

RESEARCH ARTICLE

Double U-Net: Improved multiscale modeling via fully convolutional neural networks

Julian Lißner  | Felix Fritzen 

Data Analytics in Engineering, Institute of Applied Mechanics, University of Stuttgart, Stuttgart, Germany

Correspondence

Julian Lißner, Data Analytics in Engineering, Institute of Applied Mechanics, University of Stuttgart, 70569 Stuttgart, Germany.
Email: lissner@mib.uni-stuttgart.de

Funding information

Deutsche Forschungsgemeinschaft, Grant/Award Numbers: EXC2075-390740016, DFG-FR2702/8-406068690, DFG-FR2702/10-517847245

Abstract

In multiscale modeling, the response of the macroscopic material is computed by considering the behavior of the microscale at each material point. To keep the computational overhead low when simulating such high performance materials, an efficient, but also very accurate prediction of the microscopic behavior is of utmost importance. Artificial neural networks are well known for their fast and efficient evaluation. We deploy fully convolutional neural networks, with one advantage being that, compared to neural networks directly predicting the homogenized response, any quantity of interest can be recovered from the solution, for example, peak stresses relevant for material failure. We propose a novel model layout, which outperforms state-of-the-art models with fewer model parameters. This is achieved through a staggered optimization scheme ensuring an accurate low-frequency prediction. The prediction is further improved by superimposing an efficient to evaluate U-net, which captures the remaining high-level features.

1 | INTRODUCTION

When considering the size effect [1] for larger structures, a safety factor of up to 2 has to be adopted in order to ensure the material of a component will not break. This safety factor arises due to effects occurring on the microscale of the material, that is, due to accumulated damage on the microstructured material. In multiscale modeling, the material response on the macroscopic scale is obtained by considering the effects occurring on the microscale. Thus, enabling the engineering of high-performance materials suitable to effectively and efficiently fulfill the requirements of the deployed structure. In a multiscale simulation, at every material point on the macroscale, the microscopic behavior of the material is considered and modeled [2, 3], leading to prohibitively high computational cost. To circumvent the costly and repetitive simulation of the microscale, machine learning models are deployed to obtain an efficient, yet accurate prediction of the microscopic material behavior [4]. Artificial neural networks are a popular choice to predict the quantity of interest [5, 6], often being the homogenized response of the microstructured material, either obtained via a feature transform with a subsequent machine learning model [4, 5], or via the direct prediction using convolutional neural networks (Conv Net) [6]. The microstructured material is often represented as image data, for example, obtained via a CT-scan, making it suitable for Conv Nets to directly operate on the image input and efficiently yield an accurate prediction.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Proceedings in Applied Mathematics & Mechanics* published by Wiley-VCH GmbH.

A subtype in the field of Conv Nets, that is, fully convolutional neural networks, yield another image as their prediction, which can be used to predict the full field solution in the context of microstructure modeling [7, 8]. An additional advantage of the predicted full field solution is that any quantity of interest can be extracted a posteriori, such that one is not constrained to the homogenized response, but one could also recover, for example, the peak stresses, which are relevant for material failure. The proposition of the U-net by Ronneberger et al. [9] surged the interest in the image to image prediction, which has an encoder–decoder structure, and first compresses the spatial resolution of the input image, before increasing it again to recover the prediction at the original resolution. Since the fully Conv Net operates on multiple resolutions, various approaches have emerged to utilize an approximation of the solution in a lower resolution, which contributes to the loss/model optimization and can be trivially obtained through coarse graining during training [10, 11]. Another approach simulates the material at the microscale in coarse resolution and uses the solution as input to the model while upsampling [7]. A different layout has also been considered, which incorporates information of the image on multiple spatial resolution levels [11, 12], starting at a very coarse resolution, on which the convolutional layers can consider large features within the image, and increasingly refining the prediction on each level when recovering the original resolution for its prediction.

