



Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

MAGNET: A graph U-Net architecture for mesh-based simulations

Saurabh Deshpande^a, Stéphane P.A. Bordas^{a,*}, Jakub Lengiewicz^{a,b}^a Department of Engineering; Faculty of Science, Technology and Medicine; University of Luxembourg, Luxembourg^b Institute of Fundamental Technological Research, Polish Academy of Sciences, Poland

ARTICLE INFO

Dataset link: <https://github.com/saurabhdeshpande93/MAGNET>, <https://doi.org/10.5281/zenodo.7784804>

Keywords:

Geometric deep learning
Mesh based simulations
Finite element method
Graph U-Net
Surrogate modeling

ABSTRACT

In many cutting-edge applications, high-fidelity computational models prove to be too slow for practical use and are therefore replaced by much faster surrogate models. Recently, deep learning techniques have increasingly been utilized to accelerate such predictions. To enable learning on large-dimensional and complex data, specific neural network architectures have been developed, including convolutional and graph neural networks. In this work, we present a novel encoder–decoder geometric deep learning framework called MAGNET, which extends the well-known convolutional neural networks to accommodate arbitrary graph-structured data. MAGNET consists of innovative Multichannel Aggregation (MAg) layers and graph pooling/unpooling layers, forming a graph U-Net architecture that is analogous to convolutional U-Nets. We demonstrate the predictive capabilities of MAGNET in surrogate modeling for non-linear finite element simulations in the mechanics of solids.

1. Introduction

Computational models are essential tools for studying, designing, and controlling complex systems in many fields, including engineering, physics, biology, economics, and social networks. These models are often based on physical laws and mathematical equations, with partial differential equations (PDEs) being a common tool for describing how quantities change over space and time. In mechanics and physics, the PDEs are most commonly solved with numerical methods upon earlier space- and time- discretization, and a large number of domain-specific computational models have been developed so far, with the finite element method (FEM) and the finite volume method (FVM) being the most commonly used approaches in solid- and fluid mechanics, respectively. However, despite significant advances in computational performance over the last decade, such high-fidelity numerical simulations remain prohibitively expensive for many important applications, including emerging areas such as real-time feedback/control in the computer-assisted surgery (Johnsen et al., 2015; Bui et al., 2018) or soft robotics (Rus and Tolley, 2015; Gouy and Duriez, 2018). Speeding up such models whilst maintaining the desired accuracy is an active area of research, and one of the main motivations of the present work.

Recently, deep learning (DL) techniques have taken a center stage across many disciplines. The DL models have proven to be accurate and efficient in predicting non-trivial nonlinear relationships in data. As such, they have been used for a variety of applications, including surrogate modeling in mechanics (Mendizabal et al., 2019; Nikolopoulos et al., 2022; Deshpande et al., 2022; Krokos et al., 2022; Zhang

et al., 2023; Šarkić Glumac et al., 2023), or model discovery and calibration (Huang et al., 2020; Thakolkaran et al., 2022). The deep neural network approaches can be categorized with respect to how they use the data and *a priori* knowledge about the modeled system. In purely data-driven approaches, DL models rely on performing supervised learning on either experimental or numerically generated data and are agnostic to the underlying physics or model. As such, they are able to reproduce the physics-based relationship by implicitly learning on a relatively large amount of data (Runge et al., 2017; Aydin et al., 2019; Daniel et al., 2020; Kochkov et al., 2021; Hoq et al., 2023). If the *a priori* information about the modeled system is introduced, such networks are termed as Physics Informed Neural Networks (PINNs) (Raissi et al., 2019; Samaniego et al., 2020; Henkes et al., 2022; Nguyen et al., 2022; Roy and Guha, 2023). With respect to the purely data-driven approaches, PINNs are generally more accurate, require less data for training, and possess better generalization capabilities. The framework presented in the present work is generally applicable to both cases, however, for the sake of clarity, we will later only focus on purely data-driven types of networks. In any case, once trained, the DL models can be used as fast surrogates for computationally expensive high-fidelity numerical methods.

