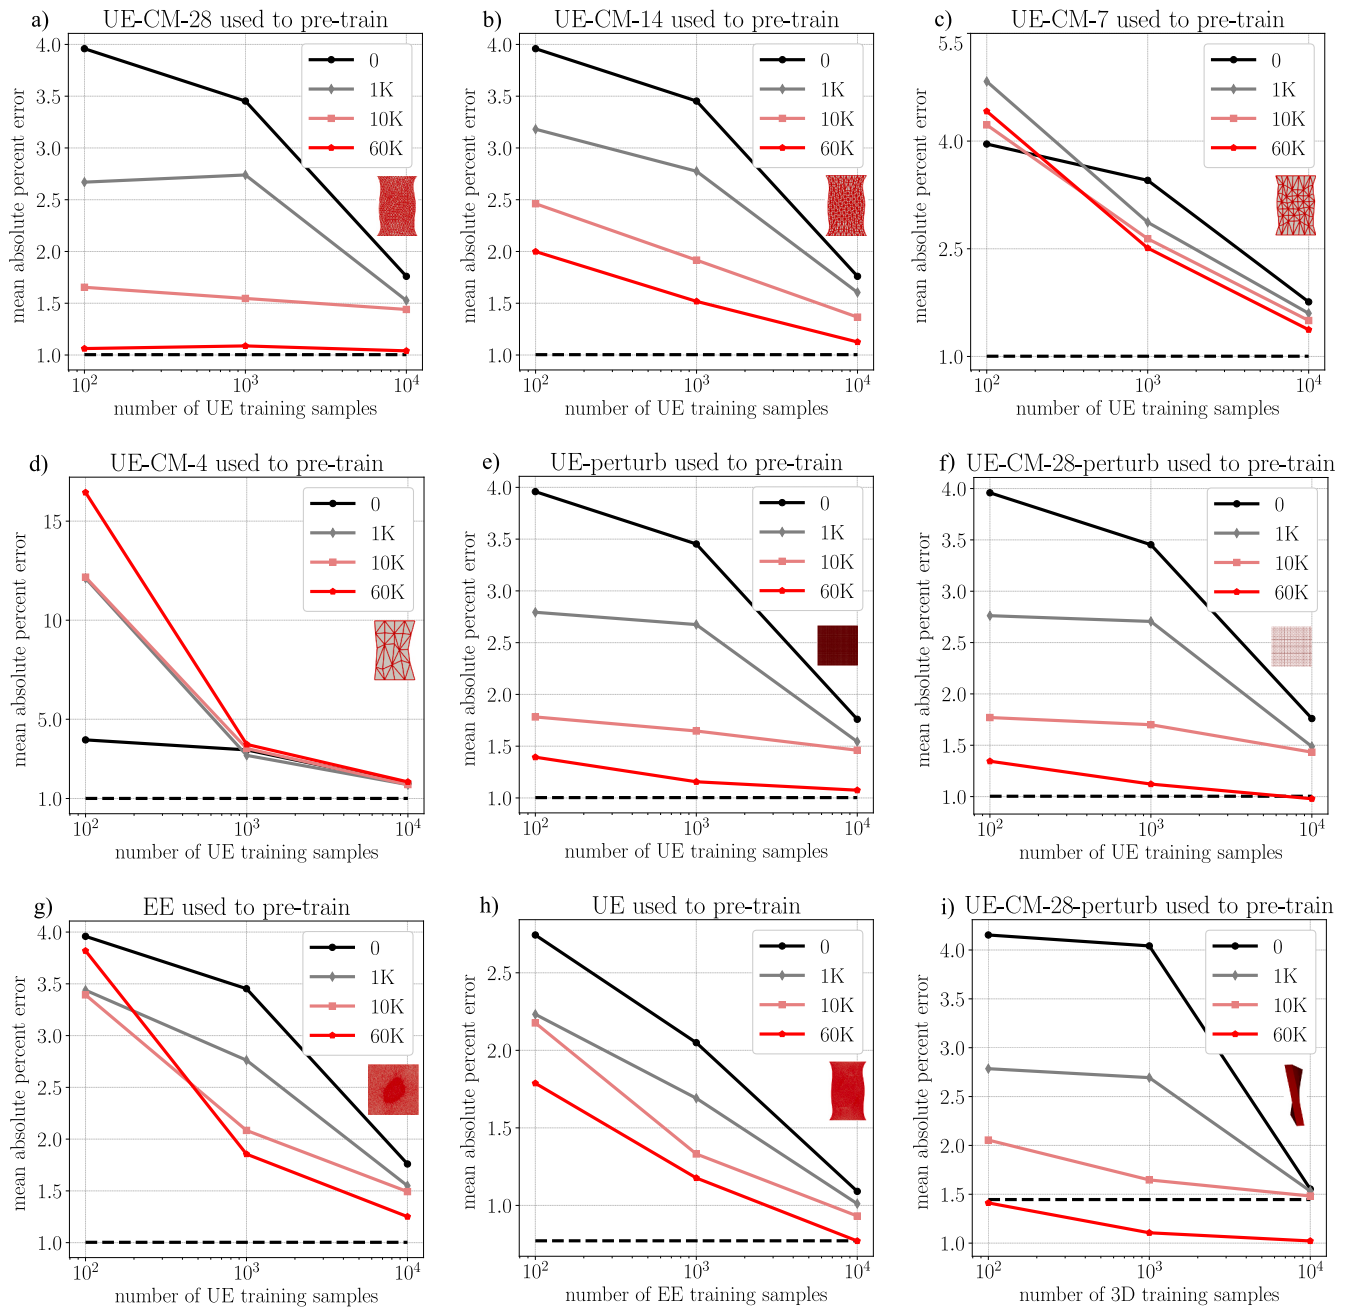


Fig. 5. These plots show the test performance of pre-trained metamodels with respect to the number of training samples. In all plots, the black curve labeled “0” corresponds to metamodel performance without pre-training and the black dashed line corresponds to metamodel performance without pre-training for 60,000 training points which represents the best performance achieved by the metamodel given the available data (see Section 3.1). The other curves, labeled “1K”, “10K”, and “60K”, correspond to pre-trained metamodel performance where each metamodel is pre-trained with the indicated number of points from the related dataset. The plots show: a) UE prediction with UE-CM-28 used for pre-training; b) UE prediction with UE-CM-14 used for pre-training; c) UE prediction with UE-CM-7 used for pre-training; d) UE prediction with UE-CM-4 used for pre-training; e) UE prediction with UE-perturb used for pre-training; f) UE prediction with UE-CM-28-perturb used for pre-training; g) UE prediction with EE used for pre-training; h) EE prediction with UE used for pre-training; i) 3D prediction with UE-CM-28-perturb used for pre-training. Note that the y-axis scaling is not the same for each plot.