This article diverts its attention to fully Conv Nets. We utilize the U-net structure and coarse grained solutions during the training for optimization. The model is designed to have comparatively few parameters to optimize, while being significantly more accurate than recent state-of-the-art models. This is achieved through a model layout, which ensures very accurate predictions on lower resolutions. To further refine the prediction of a model, a second efficient to evaluate U-Net is superimposed to capture high-level features. Due to the nature of the multiple interdependent prediction contributions, a staggered training scheme is proposed. The code of the model's implementation is made freely available in Lissner [13].

2 | MULTISCALE MODELING

In multiscale modeling, the behavior of each point at the macroscopic scale \bar{x} is governed by its underlying microscale with the domain Ω . For arbitrarily different length scales, that is, $L_{\text{micro}} \ll L_{\text{macro}}$, a separation of length scale can be safely assumed. Further, it is assumed that the micro- and the macroscale follow the same material laws, which is described in more detail in Leuschner [14]. Thus, we will focus on the microscale and consider a microstructural unit cell Ω to motivate the materials behavior. For the sake of brevity, consider the equilibrium condition for the steady state heat equation

$$\operatorname{div}(\mathbf{q}(\mathbf{x})) = 0 \quad (\mathbf{x} \in \Omega), \quad (1)$$

with the heat flux \mathbf{q} . The constitutive equation relating the heat flux \mathbf{q} to the temperature gradient $\mathbf{g} = \nabla\theta$ is given as

$$\mathbf{q}(\mathbf{x}) = -\boldsymbol{\kappa}(\mathbf{x}) \mathbf{g}(\mathbf{x}) \quad (\mathbf{x} \in \Omega), \quad (2)$$

when following Fourier's law. Here the heterogeneous thermal conductivity tensor $\boldsymbol{\kappa}$ is introduced. The investigated microstructured material is subsided to periodic and antiperiodic boundary conditions with the fluctuating terms $\tilde{\cdot}$, being expressed as

$$\tilde{\theta}_+ = \tilde{\theta}_- \quad \tilde{\mathbf{q}} \cdot \mathbf{n}|_+ = -\tilde{\mathbf{q}} \cdot \mathbf{n}|_-, \quad (3)$$

where each quantity \cdot_{\pm} belongs to the periodic point set of the boundary $\partial\Omega_{\pm}$, respectively. The loading on the microscale is prescribed by its macroscopic counterpart. Once the materials response on the microscale is calculated, the homogenized quantity of interest $\langle \cdot \rangle$ is recovered by the homogenization operation

$$\langle \cdot \rangle = \frac{1}{|\Omega|} \int_{\Omega} \cdot \, d\Omega. \quad (4)$$

The Hill–Mandel condition is expressed as

$$\bar{\mathbf{q}} = \langle \mathbf{q} \rangle \quad \bar{\mathbf{g}} = \langle \mathbf{g} \rangle \quad \bar{\mathbf{q}} \cdot \bar{\mathbf{g}} = \langle \mathbf{q} \cdot \mathbf{g} \rangle, \quad (5)$$

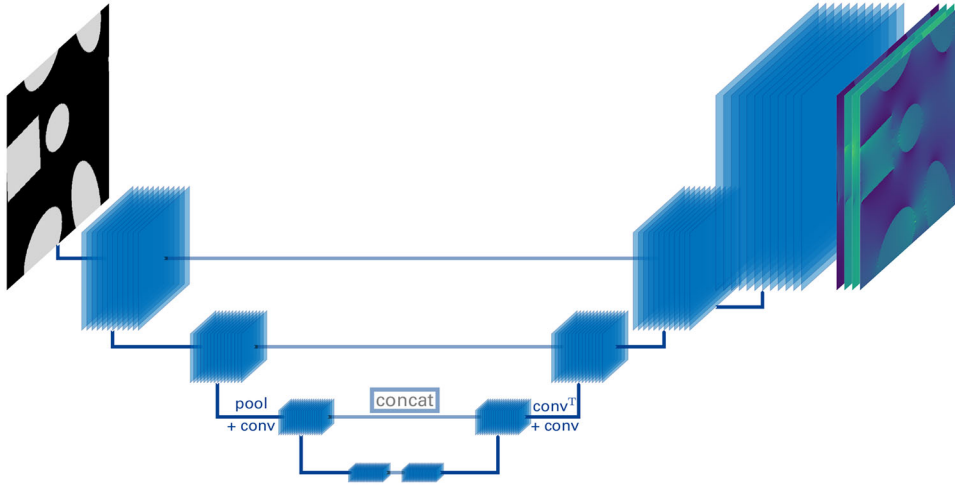


FIGURE 1 A schematic illustration of the U-Net is shown. The operations between the connections are typically modularized on the down/up path, such that each connecting arrow reflects the same operations. Each blue rectangle represents one channel at current spatial resolution.