The focus of the present work is on high-dimensional relationships in which the sizes of inputs and/or outputs are large. Examples of such relationships can be found, for instance, in experimental full-field measurement data, such as our recent work on medical imaging (Lavigne et al., 2022), or in synthetic mesh data generated from finite

* Corresponding author.

E-mail address: stephane.bordas@alum.northwestern.edu (S.P.A. Bordas).

element simulations, Lorente et al. (2017) and Pellicer-Valero et al. (2020). Although DL techniques have generally shown great success as efficient surrogates to computationally expensive numerical methods in scientific computing, some of the popular existing machine learning approaches are still based on fully-connected deep networks which are not suitable for high-dimensional inputs/outputs. As an alternative, the application of Convolutional Neural Networks (CNNs) has proven a promising performance in a wide variety of applications, also including accelerating non-linear finite element/volume/difference simulations (Obiols-Sales et al., 2020; Rao and Liu, 2020; Deshpande et al., 2022; Zhao et al., 2023). CNNs are designed to learn a set of fixed-size trainable local filters (convolutions), thus reducing the parameter space while being capable to capture non-linearities. In the context of computational mechanics, local convolutions leverage the natural local correlation of nearby nodes, which leads to more efficient neural network architectures, both in terms of training- and prediction times. Moreover, one can observe that the CNN architectures have a close analogy to some iterative solution schemes known in scientific computing (Wang et al., 2020; Brenner and Scott, 2008). This provides them with an additional interpretation of being trainable iterative computational schemes to solve sets of non-linear equations, rather than general-purpose black-box approximators.

However, there is one important limitation that prevents CNNs from being of general purpose. The problem is that they only work well with grid-like structure data, such as images or structured meshes, which greatly hinders their use for many real-world applications where data is structured differently. Although there are some attempts to alleviate that problem in the context of FEM data, for instance, combining finite elements with an immersed-boundary method (Brunet et al., 2019), or embedding a precomputed coordinate mapping into the classic grid (Gao et al., 2021), the effectiveness of those methods is limited to simple irregular domains and remains challenging for complex geometries in general. A definitive solution to that problem has only been brought by Graph Neural Networks (GNNs)—architectures that directly handle arbitrarily-structured inputs/outputs. They belong to the recently emerged family of Geometric Deep Learning (GDL) methods which focus on neural networks that can learn from non-Euclidean input such as graphs and, more generally, manifolds (Bronstein et al., 2017; Wu et al., 2021). Because of their ability to handle more general structured data, GNNs are gaining increasing importance also in surrogate modeling in scientific computing (Sanchez-Gonzalez et al., 2020; Vlassis et al., 2020; Pfaff et al., 2021; Gao et al., 2022; Shivaditya et al., 2022; Seo and Min, 2023; Jiang and Chen, 2023; Krokos et al., 2024). However, these approaches are based on relatively simple message passing schemes, which are sub-optimal for learning on high non-linear regression tasks. In this work, we propose a novel local aggregation technique, which we denote as Multichannel Aggregation layer, MAG, which performs multichannel localized weighted aggregations, that can be seen as a direct extension of the traditional convolution layer in CNNs. Thanks to that, we are able to directly adapt some of the mechanisms/layers developed for CNNs to create efficient graph neural network architectures.

One mechanism that can improve the efficiency and predictive capabilities of convolutional and graph neural networks is the application of down-sampling (coarsening) and up-sampling (refinement) layers. In the context of CNNs, the focus is on encoder–decoder architecture frameworks, such as U-Net, which has been successfully implemented in various applications, including computer vision (Ronneberger et al., 2015; Çiçek et al., 2016), signal processing (Hennequin et al., 2020; Ren et al., 2021), and scientific machine learning (Mendizabal et al., 2019; Wang et al., 2020; Pant et al., 2021; Le et al., 2022; Nikolopoulos et al., 2022; Fernández-León et al., 2023). While the CNN-based U-Net approaches are limited to grid data, their graph-based version, known as graph U-Net, can provide the desired generality. Recently, various graph coarsening approaches have been proposed (Bianchi et al., 2019; Lee et al., 2019; Gao and Ji, 2019; Cai et al., 2021), which