2.3. Transfer learning

The essential idea behind transfer learning is that knowledge can be transferred between different but related problems. We note that making predictions based on each dataset listed in Table 1 can be thought of as a different but related problem. The relationship between these datasets is illustrated in Fig. 2, where each datapoint corresponds to one input bitmap in the Mechanical MNIST test set, and each plot shows the relationship between “UE” and another dataset. In all cases, the relationship between the two datasets is clearly not entirely random. With these datasets, we are interested in understanding the potential for transfer learning for three key applications:

1. Leveraging low fidelity data to predict high fidelity results (UE-CM-28, UE-CM-14, UE-CM-7, UE-CM-4 → UE), (UE-CM-28-perturb → 3D)
2. Leveraging perturbation results to predict final simulation results (UE-perturb, UE-CM-28-perturb → UE), (UE-CM-28-perturb → 3D)
3. Leveraging one load case results to predict another load case results (EE → UE), (UE → EE), (UE-CM-28-perturb → 3D)

The exploration of UE-CM-28-perturb → 3D is related to all three applications. As discussed subsequently in Section 2.4, there are multiple different approaches to transfer learning. In this paper, we focus exclusively on transfer learning through model pre-training. With this approach, the model is “pre-trained” on the first dataset, and then further trained on the actual dataset of interest (Yosinski et al., 2014). Typically, the size of the first dataset used for pre-training will be much larger than the actual dataset of interest (Kawano and Yanai, 2014; Yanai and Kawano, 2015). With this method, transfer learning is worthwhile when the metamodel with initialized weights based on pre-training is able to outperform a metamodel with randomly initialized initial weights at a lower total computational cost for generating new simulations. We note that in PyTorch implementing model pre-training is essentially trivial. The only additional steps beyond what is necessary to implement the metamodel described in Section 2.2 are: (1) train an initial metamodel (pre-train) using the same network architecture, and (2) load the initial metamodel weights before training the actual metamodel of interest (Paszke et al., 2017).

2.4. Note on multi-fidelity modeling

In the broader computational mechanics literature, there has been substantial recent interest in “multi-fidelity modeling” where, similar to what we describe in Section 2.3, metamodels are constructed based on simulation data sources of varying fidelity (Bonfiglio et al., 2018). For example, computationally cheap one-dimensional models have been leveraged to make predictions about three-dimensional full-fidelity models (Costabal et al., 2019). The multi-fidelity modeling paradigm has also been used to inform predictive models in the experimental setting with larger amounts of “low-fidelity” information obtained from simulations paired with comparatively few physical experiments (Lu et al., 2020). Multi-fidelity models are not restricted to convolutional neural networks, we note that Gaussian Process Regression is also a popular techniques for this approach (Lee et al., 2020). Conceptually, multi-fidelity models are a form of transfer learning. Notably, many multi-fidelity modeling methods presented in the literature involve specialized model architecture and/or problem specific methods for data integration (Zhang and Garikipati, 2019). In this work, our goal is to explore the efficacy of transfer learning exclusively through straightforward model pre-training. We favor this approach because it is both simple to implement, and, as demonstrated in Section 3, highly effective for the problems that we are interested in. We note that another approach popular in the broader machine learning literature is to use a pre-trained model as a fixed feature extractor and only perform additional training on part of the model (Salberg, 2015). For example, it is

common to hold the weights of all layers but the final layer fixed during the second round of training, or add additional layers to the pre-trained model and only adjust the added layers to the new dataset. In the approach presented in this paper, all model weights are free to update when the new dataset is added. Because we have made our datasets entirely open source, it is possible for other researchers to implement alternative methods and potentially exceed the baseline performance that we report in Section 3.

3. Results and discussion

In this Section, we report the results of training convolutional neural networks to predict change in strain energy $\Delta\psi$ at fixed levels of applied displacement. We begin in Section 3.1 by showing the performance of our metamodel on each of the datasets listed in Table 1. Then, in Section 3.2, we show the efficacy of pre-training on low-fidelity simulation, in Section 3.3 we show the efficacy of pre-training on perturbation results, and in Section 3.4 we show the efficacy of pre-training on a different load case. We end in Section 3.5 with a brief discussion of the limitations of this study and how these limitations could be addressed in future work.

3.1. Metamodel without pre-training

Here we show the performance of the metamodel described in Section 2.2 on the datasets listed in Table 1. In Fig. 3, mean absolute percent error on the test set is plotted with respect to the number of samples used to train the metamodel. With our chosen metamodel architecture and PyTorch implementation (see the Additional Information Section for details for information on accessing the code), the UE dataset metamodel has 3.96% mean absolute percent error when trained with 100 samples and 1.00% error when trained with 60,000 samples. We note that this is an incremental improvement over previously published baseline model performance of 1.9% mean absolute percent error obtained with a CNN using 60,000 training points (Lejeune, 2020b). This metamodel architecture has comparable performance on all other datasets except for UE-CM-4 which, despite improved performance as more training data is added, has 37.5% test error with 100 training points and 11.5% test error with 60,000 training points.

In Fig. 4a–d, we show a visualization of the metamodel as more samples are used for training. Specifically, we show the first layer activations for all 20 feature maps for a representative input bitmap using guided backpropagation (Springenberg et al., 2014). This visualization allows us to qualitatively compare different metamodels as a supplement to our quantitative error comparisons. We note that as the number of training samples increases, the first layer activations change substantially. We will return to this qualitative comparison when we show the results of metamodel pre-training. With pre-training, the first layer activations will look much more like Fig. 4d, despite being trained with the same small number of points as the metamodel shown in Fig. 4a.

3.2. Metamodel pre-trained on a low-fidelity simulation dataset

One of the most appealing applications of transfer learning for metamodel creation is the prospect of leveraging the results of computationally cheap simulations. In our Mechanical MNIST dataset, each of the simulations has quadratic elements and a highly refined mesh and thus takes several minutes to run on a single CPU (see Appendix B for more details). For many applications relevant to biomedical engineering, for example whole organ simulation, it is common for simulations to take tens to thousands of CPU hours to run (Rausch and Humphrey, 2017; Terahara et al., 2020). Changes such as switching from quadratic to linear elements and coarsening the finite element mesh can dramatically reduce computational cost, but these changes will also introduce error. For our UE-CM datasets, simulations take seconds to run on a single CPU, and yield low-fidelity predictions that, as illustrated in

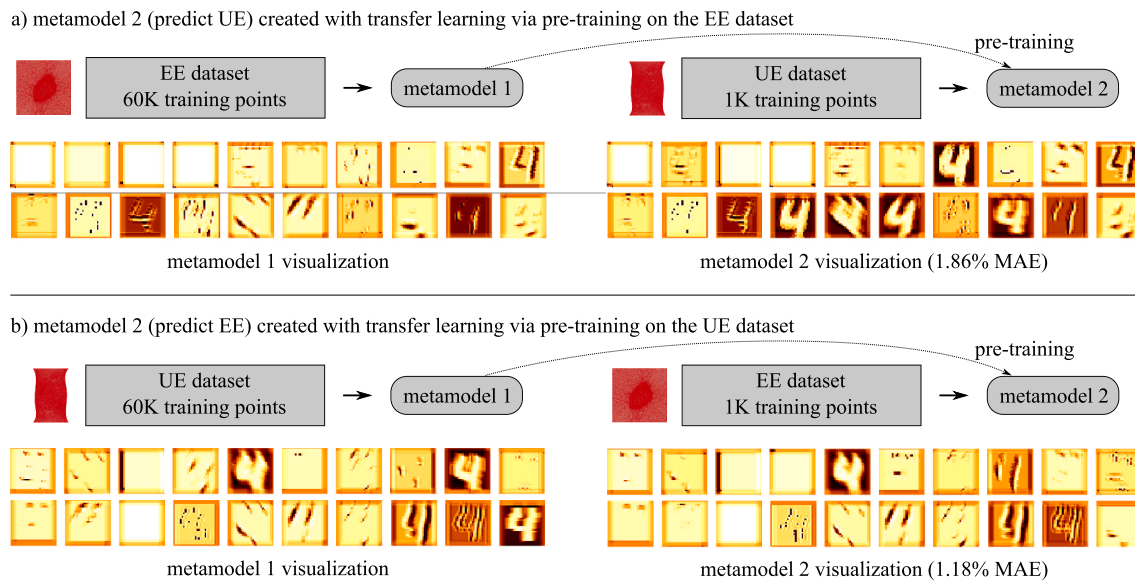


Fig. 6. Illustration of first layer activations for an example input bitmap for pre-trained metamodels: a) a metamodel pre-trained on the EE dataset is then transferred to make predictions on the UE dataset; b) a metamodel pre-trained on the UE dataset is then transferred to make predictions on the EE dataset. The “metamodel visualization” shows the first layer activations visualized with guided backpropagation (Springenberg et al., 2014; Ozbulak, 2019). Mean absolute percent error (MAE) is reported.