which completes the scale transition from micro- to macroscale. Considering the material at the microscale for each material point of the macroscopic body leads to a prohibitive computational cost. Machine learning surrogates are deployed to efficiently predict the microstructural behavior. There are numerous studies, which directly predict the homogenized quantity via machine learning [4–6], and others which predict the full field solution [11, 15], and recover the quantities of interest thereof, for example, via Equation (4). This study falls into the latter subcategory, where the full-field solution of the heat flux \mathbf{q} is predicted.

3 | CONVOLUTIONAL NEURAL NETWORKS

3.1 | Introduction

Convolutional neural networks (Conv Nets) are originally inspired by dense feed forward neural networks and developed to handle high-dimensional input data. This is realized by replacing the dense connection between two layers by a convolution operation, such that the forward propagation reads as

$$A_{l+1} = f_{l+1}(k_{l+1} * A_l + b_{l+1}) \quad (6)$$

where generally square kernels k of size 3×3 are deployed to globally operate on the previous layer's output A_l . After the addition of the bias term b , an activation function f is deployed to introduce nonlinearity between layers. To improve information processing between two layers, *channels* are introduced (represented as multiple rectangles in Figure 1), which are comparable to the number of hidden neurons in dense neural networks. Here, a brief motivation for convolutional neural networks is given to familiarize the reader with the relevant notion for the effective design of Conv Nets. Basic algorithmic operations are nicely illustrated in Dumoulin and Visin [16].

The forward pass via a convolution operation implies that each data point in the input image is processed by the same machine learned weights k . The underlying feature transform between layers is a global convolution operation, affected by the spatially local neighborhood. Therefore, we introduce the notion of the receptive field, that is, the number of adjacent pixels, which are at most considered by a single convolution operation. Since each 3×3 convolution increases the receptive field for all subsequent layers by two pixels, a prohibitive number of convolution operations would be required to achieve large receptive fields. When reducing the spatial resolution, which is implemented by evaluating the kernel each *stride* increment, the receptive field is virtually increased. After applying an operation with a stride of two, every subsequent operation has virtually doubled the receptive field with respect to the original spatial resolution of the input image. After resolution reduction, the number of feature channels can be increased while keeping memory usage constant. In terms of




FIGURE 2 The VVE-net is illustrated in a simplified manner. The two different model contributions in dark and light blue are evaluated independently and the final prediction is given via summation. Each arrow represents one module of convolution operations given in Table 1. The predicting module on each coarse grained level is graphically illustrated to the right.

fully convolutional neural networks, the spatial resolution of the Conv Nets output has to match the resolution of the input image. To increase the resolution, either upsampling, or *transpose convolutions*[16] are used. In this end to end setting, Ronneberger et al. [9] proposed the *U-net* (Figure 1), which reuses the feature channels from the *left* while decreasing the resolution, and concatenates them to later layers on the *right* at the respective spatial resolution. This improves information processing in the model and provides low-level information to latter layers of the model.

Recent research has progressed in the direction of improved information processing [17–19], where *shortcuts* are used between two modules/layers. For the design of Conv Nets, modular blocks are generally used, which replace a single convolution operation between two layers by multiple operations. In Szegedy et al. [17], multiple parallel convolutions of differently sized kernels are deployed as *inception modules*. The *ResNet* [18] deploys blocks of two 3×3 convolutions, augmented by an identity mapping, which adds the previous features to the module’s output, improving the *gradient flow* [19] and enabling deeper models to reliably converge.