serve the same function as pooling layers in CNNs, helping to reduce the size of a graph while maintaining essential properties of the processed data. In this work, we propose a novel graph pooling/unpooling operation (coarsening/refinement), that enables us to create a graph U-Net architecture, MAGNET, that can operate on arbitrary graphs. Our pooling layers are directly inspired by CNNs, where we extend the concept of pooling over local patches in regular grids to variable size non-overlapping cliques in graphs. This allows us to precompute coarsened graphs that are only based on the input graph topology, which is independent of data (i.e., node features). In the context of GNN-accelerated FEM simulations, a similar concept has been proposed by Black and Najafi (2022), however, their implementation is limited to regular meshes for simple two-dimensional geometries and linear elastic problems. Our approach enables computationally efficient deep learning models for non-linear problems involving arbitrary meshes, which is an important advancement for this field.

In summary, we introduce a novel graph U-Net framework comprising the proposed MAG and graph pooling/unpooling layers. The MAG layer captures local regularities in the input data, while the interleaved pooling layers reduce the graph representation to a smaller size while preserving important structural information. This enables us to efficiently implement our framework for large-scale problems. The proposed MAG and graph pooling layers are direct analogues of respective CNN U-Net layers and are also compatible with many state-of-the-art graph neural network layers. We elaborate on this point in the paper, providing a qualitative comparison of the proposed MAG layer with several existing graph layers. To validate the predictive capabilities of our framework, we apply it to several non-linear relationships obtained through finite element analysis and cross-validate it with predictions given by our CNN U-Net architecture (Deshpande et al., 2022). To increase the impact of our work, we provide source codes, datasets, and procedures to generate the datasets utilized in this work, which can be found in the MAGNET repository: <https://github.com/saurabhdeshpande93/MAGNET>.

The paper is organized as follows. In Section 2 we present the novel MAGNET framework, as well as its particular application to the hyperelastic FEM-based datasets. Then, in Section 3, we provide details of implementation and a thorough study of MAGNET for several 2D and 3D benchmark non-linear FEM examples. The conclusions and future research directions are summarized in Section 4.

2. MAGNET deep learning framework

In this section, we will propose a novel graph-based encoder–decoder (U-Net) deep-learning framework. We will provide a general formulation, in which inputs and outputs follow a certain graph topology (that is expressed by an adjacency matrix A). The graph expresses an assumed structure of correlations within input/output data and allows us to devise a *robust* DNN architecture defining a non-linear mapping between inputs and outputs. We will apply this general framework to *mesh-based* graphs. Such mesh topology of data is characteristic to spatially discretized numerical solution schemes for PDEs in scientific computing. In particular, we will focus on hyperelastic problems in solid mechanics, for which the training/testing data is obtained through the finite element method (see also the schematics in Fig. 1).

In Section 2.1 we will provide an overview of the proposed Graph U-Net framework MAGNET. Next, in Sections 2.2–2.4 we will introduce the building blocks of MAGNET. In particular, in Section 2.2, we will introduce the adjacency matrix representation of the mesh-based graph, which will be utilized later in this paper; and in Sections 2.3–2.4 we will specify a new graph Multi-channel Aggregation (MAG) layer, as well as new graph pooling/unpooling layers. Afterwards, in Section 2.5, we will provide an information-passing interpretation of the proposed Graph U-Net architecture. Finally, in Section 2.6 we will introduce a particular application of the framework to mesh-based datasets that are generated from FEM solutions of problems in hyper-elasticity.




Fig. 1. A novel graph U-Net neural network surrogate model for mesh-based simulations. MAGNET accurately captures non-linear FEM responses.