Fig. 2, are imperfect but correlated with the high-fidelity simulation results. Here we show what happens to our UE metamodel when it is pre-trained with our low-fidelity datasets.

In Fig. 5a–d, we show metamodel test error with respect to number of training samples for the UE metamodel (black curve) and then the UE metamodel pre-trained with different numbers of training points from the UE-CM datasets. As illustrated in Fig. 2, the UE-CM-28 dataset has the closest relationship to the UE dataset while the relationship degrades as the mesh coarsens. In Fig. 5a, we see that pre-training with UE-CM-28 allows the UE pre-trained metamodel to achieve lower test error with fewer training samples. And, increasing the number of UE-CM-28 datapoints used for training improves the pre-trained metamodel performance. In Fig. 5b we see similar but less dramatic results for the UE-CM-14 dataset. Then, in Fig. 5c, we see an initially worse performance at 100 training points for pre-training with UE-CM-7 followed by a marginal reduction in error for the UE pre-trained metamodels trained with 1,000 and 10,000 points. In Fig. 5d, we see that pre-training with the UE-CM-4 dataset is detrimental to model performance. For UE-CM-4, the combination of such an extremely coarse mesh and linear elements led to erroneous simulation results that could not be leveraged with our current method. In Appendix A, we provide some additional results with a coarse mesh and quadratic elements to separate the two sources of lower data quality.

These results show that with the metamodel architectures described in this paper it is possible to use transfer learning to achieve improved performance by pre-training on low-fidelity data. Pre-training with UE-CM-28 and UE-CM-14 showed clear and consistent benefit, while pre-training with UE-CM-7 and UE-CM-4 did not. All of the datasets are freely available, thus we welcome others to implement methods that surpass this baseline performance and are potentially able to leverage

the UE-CM-7 and UE-CM-4 datasets. In Fig. 4a–d, we visualize the first layer activations for the UE metamodel without pre-training for different numbers of training points, and in Fig. 4e we visualize both the UE-CM-28 60,000 training points metamodel and the subsequent pre-trained UE metamodel with 100 high-fidelity training points. The comparison between Fig. 4a–d and Fig. 4e shows the qualitative similarity between the UE full-fidelity metamodel with a large number of training points and the pre-trained metamodel with a small number of training points. Essentially, the metamodels are qualitatively similar and the pre-training process makes it possible to leverage this similarity.

3.3. Metamodel pre-trained on a perturbation simulation dataset

Building on the theme of leveraging the results of computationally cheap simulations, here we explore the potential benefits of pre-training on the results of simulated perturbations. In this case, this means simulation results from applied displacements orders of magnitude smaller than the actual applied displacement of interest. As illustrated in Fig. 2, both the UE-perturb and the UE-CM-28-perturb datasets are clearly related to the UE dataset. We note that in order to avoid numerical errors during the metamodel training process, we multiply the perturbation $\Delta\psi$ values by 10^8 before training.¹ In Fig. 5e and f, we show metamodel test error with respect to the number of training samples for the UE metamodel (black curve) and then the UE metamodel pre-trained with different numbers of training points from the UE-perturb and the UE-CM-28-perturb datasets. In both cases, pre-training leads to the UE pre-trained metamodel achieving lower test error with fewer training samples. And, increasing the number of datapoints used for pre-training improves the metamodel performance. In Fig. 4f, we also qualitatively

¹ without this step, QoIs are 10^{-6} and, at the time of publication, training on the 10^{-6} values in PyTorch led to a metamodel that predicted 0.0 for all values (Paszke et al., 2017).

compare the metamodel first layer activations for the UE-CM-28-perturb dataset and the UE dataset pre-trained with CM-28-perturb.

3.4. Metamodel pre-trained on an alternative load case dataset

Switching between load cases exemplifies the transfer learning concept of a “different but related problem.” Practically, the ability to leverage data acquired for one load case and apply it to another with minimal additional simulation runs required is a powerful tool for exploring the model parameter space beyond solely the effects of heterogeneous material properties. In Fig. 2, we show that the UE and EE datasets are related, but perhaps less similar than some of the UE and UE-CM datasets. In Fig. 5g, we show metamodel test error for UE metamodels pre-trained on the EE datasets. In Fig. 5h, we show metamodel test error for EE metamodels pre-trained on the UE datasets. And, in Fig. 6, we qualitatively compare the metamodel first layer activations. The plots in Fig. 5 show that pre-training the metamodel on data from another load case can be an effective way to reduce metamodel error. And, using more points for pre-training corresponds to lower metamodel error. The only exception to this is the UE metamodel pre-trained on the EE metamodel trained with 60,000 data points. We anticipate that it would be possible to adjust model architecture or training parameters to overcome what is likely overfitting to the newly provided data, but we chose not to in order to maintain consistency throughout the paper. Overall, we show that metamodel pre-training for transferring between load cases is viable.

In Fig. 5i, we show metamodel test error for 3D metamodels pre-trained on the UE-CM-28-perturb dataset. In Fig. 5j, we see that a metamodel pre-trained with 60,000 UE-CM-28-perturb datapoints and subsequently trained with 100 3D datapoints achieves a slightly lower test error than a metamodel trained with 60,000 3D datapoints. This example ties together the exploration of low fidelity datasets, perturbation datasets, and alternative load case datasets. Notably, the computational cost of generating the simulations for the transfer learning based metamodel is $\approx 0.13\%$ of the computational cost for generating the 60,000 3D simulations, see Table B.3 for additional details. This shows the potential for transfer learning to have an important impact in computational modeling of heterogeneous materials undergoing large deformation.

3.5. Limitations of the present study

In our opinion, there are three main limitations to the present study. First, we are not able to claim that any of our metamodels are the best possible metamodels for a given dataset or that pre-training will *always* improve metamodel performance. It is possible that other researchers would be able to design metamodels that vastly exceed the performance reported here. And, it is possible that other approaches that leverage different but related data would be able to exceed the performance of our model pre-training based approach with either comparable or smaller data requirements. To address this limitation, we have released all of the data that our results are based on and welcome others to create models that exceed the baseline performance shown here (Lejeune, 2020d). This speaks to the broader problem of lack of benchmark data for method comparison within the computational mechanics research community (Lejeune, 2020e; Dolbow et al., 2020).