In the field of fully convolutional neural networks, additional improvements have been found by infusing some knowledge of the approximate solution on lower resolutions into the model while the U-net structure increases the resolution. A popular option to infuse the knowledge is to assist the neural network during training by fitting the solution on multiple resolutions [11] to the coarse grained (average pooled) solution. Another interesting approach has been proposed by Zhou et al. [7], evaluating the simulation on a very coarse resolution, serving as a feature input on the respective spatial resolution.

3.2 | Model proposition

The general structure of the U-net [9] is adopted, and the data usage of the model is optimized to reduce the computational overhead of the model while retaining/increasing accuracy. A main idea of the model is to have two prediction contributions, where the **blue part**, the \surd -net, aims to predict a good *low-frequency* field, while the **lightblue part**, the \surd -net, is superimposed to capture the remainder of the high-level/frequency features.

The structure and information flow of the proposed layout is illustrated in Figure 2 (left), where each arrow indicates one module. The multiple feature types, which contribute to the prediction in the proposed U-net structure, are denoted as

- I—the input image, also coarse grained to each spatial resolution.
- II—features obtained after decreasing the spatial resolution in the *down path*.

TABLE 1 Each repeated module in the \checkmark -net is shown. The “conv” keyword before each layer has been omitted for readability. Each operation is read as “(conv) kernel size/stride.” “Max” refers to Max-pooling. Operations displayed in parallel are evaluated in parallel as an inception module.

| \checkmark -net | | | \checkmark -net | | |
|-----------------------------|------------------------------------|-----------------------------|---------------------------|------------------------------------|------------------------|
| Down path: II ^{VE} | Up path III ^{VE} | Predictor ^{VE} | Down path II ^V | Up path III ^V | Predictor ^V |
| 3/2-5/2-Max 3/2 | 2/2 ^T -4/2 ^T | 2·3/1-1/1·2n _{out} | 3/2-5/2-Max/2 | Add ← II ^V | 3·3/1 |
| Concat | Concat | Add | Concat | 2/2 ^T -4/2 ^T | 1/1 |
| 1/1 | 1/1 | 1/1·n _{out} | 1/1 | Concat | |
| | Concat ← g(x, L _{i-1}) | | 3·3/1 | 1/1 | |
| | | | | 3·3/1 | |

- III—features while increasing the spatial resolution in the *up path* (III* + the upsampled prediction of the previous level).
- IV—the prediction at current spatial resolution.

The model reuses modular blocks of convolutions on each operation in the same *path*, shown in Table 1, where the more intricate predictor on each spatial resolution of the \checkmark -net is shown in Figure 2 (right). Note that the coarse grained prediction (IV) is not directly added to the next levels prediction, but combined with the feature channels used for upsampling and followed by a few convolution operations. This improved the models prediction quality significantly, since a Conv Net struggles to remove the introduced upsampling artifacts with a single convolution, or by adding a *correction* term after each upsampling operation. Such upsampling artifacts can be spotted by a critical eye in Marcato et al. [11].

The final prediction of the model is given as

$$f(x) = g(x) + \Delta^h(x). \quad (7)$$

Since the two contributions do interact with each other, a multistage training scheme is proposed. First, the \checkmark -net is optimized, and frozen thereafter. As a next step, the \checkmark -net is activated and trained to refine the prediction. In a final step, both predictors of the \checkmark -net are optimized in a combined manner.

When optimizing the \checkmark -net, the prediction $g(x)$ on each level $L_* = \{1, 2, \dots, n_{\text{levels}}\}$ contributes to the total loss $\tilde{\Phi}$, that is,

$$\tilde{\Phi}(g(x), y, L_*) = \sum_{i \in L_*} c_i \cdot \Phi(g(x, i), \text{pool}(y, i)), \quad (8)$$

where multiple contributions of the loss Φ are weighted with predefined constants c_i . The solution y is coarse-grained with the average pooling operation to match the current levels spatial resolution. In general, the smaller the spatial resolution, the more channels are used. Using the base number of channels of $s_0 = 15$, the number of channels per level are scaled as

$$s_i = \left(1 + \frac{i}{3}\right) \cdot s_0. \quad (9)$$

In addition to the training split of the \checkmark -net and \checkmark -net, the \checkmark -net is optimized in a staggered manner, while increasing the levels and the number of active modules, during a *pretraining* loop. One pretraining loop consists of a first optimization loop where the constant c_j is zero for all but the current levels, that is,

$$c_j = 0 \quad \forall j \neq i \quad i, j \in L_* \quad (10)$$

optimized in an incremental manner for each level i in L_* . Since the unconstrained optimization only ensures that the current level is a good approximation of the solution, in a second step, the approximation on each level is ensured with

$$0 \leq c_j \leq c_i \quad \forall i \in L_i \quad \text{with} \quad j < i, \quad (11)$$