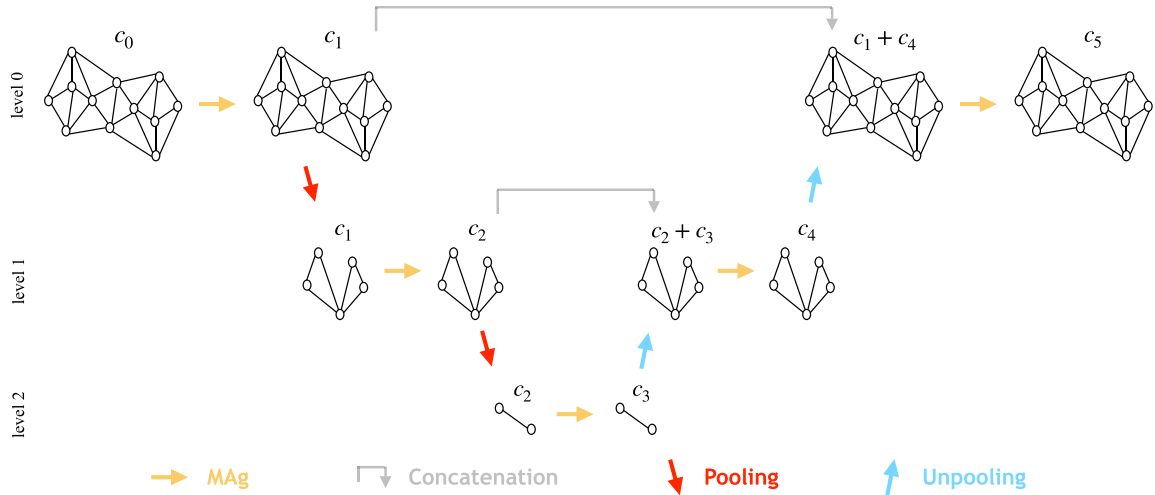


Fig. 2. A schematic of Graph U-Net architecture for mesh based inputs. Colors indicate different types of layers. c_1, c_2, \dots, c_5 stand for channel dimensions. Different arrows indicate different layers: the graph Multi-channel Aggregation (MAG) layer, the graph pooling/unpooling layers, and the concatenation layer.

2.1. MAGNET architecture overview

The MAGNET graph neural network architecture can be classified as a graph U-Net network and is an extension to the well-known class of convolution-based U-Net architectures (see Ronneberger et al. (2015)). As such, the graph U-Net comprises of aggregation (‘convolution’), pooling, unpooling, and concatenation layers (see the schematics in Fig. 2), which are here suitably adjusted to work with general (non-grid) topologies of inputs/outputs.

The graph U-Net architecture has two stages: encoding and decoding. In the encoding stage, first, we apply one or more aggregation (MAG) layers, which are analogues of convolution layers in non-graph U-Net networks. Next, we apply a single graph pooling layer, which is a particular contraction of the graph, and which downsamples (coarsens) the problem. This aggregation-pooling sequence is repeated several times to achieve the desired level of contraction (coarsening). At the coarsest level, the MAG aggregation is performed one or more times, after which the decoding stage begins, which is the opposite to the encoding stage. At each level of decoding, the graph unpooling layer

is followed by one or more MAG layers. At the upmost level, the last MAG layer is applied with linear activation to get the output.

More formally, the Graph U-Net network, \mathcal{G} , is constructed as follows. First, we set the input layer d^0 as a vector of N nodes, each of which being a vector of input values (also known as features or channels) of a constant length c_0 . (Further on, we will refer to the features as the *channels*.) Next, we subsequently add layers, d^l , to form a U-Net architecture. The subsequent layers, d^{l-1} and d^l , are linked by the following relationship

$$d^l = \mathbf{T}^l(d^{l-1}; \theta^l), \quad (1)$$

where θ^l is a vector of trainable parameters (e.g., weights and biases, $\theta^l = \mathbf{k}^l \cup \mathbf{b}^l$), and $\mathbf{T}^l(\cdot)$ is one of three already introduced transformations: MAG(), gPool() or gUnpool(), which will be more precisely defined in the following sections. Additionally, we also consider remote concatenation links between respective layers from the encoding and decoding stages, see Fig. 2. The output layer, d^L , is assumed to be of the same mesh format as the input layer but can have a different number of channels (features), c_L . Finally, we define the Graph U-Net network




Fig. 3. Adjacency matrices for the (a) square and (b) triangular meshes. The dashed lines in (a) represent additional edges that are added to the original mesh.

as a parameterized transformation

$$\mathcal{G}(d^0, \theta) = d^L = \mathbf{T}^L(\mathbf{T}^{L-1}(\mathbf{T}^{L-2}(\dots); \theta^{L-1}); \theta^L), \quad (2)$$

where $\theta = \bigcup_{l=1}^L \{\theta^l\}$ is a concatenated vector of network parameters.