Second, the Mechanical MNIST dataset that we used does not represent all possible qualities of heterogeneous soft tissue. For example, the current form of the dataset lacks a version with an anisotropic material model, a version with coupled fields, and a version where the

spatially heterogeneous distribution of material properties reflects specific types of biological tissue. To address this limitation, we plan to add to the dataset in the future and we welcome others to either build on our open source software (see the Additional Information Section) to generate simulations or create benchmark datasets with their own software that best address domain specific problems of interest. We note that our proposed straightforward transfer learning approach is likely successful because, as shown in Fig. 2, the QoIs in each datasets are correlated. Specifically, the features identified by the CNNs are relevant across datasets. We note that transfer learning is an active area of research (Pan and Yang, 2009; Tan et al., 2018; Torrey and Shavlik, 2010), and that future exploration will involve: (1) developing protocols for generating the most effective low fidelity datasets and (2) benchmarking alternative transfer learning approaches applied to mechanical data.

Finally, our study is limited by the relative simplicity of our datasets compared to the complexity of potential applications of this method such as whole organ simulations with heterogeneous material properties. Our investigation suggests that other researchers will be able to use this method to leverage data from whole organ simulations performed on an un-refined mesh to help make predictions about the outcomes of whole organ simulations performed on a fully refined mesh. However, we have yet to directly demonstrate this claim. And, to our knowledge, this has not been directly demonstrated in the literature to date. Future work will directly evaluate the efficacy of this approach for creating metamodels of whole organ simulation.

4. Conclusion

In this paper, we demonstrate that transfer learning via metamodel pre-training is a powerful tool for creating metamodels of heterogeneous material behavior. We demonstrate the efficacy of the approach by example using our open source Mechanical MNIST benchmark dataset. First, we show that convolutional neural networks are an effective tool for creating metamodels based on our datasets. Then, we demonstrate that pre-training on data from simulations conducted with a coarse mesh, pre-training on data from simulated perturbations, and pre-training on alternative load cases can all improve metamodel performance. In the most dramatic example of the efficacy of transfer learning, we show that pre-training with the UE-CM-28 (coarse mesh) and the UE-CM-28-perturb (coarse mesh perturbation simulation) datasets makes metamodels trained with 100 high-fidelity simulation results perform almost as well as metamodels trained with 60,000 high-fidelity simulation results for both two-dimensional and three-dimensional high fidelity simulations ($\approx 1\%$ error and $\approx 1.5\%$ error respectively). In addition to this exploration of transfer learning, we have also contributed both low-fidelity and three-dimensional simulation results to the Mechanical MNIST benchmark dataset. This contribution will enable others to directly compare competing methods that may exceed the baseline performance shown here. Dataset and software access details are provided in the Additional Information Section. Looking forward, we anticipate that other researchers interested in creating metamodels of heterogeneous materials will be able to directly implement the methods presented in this paper. These methods will make it computationally tractable to explore the true behavior of spatially heterogeneous soft tissue and thus better understand and predict the true mechanical behavior of biological materials. Overall, we anticipate that transfer learning will enable substantial future advances in soft tissue simulation for both basic research and clinical applications.

Author statement

Emma Lejeune: Conceptualization, Methodology, Software, Data Curation, Writing, Supervision. Bill Zhao: Software, Writing.

Additional information

The entire Mechanical MNIST collection is available through the OpenBU Institutional Repository: <https://open.bu.edu/handle/2144/39371> (Lejeune, 2019, 2020c). The data used specifically for this paper is available at: <https://hdl.handle.net/2144/41357> (Lejeune, 2020d). All code used to generate the data and metamodels presented in this paper is available through GitHub: <https://github.com/elejeune11/Mechanical-MNIST-Transfer-Learning>.

Declaration of competing interest

The authors declare that they have no known competing financial

interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We would like to thank the staff of the Boston University Research Computing Services and the OpenBU Institutional Repository (in particular Eleni Castro) for their invaluable assistance with generating and disseminating the Mechanical MNIST datasets. This work was made possible through the Engineering Research Centers Program of the National Science Foundation (No. EEC-1647837), the Boston University Hariri Institute Junior Faculty Fellows program, and start up funds from the Boston University Department of Mechanical Engineering. We gratefully acknowledge this support.

Appendix A. Additional datasets

Table A.2

Summary of two additional datasets, both within Mechanical MNIST. Both datasets are freely available with access details provided in the Additional Information Section.

Name	Description
UE-CM-7-quad	uniaxial extension, $7 \times 7 \times 2$ quadratic triangular elements
UE-CM-4-quad	uniaxial extension, $4 \times 4 \times 2$ quadratic triangular elements

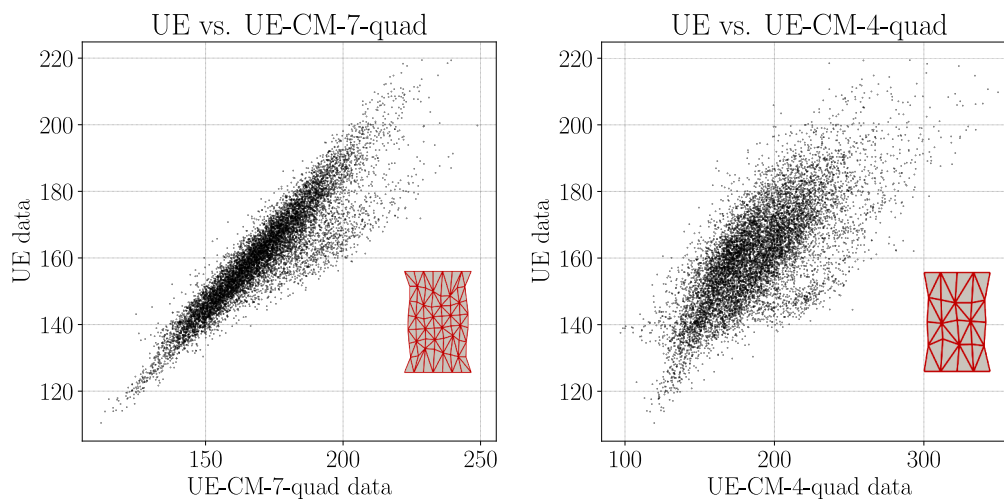


Fig. A.7. Both plots illustrate the correlation between the standard Uniaxial Extension (UE) dataset and another dataset: UE-CM-7-quad and UE-CM-4-quad. All axis represent the QoI, change in strain energy $\Delta\psi$ at the end of the simulation, for each dataset. Each point corresponds to an input bitmap in the test set (10,000 points total per plot).