FIGURE 3 The predicted distribution of the heat flux q_{xx} is compared using the models best predicting \checkmark -net and the U-ResNet for one random sample of the test set (left). The red bars denote the overestimation of the model in the current bin, and the orange bars denote the under estimation, respectively. On the right, the relative cumulative bin errors are shown, for the plotted sample (on the left) in dashed lines, and averaged over the entire test set in full lines, where the colors associate to each model shown on the left.

such that the current level is weighted as the most important. After one complete pretraining loop, that is, both incremental optimization schemes (10), (11), the \checkmark -net is frozen except for the final predicting level, and the remainder of the \checkmark -net is optimized as described above.

4 | RESULTS

The proposed model is applied for the full-field prediction of the microstructured material characterized by representative volume elements (RVE), which is subsided to periodic boundary conditions (3). Thus, a single frame, that is, the input image to the fully convolutional neural network (Conv Net), suffices to fully characterize the material at the microscale. We consider the stationary heat equation outlined in Section 2 in a linear setting under constant loading conditions. Through superpositioning, the effective heat conduction tensor $\bar{\kappa}$ from Equation (2) can be recovered and generalized to arbitrary loading conditions.

The proposed model is compared to the *U-ResNet* used in Santos et al. [15]. It has also been compared to the more recent models [11, 12], which have significantly underperformed the presented models. We deploy a data augmentation scheme during runtime outlined in Lißner and Fritzen [20], which utilizes the periodic boundary conditions to artificially generate new samples. In total, 1500 samples have been used for training, where the validation set contained 300 samples thereof. For optimization, all models have been trained using the mean squared error (MSE) as loss and the adam optimizer, with constant initial learning rate of 10^{-3} and early stopping. Each model is set out to predict all four components of the heat flux \mathbf{q} , that is, x and y response under x and y loading, respectively. The RVE image is the input to the model of fixed resolution of 128×128 voxels, where different input features following [15] did not improve the prediction. All error measures below refer to a test set consisting of 500 samples.

The effectiveness of the proposed optimization scheme of Section 3.2 has been investigated in an ablation study, where the results are summarized in Table 2. The table outlines statistical moments of the resulting heat flux distribution q_{xx} , as well as global MSE metrics. There it can be observed that the \checkmark -net significantly outperforms the UResNet when utilizing the optimization scheme. Adding the \checkmark -net on top of the \checkmark -net significantly improves the prediction, especially the errors on the peak values of the heat flux, which supports the design purpose of the \checkmark -net, delivering a good field prediction.

In the subsequent figures, the U-ResNet, which has been reimplemented from Santos et al. [15], is compared to the previously best performing model, that is, the \checkmark -net pretrained multiple times. In Figure 3, it is quantified how well the machine learned VVE model recovers the actual distribution of the heat flux. As can be seen, the proposed \checkmark -net misclassifies only $\approx 6\%$ of the pixels, whereas the comparable U-ResNet has more than $\approx 12\%$ of pixels wrongly predicted over the entire test set. The \checkmark -net captures the underlying physics of the problem, delivers a smooth field prediction as well as accurate peak value predictions.