The calibration of the Graph U-Net parameters is done through a supervised learning, by fitting a given known input–output training dataset. The training dataset,

$$\mathcal{D}_{\text{tr}} = \{(\mathbf{f}_1, \mathbf{u}_1), \dots, (\mathbf{f}_{M_{\text{tr}}}, \mathbf{u}_{M_{\text{tr}}})\}, \quad (3)$$

is in the mesh format, and the training is done by minimizing the following mean squared error loss function

$$\mathcal{L}(\mathcal{D}_{\text{tr}}, \theta) = \frac{1}{M_{\text{tr}}} \sum_{m=1}^{M_{\text{tr}}} \|\mathcal{G}(\mathbf{f}_m, \theta) - \mathbf{u}_m\|^2 \quad (4)$$

which gives the optimal parameters

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\mathcal{D}_{\text{tr}}, \theta). \quad (5)$$

2.2. Adjacency matrix of the mesh-based graph

For the purpose of this work, we will focus on sparse graphs that derive from data that is spatially organized in the form of meshes. Those can be 1D, 2D, or higher-dimensional meshes, of an arbitrary connection topology (see Fig. 3 for examples of 2D meshes). The graph can be conveniently represented by a symmetric, square, Boolean adjacency matrix, A , whose order is equal to the number of nodes in the original mesh. To simplify the further notation, all nodes (vertices) are self-connected (have loops), which results in having 1 on the diagonal of A . This allows us to more easily express certain graph operations that are used in this work, for instance, the k th power of a graph A , and the selection of pooling sub-graphs that is presented in Section 2.4.

It is fairly straightforward to generate an adjacency matrix from an element connectivity matrix of the mesh. For that reason we will not discuss it in detail. The only point to be emphasized is that we make all nodes belonging to a given element mutually inter-connected in the resulting graph (as they can be assumed to be strongly inter-related). We can visually represent it by adding more links as compared to a standard wire-frame visualization of finite-elements (see, e.g., the dashed lines in Fig. 3a).

Remark. In our work we do not consider any attributes for the edges of a graph. Therefore, the data is only represented through nodal features and node–node connections which are defined through the adjacency matrix A .

2.3. Multi-channel Aggregation (MAg) layer

The proposed novel neural network layer, MAg, is a multi-channel local aggregation layer that can operate on graph-structured data. Its architecture is a direct extension to the standard convolutional layer in CNNs, in which a shareable convolution window is used, making CNNs restricted to grid-structured data. In the MAg layer, instead, we propose to use fully-trainable local weighted aggregations (the so-called message passing scheme), where the aggregation neighborhood of a given node is prescribed through the graph connectivity (the adjacency

matrix). As such, the scheme is very well suited for sparse graphs and can be directly applied to graphs that derive from arbitrary 2D or 3D meshes.

The use of multiple channels aims to improve the capabilities of the network to capture non-linearities. In the multi-channel version, each node represents a vector of values (features), which can be visualized as multiple layers (channels) of the same graph structure (see the schematics in Fig. 4(a)). The transformation between the input- and output multi-channel graphs is realized by applying multiple MAg aggregations on vector data to produce respective multiple components of output vectors. Note that the input/output channels of the whole network have usually a certain meaning, and their sizes are fixed (e.g., three RGB channels of a color image at the input and a single channel of a segmented image at the output). The number of channels in the latent layers can be chosen arbitrarily, which is up to the choice of a designer of a particular graph U-Net architecture.