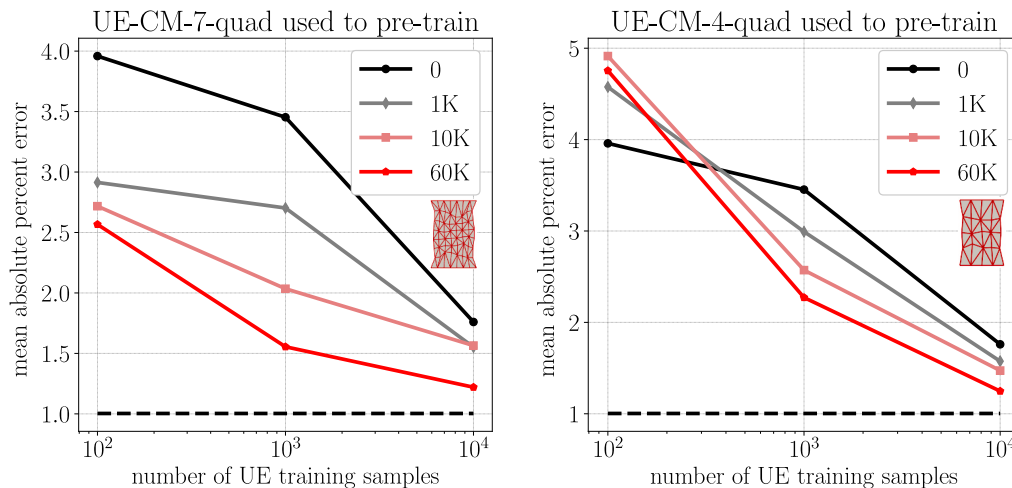


Fig. A.8. These plots show the test performance of pre-trained metamodels with respect to the number of training samples. In all plots, the black curve labeled “0” corresponds to metamodel performance without pre-training and the black dashed line corresponds to metamodel performance without pre-training for 60,000 training points which represents the best performance achieved by the metamodel given the available data (see Section 3.1). The other curves, labeled “1K”, “10K”, and “60K”, correspond to pre-trained metamodel performance where each metamodel is pre-trained with the indicated number of points from the related dataset. Left: UE prediction with UE-CM-7-quad used for pre-training; Right: UE prediction with UE-CM-4-quad used for pre-training.

In Table A.2, we describe two additional datasets: “UE-CM-7-quad” and “UE-CM-4-quad.” These datasets are identical to UE-CM-7 and UE-CM-4 respectively but with quadratic elements rather than linear elements. Details for accessing these datasets are provided in the Additional Information Section. In Fig. A.7, we illustrate the relationship between each dataset and the UE dataset. These plots are consistent with Fig. 2 in the main body of the text. Then, in Fig. A.8, we plot the results of predicting UE from metamodels pre-trained on each dataset. These plots are consistent with Fig. 5 in the main body of the text.

In Table B.3, we include these datasets in our evaluation of *approximate* computational cost. When making choices about the type of simulations to run for pre-training, the key tradeoff is between the computational cost of generating the finite element simulation dataset, and the potential increase in metamodel performance. We note that, for this example, a finer discretization of linear elements provides a better tradeoff compared to a coarser discretization of quadratic elements. However, computational cost may vary with different choices of hardware and software, which may influence the balance of this tradeoff. Ultimately, the optimal parameters of the low fidelity model will be problem specific.

Appendix B. Additional information on dataset generation

Essential information about dataset generation is provided in Section 2.1 of this paper and in our previously published work (Lejeune, 2020b). For reference, we use a compressible Neo-Hookean material model:

$$\psi = \frac{1}{2}\mu[\mathbf{F} : \mathbf{F} - 3 - 2\ln(\det\mathbf{F})] + \frac{1}{2}\lambda\left[\frac{1}{2}((\det\mathbf{F})^2 - 1) - \ln(\det\mathbf{F})\right] \quad (\text{B.1})$$

where ψ is the strain energy, \mathbf{F} is the deformation gradient, and μ and λ are the Lamé parameters equivalent to Poisson’s ratio ν and Young’s modulus E :

$$E = \frac{\mu(3\lambda + 2\mu)}{\lambda + \mu} \quad \nu = \frac{\lambda}{2(\lambda + \mu)}. \quad (\text{B.2})$$

We convert the MNIST bitmap images to material properties by dividing the material domain such that it corresponds with the grayscale bitmap and then we specify E as

$$E = \frac{c}{255.0} (100.0 - 1.0) + 1.0 \quad (\text{B.3})$$

where c is the corresponding value of the grayscale bitmap that can range from 0 – 255. For all simulation, we keep Poisson’s ratio fixed at $\nu = 0.3$ throughout the domain. Each simulation is thus a soft background material with “digits” that are two orders of magnitude stiffer than the background material. In Fig. B.9, we show change in strain energy $\Delta\psi$ as a function of applied displacement for five randomly selected curves from each dataset. The same five input bitmaps are used to generate the curves in each plot. Consistent with the plots shown in Fig. 2 and Fig. A.7, many of the datasets are both qualitatively and quantitatively similar. We note that all code used to generate the dataset is freely available, with access details given in the Additional Information Section.