TABLE 2 Different global error measures are given of the heat flux distribution for the differently trained models. Every error measure but the mean squared error (MSE) is given as a relative error measure as percentages [%]. Error measures in the left box refer to one component of the heat flux q_{xx} , whereas the error metrics in the right box are global over all components. The best performing model is marked with boldface in each metric.

| | \bar{q}_{xx} [%] | std $\bar{\sigma}_{xx}$ [%] | skew $\bar{\mu}_{xx}$ [%] | peak flux _{xx} [%] | MSE [$\cdot 10^{-5}$] | rel $\sqrt{\text{MSE}}$ [%](field) | rel $\sqrt{\text{MSE}}$ [%] (homogenized) | \approx params |
|-----------------------------|--------------------|-----------------------------|---------------------------|-----------------------------|-------------------------|------------------------------------|---|------------------|
| U-ResNet | 0.35 | 1.95 | 19.80 | 3.31 | 9.35 | 2.74 | 1.11 | 780 000 |
| \checkmark direct | 1.41 | 2.22 | 84.96 | 32.36 | 60.60 | 6.97 | 1.73 | 310 000 |
| \checkmark pretrained | 0.21 | 0.61 | 14.14 | 7.79 | 6.71 | 2.32 | 0.51 | 310 000 |
| \checkmark direct | 0.33 | 0.71 | 15.96 | 5.06 | 7.51 | 2.45 | 0.64 | 505 000 |
| \checkmark pretrained | 0.19 | 0.61 | 10.13 | 2.33 | 3.11 | 1.58 | 0.46 | 505 000 |
| \checkmark multi-pretrain | 0.23 | 0.57 | 9.14 | 2.14 | 2.96 | 1.54 | 0.41 | 505 000 |

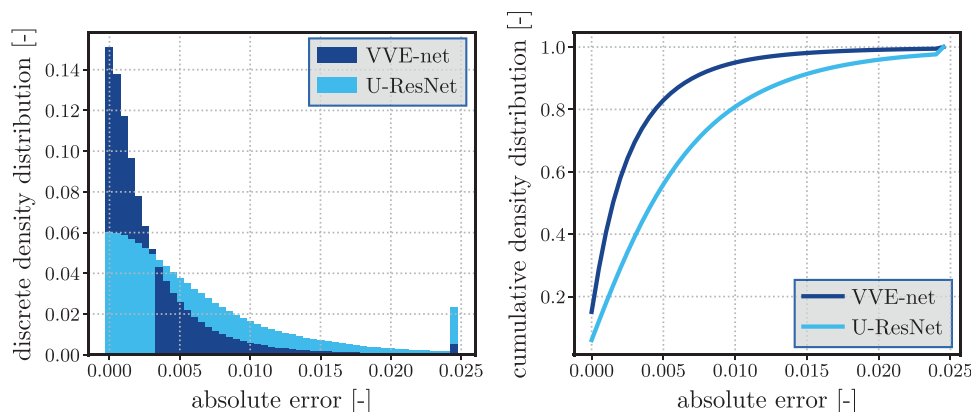


FIGURE 4 The empirical discrete density function as well as the cumulative density function of pixel errors are shown for the entire test set over all components of the heat flux.

An increasingly strict error measure is considered in Figure 4, since Figure 3 disregards the position of the pixels, for example, a random permutation of the predicted pixels would not be reflected in the error metric. The pixel error compares each pixel with the solution at its position. Once again, a distribution over the entire test set and all components of \mathbf{q} is presented. Every prediction error larger than $e \geq 0.25$ is capped and summarized to 0.25, which hides the long tail of the U-ResNet's prediction errors and ignores the single pixel maximum prediction error. In general, the left skew of the error distribution of the VVE-net implies that most of the pixels have a very low prediction error, which is also reflected in the cumulative pixel error distribution, where more than 80% of pixels have a lower absolute error than $5 \cdot 10^{-3}$.

5 | CONCLUSION

We developed a new efficient model layout of fully convolutional neural networks, which significantly outperforms recent state-of-the-art models while having fewer model parameters. There is multiple contributions improving the model, one being a staggered optimization scheme to ensure an excellent low-frequency prediction. The other improvement is found by superimposing a U-Net to refine the low-level prediction, reducing the errors on the peak values by 70%. The code of the model implementation is made freely available in Lissner [13]. The model has been applied in microstructure modeling under thermal boundary conditions, achieving a homogenization error of 0.2%. The model could be easily adopted to different physical boundary conditions once the data are readily available.