More formally, we will consider the MAg layer as a parameterized transformation between the input and output nodes, defined as

$$d_{i,\alpha}^{l+1} = \sigma(b_{i,\alpha}^{l+1} + \sum_{\beta=1}^{c^l} \sum_{j \in \mathcal{N}_i} k_{i,j,\alpha,\beta}^{l+1} d_{j,\beta}^l), \quad (6)$$

where $\mathcal{N}_i = \{j \mid A_{ij} = 1\}$ is a set of neighbors of a node i to be aggregated, α and β represent the output and input channels, respectively, while k^{l+1} and b^{l+1} are trainable weights and biases, respectively. In this multi-channel definition, for a given component, $d_{i,\alpha}^{l+1}$, of an output, a single aggregation is performed throughout the neighborhood, \mathcal{N}_i , and all the input channels, $\beta \in \{1, \dots, c^l\}$. The kernel parameters of MAg transformation, $k_{i,j,\alpha,\beta}^{l+1}$, are not shared, i.e. they can be independently trained for each aggregation window (note the free indexes i and α).

2.3.1. Comparison to existing graph aggregation/convolution layers

As already mentioned in the introduction, the very idea of generalization of convolution layers to arbitrary graph structure is not new. In fact, various concepts have emerged so far, Zhou et al. (2020) and Chen et al. (2020), most of which are compatible with the U-Net framework proposed in the present work. Below, we will discuss several of them, introducing a unified notation that will facilitate a qualitative comparison with respect to the proposed MAg layer, see Table 1.

Graph neural network layers aim to utilize the information about assumed correlations in data, with the graph structure expressing those correlations. The general approach is to specify a suitable (possibly nonlinear) trainable local transformation that can aggregate the information from a node in consideration and its neighbors. (This aggregation is followed by a chosen activation function before being propagated to the next layer.) Such transformations form a wide class of, so-called, message passing schemes, and can combine shareable (independent of a node) and non-shareable (dependent on a node, i.e., independently-trainable) sets of parameters.

The simplest and most lightweight realizations of the graph aggregation/convolution layer concept only utilize shareable weights, see, e.g., the Graph Convolutional Network (GCN), Kipf and Welling (2016). In those approaches, a non-trainable (arbitrary) weighted aggregation is performed prior to application of a shareable trainable operator — something completely opposite to our MAg layer, which

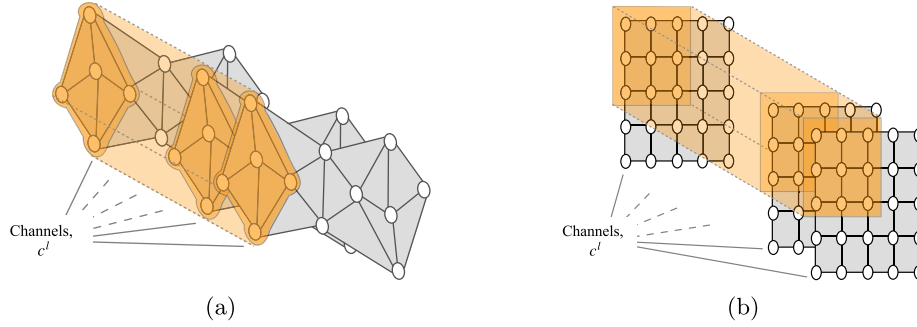


Fig. 4. Local aggregation in MAg (a) works very similar to the filter application in CNN (b). However as opposed to CNN, MAg uses different set of weights at different spatial locations with heterogeneous window size. In CNN, a constant filter slides across the channel.

Table 1

Comparison of the MAg layer with selected state of the art graph convolution layers (biases are omitted for the sake of brevity). In GAT formulation, $\alpha = (t-1)N_g + \gamma$, which represents the stacking operation for the multi-head attention mechanism, where $t \in (1, \dots, N_h^{l+1})$ is the attention head index, and $\gamma \in (1, \dots, N_g)$ is the internal channel index (cardinality of each node in the layer $l+1$).