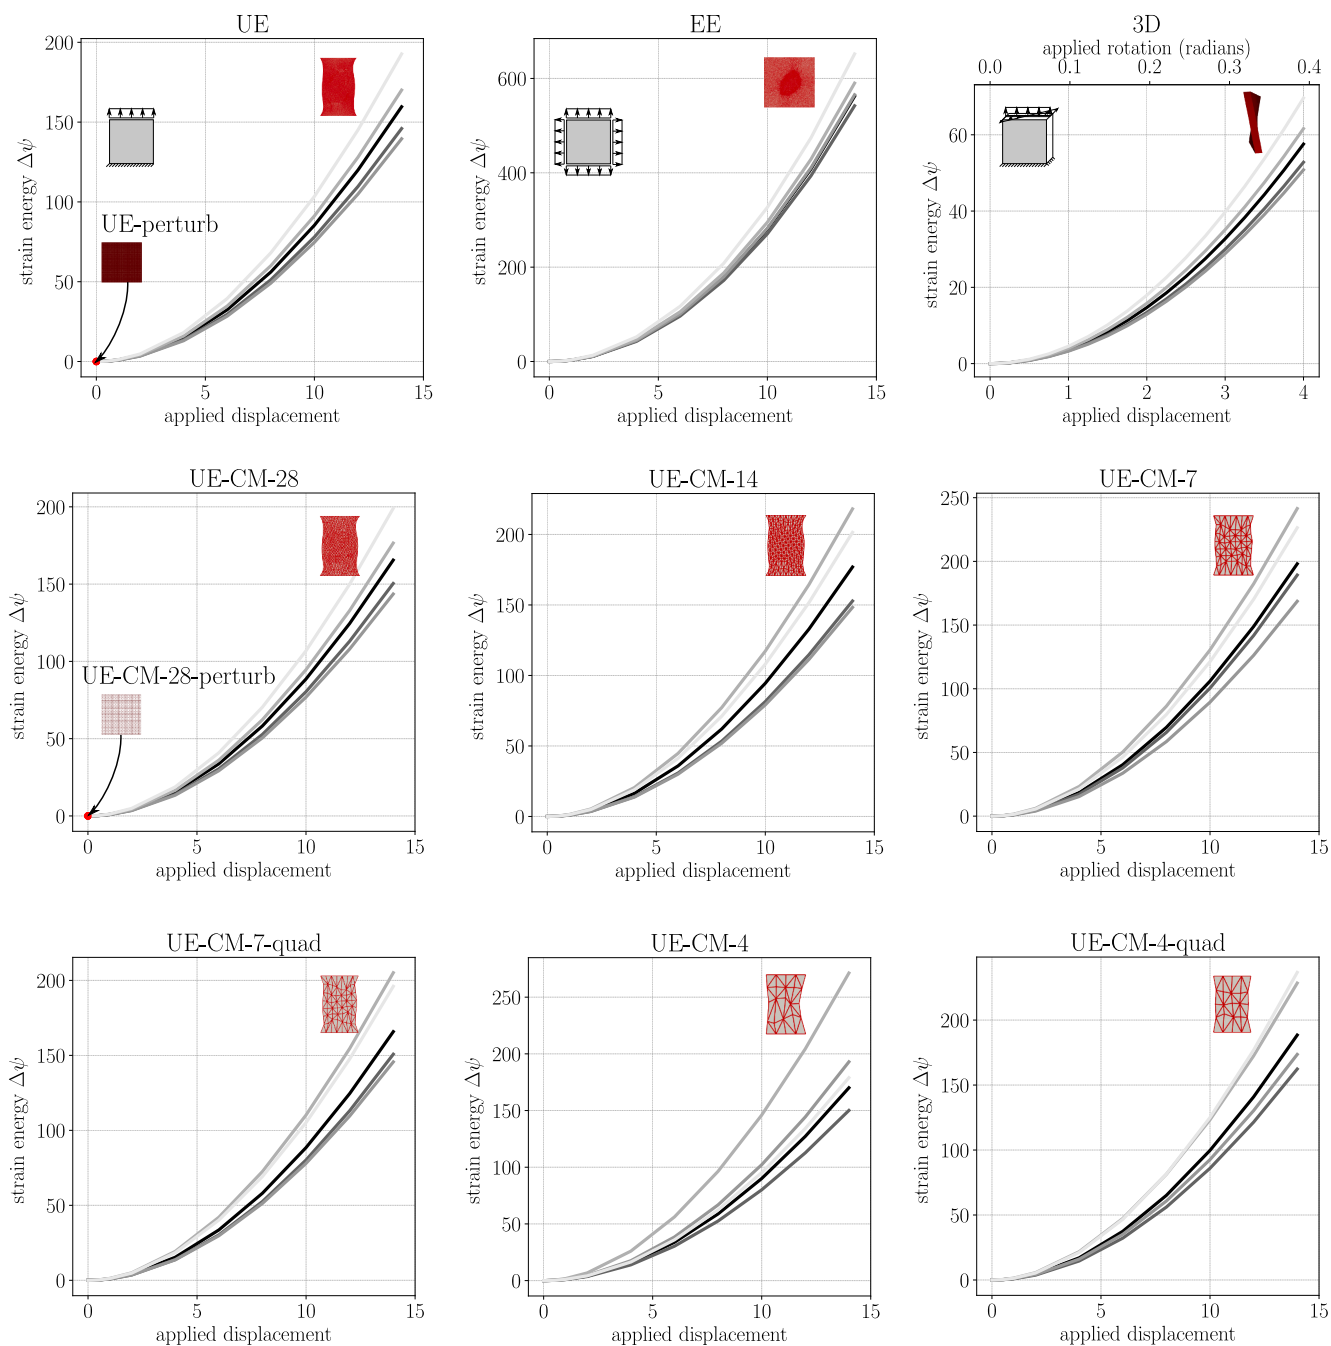


Fig. B.9. Five randomly selected change in strain energy $\Delta\psi$ vs. applied displacement curves from each dataset. Applied displacement units match the units of the 28×28 unit block. The same five input bitmaps are used for each plot. For the “perturb” datasets, the datapoint that corresponds to the perturbation simulation is shown in red. For all other datasets, $\Delta\psi$ is taken at the final step of applied displacement.

Table B.3

Summary of the *approximate* computational cost per simulation. Each simulation is run with open source software FEniCS (Alnaes et al., 2015; Logg et al., 2012) at the Massachusetts Green High Performance Computing Center on a single core. All code to generate the data is freely available with access details provided in the Additional Information Section. We note that with different hardware or software choices computational time may vary significantly.

Name	<i>Approximate</i> time for one simulation in seconds	<i>Approximate</i> time for 60,000 simulations in hours
3D	1,300 s	30,000 h
UE	400 s	6,600 h
UE-CM-28	1.8 s	30 h
UE-CM-14	1.4 s	23 h
UE-CM-7	1.4 s	23 h
UE-CM-4	1.3 s	22 h
UE-CM-7-quad	2.8 s	46 h
UE-CM-4-quad	1.6 s	26 h
UE-perturb	28 s	460 h
UE-CM-28-perturb	0.24 s	4.1 h

In Table B.3, we list the *approximate* computational cost of generating each dataset. With our chosen software and hardware, time for a single simulation ranges from 0.24 s (UE-CM-28-perturb) to 1,300 s (3D). In Fig. 5, we compare the performance of metamodels that are and are not pre-trained and argue that pre-training can lead to equivalent metamodel performance with a substantially cheaper to generate dataset. For example, in Fig. 5a, a pre-trained metamodel with 60,000 UE-CM-28 simulations and 100 UE simulations performs similarly to a metamodel trained with 60,000 UE data points. The pre-training case requires approximately 41 CPU hours whereas the case without pre-training requires approximately 6,600 CPU hours. For the example shown in Fig. 5f, the pre-training case (60,000 UE-CM-28-perturb and 100 UE) requires approximately 15 CPU hours whereas the case without pre-training requires approximately 6,600 CPU hours. From all examples shown in Fig. 5i, the most extreme example is UE-CM-28-perturb used to pre-train 3D. In this example, the pre-training case (60,000 UE-CM-28-perturb and 100 3D) requires approximately 40 CPU hours, whereas the case without pre-training requires approximately 30,000 CPU hours. Therefore, in our most extreme example, the metamodel built with transfer learning requires 0.13% of the CPU hours of the metamodel built without it. This clearly demonstrates the potential of transfer learning for enhancing metamodels of heterogeneous material behavior.

Table B.4

This table accompanies Fig. 5 and Fig. A.8. Here we compare the metamodel error on the dataset of interest (dataset 2) before and after using 100 dataset-2 samples for additional training. *As stated in the text, $\Delta\psi$ in the perturbation dataset is multiplied by 10^8 before metamodel training to avoid numerical errors during the training process.

pre-trained on dataset-1 \rightarrow additional training on dataset-2	number of pre-training points	mean absolute test error on dataset-2 before transfer	mean absolute test error on dataset-2 after transfer with 100 points additional training
UE-CM-28 \rightarrow UE	1, 000	6.08%	2.67%
	10, 000	3.34%	1.65%
	60, 000	3.16%	1.06%
UE-CM-14 \rightarrow UE	1, 000	14.1%	3.18%
	10, 000	7.98%	2.46%
	60, 000	10.3%	2.00%
UE-CM-7 \rightarrow UE	1, 000	18.4%	4.83%
	10, 000	22.6%	4.23%
	60, 000	24.1%	4.42%
UE-CM-4 \rightarrow UE	1, 000	47.9%	12.1%
	10, 000	47.7%	12.2%
	60, 000	55.1%	16.4%
UE-perturb* \rightarrow UE	1, 000	37.1%	2.79%
	10, 000	37.4%	1.78%
	60, 000	38.3%	1.39%
UE-CM-28-perturb* \rightarrow UE	1, 000	32.9%	2.76%
	10, 000	34.9%	1.77%
	60, 000	34.6%	1.34%
UE-CM-7-quad \rightarrow UE	1, 000	10.9%	2.91%
	10, 000	8.32%	2.72%
	60, 000	9.99%	2.57%
UE-CM-4-quad \rightarrow UE	1, 000	16.6%	4.58%
	10, 000	19.7%	4.91%
	60, 000	18.1%	4.75%
EE \rightarrow UE	1, 000	258%	3.44%
	10, 000	252%	3.39%
	60, 000	252%	3.82%
UE \rightarrow EE	1, 000	71.0%	2.23%
	10, 000	71.4%	2.18%
	60, 000	71.7%	1.79%
UE-CM-28-perturb* \rightarrow 3D	1, 000	85.5%	2.79%
	10, 000	80.2%	2.05%
	60, 000	81.1%	1.41%