ACKNOWLEDGMENTS

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2075-390740016. Contributions by Felix Fritzen are funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) within the Heisenberg program DFG-FR2702/8-406068690 and DFG-FR2702/10-517847245. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech).

Open access funding enabled and organized by Projekt DEAL.

ORCID

Julian Lißner  <https://orcid.org/0000-0002-2286-5211>

Felix Fritzen  <https://orcid.org/0000-0003-4926-0068>

REFERENCES

1. Bažant, Z. P. (1999). Size effect on structural strength: A review. *Archive of applied Mechanics*, 69, 703–725.
2. Miehe, C. (2002). Strain-driven homogenization of inelastic microstructures and composites based on an incremental variational formulation. *International Journal for Numerical Methods in Engineering*, 55, 1285–1322.
3. Keshav, S., Fritzen, F., & Kabel, M. (2023). FFT-based homogenization at finite strains using composite boxels (ComBo). *Computational Mechanics*, 71(1), 191–212.

4. Li, H., Kafka, O. L., Gao, J., Yu, C., Nie, Y., Zhang, L., Tajdari, M., Tang, S., Guo, X., Li, G., Tang, S., Cheng, G., Liu, & W. K. (2019). Clustering discretization methods for generation of material performance databases in machine learning and design optimization. *Computational Mechanics*, *64*, 281–305.
5. Lißner, J., & Fritzen, F. (2019). Data-driven microstructure property relations. *Mathematical and Computational Applications*, *24*(2), 57.
6. Eidel, B. (2023). Deep CNNs as universal predictors of elasticity tensors in homogenization. *Computer Methods in Applied Mechanics and Engineering*, *403*, 115741.
7. Zhou, X. H., McClure, J. E., Chen, C., & Xiao, H. (2022). Neural network-based pore flow field prediction in porous media using super resolution. *Physical Review Fluids*, *7*(7), 074302.
8. Santos, J. E., Yin, Y., Jo, H., Pan, W., Kang, Q., Viswanathan, H. S., Prodanović, M., Pyrcz, M. J., & Lubbers, N. (2021). Computationally efficient multiscale neural networks applied to fluid flow in complex 3D porous media. *Transport in Porous Media*, *140*(1), 241–272.
9. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (pp. 234–241).
10. Wang, F., Eljarrat, A., Müller, J., Henninen, T. R., Erni, R., & Koch, C. T. (2020). Multi-resolution convolutional neural networks for inverse problems. *Scientific Reports*, *10*(1), 1–11.
11. Marcato, A., Santos, J. E., Boccardo, G., Viswanathan, H., Marchisio, D., & Prodanović, M. (2023). Prediction of local concentration fields in porous media with chemical reaction using a multi scale convolutional neural network. *Chemical Engineering Journal*, *455*, 140367.
12. Shaham, T. R., Dekel, T., & Michaeli, T. (2019). SinGAN: Learning a generative model from a single natural image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 4570–4580).
13. Lissner, J. (2023, June). TensorFlow models, plotting templates and python utilities. Computer Software, https://github.com/J-lissner/python_scripts
14. Leuschner, M. (2018). *Numerically efficient computational homogenization: Fourier-accelerated nodal solvers and reduced basis model order reduction*. Institute of Applied Mechanics.
15. Santos, J. E., Xu, D., Jo, H., Landry, C. J., Prodanović, M., & Pyrcz, M. J. (2020). PoreFlow-Net: A 3D convolutional neural network to predict fluid flow through porous media. *Advances in Water Resources*, *138*, 103539.
16. Dumoulin, V., & Visin, F. (2016). A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*.
17. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1–9).
18. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778).
19. Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4700–4708).
20. Lißner, J., & Fritzen, F. (2023). Hybrid machine-learned homogenization: Bayesian data mining and convolutional neural networks. *arXiv preprint arXiv:2302.12545*.

How to cite this article: Lißner, J., & Fritzen, F. (2023). Double U-Net: Improved multiscale modeling via fully convolutional neural networks. *Proceedings in Applied Mathematics and Mechanics*, *23*, e202300205.
<https://doi.org/10.1002/pamm.202300205>