Layer	Transformation
GCN (Kipf and Welling, 2016)	$d_{i,\alpha}^{l+1} = \sigma(\sum_{\beta=1}^{c^l} w_{\alpha,\beta}^{l+1} \sum_{j \in \mathcal{N}_i} \frac{A_{j,\alpha}}{\sum_{x \in \mathcal{N}_i} A_{j,x}} d_{j,\beta}^l)$
GAT (Veličković et al., 2017)	$d_{i,\alpha}^{l+1} = \sigma(\sum_{\beta=1}^{c^l} w_{i,\gamma,\beta}^{l+1} \sum_{j \in \mathcal{N}_i} \text{softmax}_j(\text{attn}(w_i^{l+1} d_i^l, w_j^{l+1} d_j^l, \theta_i)) d_{j,\beta}^l)$
SemGCN (Zhao et al., 2019)	$d_{i,\alpha}^{l+1} = \sigma(\sum_{\beta=1}^{c^l} w_{\alpha,\beta}^{l+1} \sum_{j \in \mathcal{N}_i} \text{softmax}_j(k_{i,j,\alpha,\beta}^{l+1}) d_{j,\beta}^l)$
MAg [present work]	$d_{i,\alpha}^{l+1} = \sigma(\sum_{\beta=1}^{c^l} \sum_{j \in \mathcal{N}_i} k_{i,j,\alpha,\beta}^{l+1} d_{j,\beta}^l)$

is fully trainable. This enables to keep the number of trainable parameters low, which is achieved at the cost of relatively low capacity of such networks. This low capacity cannot be straightforwardly increased by simply deepening the network because of the well-known over-smoothing phenomenon.

We will discuss two out of many available approaches to increase the capacity of graph neural networks. The first approach relies on the multi-head attention mechanism which allows to assign different importance to nodes in the neighborhood, see, e.g, the Graph Attention Network (GAT), Veličković et al. (2017). In the attention mechanism, the weights used in local aggregation depend on input nodal features, which makes the concept qualitatively different from all approaches (including the MAg layer) that use input-independent aggregation weights. The second class of approaches resemble the MAg layer more closely. Particular notable examples of that approach are the Spatial-Temporal Graph Convolution Network (ST-GCN), Yan et al. (2018), and the Semantic Graph Convolution Network (SemGCN), Zhao et al. (2019), which have been introduced in the particular context of human pose recognition problem (the computer vision domain). The common features of MAg and SemGCN layers are the input-independent learnable weighted aggregation and the use of channels to increase the model capacity. The difference is that the MAg does not use a shared transformation matrix (w) nor the softmax normalization — both used in the case of SemGCN.

To summarize, the proposed MAg layer relies on one of the most flexible message-passing schemes, with no shareable parameters. This promises a very high capacity of the MAg network. Also, as shown above, the proposed MAg layer is compatible with other graph convolution/aggregation layer concepts, and thus can be straightforwardly exchanged, if needed.

2.4. Graph pooling- and unpooling layers

Pooling and unpooling are two fundamental operations allowing U-Nets to encode (compress) and decode (decompress) information, respectively, see Fig. 2. The pooling layers are composed of local contracting operations over the mesh-structured data, and are used to coarsen the data at the encoding part of the network. At the decoding part, the original refined mesh structure is restored by the unpooling

layers (upsampling operations). In U-Nets, the unpooling layer is usually combined with the concatenation operation, which creates a direct link between the encoding and decoding part of the network (this will be explained later).

Graph pooling

In this work, we propose a novel clustering-based graph pooling layer that can be applied to arbitrary graph-structured data. It can be seen as an extension to the pooling layers known from CNN U-Nets that are limited to grid-structured data. In our graph pooling approach, we split the graph into disjoint cliques (fully-connected subgraphs), and perform the contraction of all the identified cliques (i.e., every clique is replaced by a vertex, and new edges represent formerly connected cliques), see Fig. 5. The split into cliques is done statically, i.e., at the graph U-Net construction phase. In particular, the split does not depend on the input data.

Below, we will provide a more formal construction of the pooling layer. For a given input graph G that is represented by the vertices S and the connectivity matrix A , we first generate an arbitrary set of \tilde{N} non-overlapping fully-connected subgraphs (cliques) $G_1, G_2, \dots, G_{\tilde{N}}$, i.e.,

$$S = \bigcup_{i=1}^{\tilde{N}} S_i, \quad \forall S_i \forall j,k \in S_i A_{jk} = 1 \quad \text{and} \quad \forall i \neq j S_i \cap S_j = \emptyset, \quad (7)$$

where the sets S_i represent nodes of the respective subgraphs G_i . The procedure to generate these subgraphs and the respective pooled adjacency matrix \tilde{A} is described in Algorithm 1. The pooled graph \tilde{G} is composed of vertexes $\tilde{S} = \{1, \dots, \tilde{N}\}$ with edges defined by the adjacency matrix \tilde{A} . The pooling layer is described as:

$$d_{i,\beta}^{l+1} = \text{aggr}_{j \in S_i} d_{j,\beta}^l, \quad (8)$$

where the 'aggr' operation can be a max/min/avg, etc. Note that graph pooling layers do not modify the number of channels, i.e., the pooling is performed individually per each channel of the input.