Appendix C. Metamodel performance before and after transfer

In Fig. 5 and Fig. A.8, we plot metamodel performance with respect to the number of “dataset-2” training samples (see schematic in Fig. 1). In each plot, each curve corresponds to a different number of “dataset-1” samples used for metamodel pre-training. In Fig. 5 and Fig. A.8, there is a clear comparison between metamodels with pre-training (curves labeled 1K, 10K, and 60K) and metamodels with no pre-training (black curve labeled 0). We note that in most cases investigated, pre-training reduces metamodel error. An additional useful comparison is the metamodel performance without additional training on dataset-2 (i.e. transfer). If dataset-1 and dataset-2 are identical, the additional training will make little difference. In Table B.4, we compare mean absolute test error before transfer (column 3) and mean absolute test error after transfer with 100 points of additional training (column 4). We note that the values in column 4 are datapoints in the Fig. 5 and Fig. A.8 plots. In all cases, the additional training decreases error when making predictions on dataset-2. In many cases, the metamodel without additional training would do a poor job making predictions on dataset-2.

References

- Alber, M., Tepole, A.B., Cannon, W.R., De, S., Dura-Bernal, S., Garikipati, K., Karniadakis, G., Lytton, W.W., Perdikaris, P., Petzold, L., et al., 2019. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *npj Dig. Med.* 2 (1), 1–11.
- Alnas, M., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M.E., Wells, G.N., 2015. The fenics project version 1.5. *Arch. Numer. Software* 3 (100).
- Ambrosi, D., Ateshian, G.A., Arruda, E.M., Cowin, S., Dumais, J., Goriely, A., Holzapfel, G.A., Humphrey, J.D., Kemkemer, R., Kuhl, E., et al., 2011. Perspectives on biological growth and remodeling. *J. Mech. Phys. Solid.* 59 (4), 863–883.
- Bersi, M.R., Santamaría, V.A.A., Marback, K., Di Achille, P., Phillips, E.H., Goergen, C.J., Humphrey, J.D., Avril, S., 2020. Multimodality imaging-based characterization of regional material properties in a murine model of aortic dissection. *Sci. Rep.* 10 (1), 1–23.
- Bessa, M., Bostanabad, R., Liu, Z., Hu, A., Apley, D.W., Brinson, C., Chen, W., Liu, W.K., 2017. A framework for data-driven analysis of materials under uncertainty: countering the curse of dimensionality. *Comput. Methods Appl. Mech. Eng.* 320, 633–667.
- Bessa, M.A., Glowacki, P., Houlder, M., 2019. Bayesian machine learning in metamaterial design: fragile becomes supercompressible. *Adv. Mater.* 31 (48), 1904845.
- Bonfiglio, L., Perdikaris, P., Brizzolara, S., Karniadakis, G., 2018. Multi-fidelity optimization of super-cavitating hydrofoils. *Comput. Methods Appl. Mech. Eng.* 332, 63–85.
- Costabal, F.S., Perdikaris, P., Kuhl, E., Hurtado, D.E., 2019. Multi-fidelity classification using Gaussian processes: accelerating the prediction of large-scale computational models. *Comput. Methods Appl. Mech. Eng.* 357, 112602.
- Dolbow, J., Garikipati, K., Oskay, C., Park, H., 2020. The computational mechanics vision workshop. <https://micde.umich.edu/wp-content/uploads/sites/6/2020/07/NSF-Workshop-Report-Final-1.pdf>.
- D’souza, R.N., Huang, P.-Y., Yeh, F.-C., 2020. Structural analysis and optimization of convolutional neural networks with a small sample size. *Sci. Rep.* 10 (1), 1–13.
- Genet, M., Rausch, M., Lee, L.C., Choy, S., Zhao, X., Kassab, G.S., Kozerke, S., Guccione, J.M., Kuhl, E., 2015. Heterogeneous growth-induced prestrain in the heart. *J. Biomech.* 48 (10), 2080–2089.
- Gokhale, N.H., Barbone, P.E., Oberai, A.A., 2008. Solution of the nonlinear elasticity imaging inverse problem: the compressible case. *Inverse Probl.* 24 (4), 045010.
- Gu, G.X., Libonati, F., Wettermark, S.D., Buehler, M.J., 2017. Printing nature: unraveling the role of nacre’s mineral bridges. *J. Mech. Behav. Biomed. Mater.* 76, 135–144.
- Gu, G.X., Chen, C.-T., Buehler, M.J., 2018. De novo composite design based on machine learning algorithm. *Extreme Mech. Lett.* 18, 19–28.
- Hanakata, P.Z., Cubuk, E.D., Campbell, D.K., Park, H.S., 2018. Accelerated search and design of stretchable graphene kirigami using machine learning. *Phys. Rev. Lett.* 121 (25), 255304.
- Hsu, Y.-C., Yu, C.-H., Buehler, M.J., 2020. Using deep learning to predict fracture patterns in crystalline solids. *Matter* 3 (1), 197–211. <https://www.sciencedirect.com/science/article/pii/S2590238520301909>.
- Hugenberg, N.R., Dong, L., Cooper, J.A., Corr, D.T., Oberai, A.A., 2020. Characterization of spatially graded biomechanical scaffolds. *J. Biomech. Eng.* 142 (7).
- Javili, A., Dortdivanlioglu, B., Kuhl, E., Linder, C., 2015. Computational aspects of growth-induced instabilities through eigenvalue analysis. *Comput. Mech.* 56 (3), 405–420.
- Kawano, Y., Yanai, K., 2014. Automatic expansion of a food image dataset leveraging existing categories with domain adaptation. In: *European Conference on Computer Vision*. Springer, pp. 3–17.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al., 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86 (11), 2278–2324.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521 (7553), 436–444.
- Lee, T., Turin, S.Y., Gosain, A.K., Bilionis, I., Tepole, A.B., 2018. Propagation of material behavior uncertainty in a nonlinear finite element model of reconstructive surgery. *Biomech. Model. Mechanobiol.* 17 (6), 1857–1873.
- Lee, T., Bilionis, I., Tepole, A.B., 2020. Propagation of uncertainty in the mechanical and biological response of growing tissues using multi-fidelity Gaussian process regression. *Comput. Methods Appl. Mech. Eng.* 359, 112724.
- Lejeune, E., 2019. Mechanical MNIST – uniaxial extension. <https://open.bu.edu/handle/2144/38693>.
- Lejeune, E., 2020a. Geometric stability classification: datasets, metamodels, and adversarial attacks. *Comput. Aided Des.* 131, 102948.
- Lejeune, E., 2020b. Mechanical MNIST: a benchmark dataset for mechanical metamodels. *Extreme Mech. Lett.* 100659.
- Lejeune, E., 2020c. Mechanical MNIST – equibiaxial extension. <https://open.bu.edu/handle/2144/39428>.
- Lejeune, E., 2020d. Mechanical MNIST – multi-fidelity. <https://open.bu.edu/handle/2144/41357>.
- Lejeune, E., 2020e. Geometric stability classification: datasets, metamodels, and adversarial attacks. *Comput. Aided Des.* 102948.
- Lejeune, E., Linder, C., 2019. Interpreting stochastic agent-based models of cell death. *Comput. Methods Appl. Mech. Eng.* 112700.
- Lejeune, E., Sacks, M.S., 2019. Analyzing valve interstitial cell mechanics and geometry with spatial statistics. *J. Biomech.* 93, 159–166.
- Lejeune, E., Javili, A., Weickenmeier, J., Kuhl, E., Linder, C., 2016. Tri-layer wrinkling as a mechanism for anchoring center initiation in the developing cerebellum. *Soft Matter* 12 (25), 5613–5620.
- Lejeune, E., Dortdivanlioglu, B., Kuhl, E., Linder, C., 2019. Understanding the mechanical link between oriented cell division and cerebellar morphogenesis. *Soft Matter*.
- Liu, Z., Bessa, M., Liu, W.K., 2016. Self-consistent clustering analysis: an efficient multi-scale scheme for inelastic heterogeneous materials. *Comput. Methods Appl. Mech. Eng.* 306, 319–341.
- Logg, A., Mardal, K.-A., Wells, G., 2012. Automated Solution of Differential Equations by the Finite Element Method: the FEniCS Book, vol. 84. Springer Science & Business Media.
- Lu, L., Dao, M., Kumar, P., Ramamurty, U., Karniadakis, G.E., Suresh, S., 2020. Extraction of mechanical properties of materials through deep learning from instrumented indentation. *Proc. Natl. Acad. Sci. Unit. States Am.* 117 (13), 7052–7062.
- Ozbulak, U., 2019. Pytorch Cnn Visualizations. <https://github.com/utkuozbulak/pytorch-cnn-visualizations>.
- Pan, S.J., Yang, Q., 2009. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* 22 (10), 1345–1359.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A., 2017. Automatic Differentiation in Pytorch.
- Peirlinck, M., Costabal, F.S., Sack, K., Choy, J., Kassab, G., Guccione, J., De Beule, M., Segers, P., Kuhl, E., 2019. Using machine learning to characterize heart failure across the scales. *Biomech. Model. Mechanobiol.* 18 (6), 1987–2001.
- Queipo, N.V., Haftka, R.T., Shyy, W., Goel, T., Vaidyanathan, R., Tucker, P.K., 2005. Surrogate-based analysis and optimization. *Prog. Aero. Sci.* 41 (1), 1–28.
- Ramachandran, P., Zoph, B., Le, Q.V., 2017. Searching for Activation Functions arXiv preprint arXiv:1710.05941.
- Rao, C., Liu, Y., 2020. Three-dimensional convolutional neural network (3d-cnn) for heterogeneous material homogenization. *Comput. Mater. Sci.* 184, 109850.
- Rausch, M.K., Humphrey, J.D., 2017. A computational model of the biomechanics of an evolving occlusive thrombus. *J. Elasticity* 129 (1–2), 125–144.
- Sahli Costabal, F., Yao, J., Kuhl, E., 2018. Predicting drug-induced arrhythmias by multiscale modeling. *Int. J. Numer. Methods Biomed. Eng.* 34 (5), e2964.
- Salberg, A.-B., 2015. Detection of seals in remote sensing images using features extracted from deep convolutional neural networks. In: *2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE, pp. 1893–1896.
- Schwarzer, M., Rogan, B., Ruan, Y., Song, Z., Lee, D.Y., Percus, A.G., Chau, V.T., Moore, B.A., Rougier, E., Viswanathan, H.S., et al., 2019. Learning to fail: predicting fracture evolution in brittle material models using recurrent graph convolutional neural networks. *Comput. Mater. Sci.* 162, 322–332.
- Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M., 2014. Striving for Simplicity: the All Convolutional Net arXiv preprint arXiv:1412.6806.
- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., Liu, C., 2018. A survey on deep transfer learning. In: *International Conference on Artificial Neural Networks*. Springer, pp. 270–279.
- Teichert, G.H., Garikipati, K., 2019. Machine learning materials physics: surrogate optimization and multi-fidelity algorithms predict precipitate morphology in an alternative to phase field dynamics. *Comput. Methods Appl. Mech. Eng.* 344, 666–693.

- Terahara, T., Takizawa, K., Tezduyar, T.E., Bazilevs, Y., Hsu, M.-C., 2020. Heart valve isogeometric sequentially-coupled fsi analysis with the space–time topology change method. *Comput. Mech.* 1–21.
- Torrey, L., Shavlik, J., 2010. Transfer learning. In: *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*. IGI global, pp. 242–264.
- Tran, J.S., Schiavazzi, D.E., Kahn, A.M., Marsden, A.L., 2019. Uncertainty quantification of simulated biomechanical stimuli in coronary artery bypass grafts. *Comput. Methods Appl. Mech. Eng.* 345, 402–428.
- Vu, K.K., d'Ambrosio, C., Hamadi, Y., Liberti, L., 2017. Surrogate-based methods for black-box optimization. *Int. Trans. Oper. Res.* 24 (3), 393–424.
- Wang, K., Sun, W., 2018. A multiscale multi-permeability poroplasticity model linked by recursive homogenizations and deep learning. *Comput. Methods Appl. Mech. Eng.* 334, 337–380.
- Wang, K., Sun, W., 2019. Meta-modeling game for deriving theory-consistent, microstructure-based traction–separation laws via deep reinforcement learning. *Comput. Methods Appl. Mech. Eng.* 346, 216–241.
- Yanai, K., Kawano, Y., 2015. Food image recognition using deep convolutional network with pre-training and fine-tuning. In: *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, pp. 1–6.
- Yang, Y., Perdikaris, P., 2019. Adversarial uncertainty quantification in physics-informed neural networks. *J. Comput. Phys.* 394, 136–152.
- Yang, W., Feinstein, J.A., Marsden, A.L., 2010. Constrained optimization of an idealized y-shaped baffle for the fontan surgery at rest and exercise. *Comput. Methods Appl. Mech. Eng.* 199 (33–36), 2135–2149.
- Yang, C., Kim, Y., Ryu, S., Gu, G.X., 2019. Using convolutional neural networks to predict composite properties beyond the elastic limit. *MRS Communications* 9 (2), 609–617.
- Yang, C., Kim, Y., Ryu, S., Gu, G.X., 2020. Prediction of composite microstructure stress-strain curves using convolutional neural networks. *Mater. Des.* 189, 108509.
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H., 2014. How transferable are features in deep neural networks?. In: *Advances in Neural Information Processing Systems*, pp. 3320–3328.
- Yu, C.-H., Qin, Z., Buehler, M.J., 2019. Artificial intelligence design algorithm for nanocomposites optimized for shear crack resistance. *Nano Futur.* 3 (3), 035001.
- Zhang, X., Garikipati, K., 2019. Machine Learning Materials Physics: Multi-Resolution Neural Networks Learn the Free Energy and Nonlinear Elastic Response of Evolving Microstructures arXiv preprint arXiv:2001.01575.
- Zhang, Z., Gu, G.X., 2020. Finite-element-based Deep-Learning Model for Deformation Behavior of Digital Materials, *Advanced Theory and Simulations*, p. 2000031.