Graph pooling can be applied several times at the encoding part of the U-Net, e.g., see Fig. 2. For the purpose of future unpooling operations, after each pooling operation, we save the original graph, G , the adjacency matrix, A , and the pooling subgraphs, G_i . After doing so, we substitute $G \leftarrow \tilde{G}$, and $A \leftarrow \tilde{A}$.

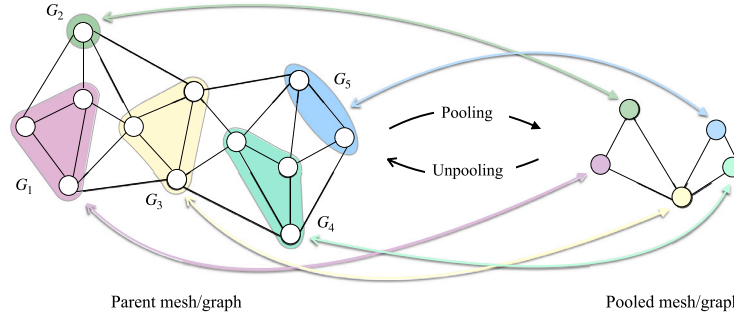


Fig. 5. One arbitrary choice of non-overlapping subgraphs to create a pooled graph. Subgraphs G_1, \dots, G_5 are represented with different colors and are generated by the Algorithm 1.

Algorithm 1: Generate a pooled graph from an arbitrary parent graph

Input: $N \times N$ adjacency matrix, A

Result: list of subgraphs, S ; $\tilde{N} \times \tilde{N}$ pooled adjacency matrix, \tilde{A}

$S \leftarrow \{\}$ /* initialisation of the subgraph list */

$P \leftarrow \{1, 2, \dots, N\}$ /* node indices of the parent graph */

$A' \leftarrow A$ /* temporary copy of matrix A */

/* Loop for generating non-overlapping subgraphs, S , see Fig. 5 */

while $P \neq \text{null}$ **do**

$p \in P$ /* randomly select a single node */

$S_i \leftarrow \{p\}$ /* initialise subgraph */

$\mathcal{N}_p \leftarrow \{m \neq p \mid A'[m, p] = 1\}$ /* nodes connected to selected node */

for n in \mathcal{N}_p **do**

if $\forall m \in S_i, A'[m, n] = 1$ **then**

$S_i \leftarrow S_i \cup \{n\}$ /* append node to the subgraph */

end

end

$P \leftarrow P \setminus S_i$ /* remove subgraph from parent graph */

$\forall m \in S_i, A'[m, :] \leftarrow 0; A'[:, m] \leftarrow 0$ /* remove subgraph from parent graph */

$S \leftarrow S \cup S_i$ /* add subgraph to subgraphs list */

end

$\tilde{N} = \text{sizeof}(S)$ /* number of pooled nodes = number of pooling subgraphs */

$\tilde{A} \leftarrow \text{zeros}(\tilde{N}, \tilde{N})$ /* zero initialisation of pooled matrix */

/* Loop for constructing pooled adjacency matrix \tilde{A} from subgraphs S */

for r in $\{1, 2, \dots, \tilde{N}\}$ **do**

for c in $\{1, 2, \dots, \tilde{N}\}$ **do**

if $\exists n \in S[r], m \in S[c] \mid A[n, m] = 1$ **then**

$\tilde{A}[r, c] \leftarrow 1$ /* if $S[r]$ and $S[c]$ are connected by an edge */

end

end

end

Graph unpooling

Structure-wise, the graph unpooling is a reverse operation to pooling. More precisely, the output graph of an unpooling layer will have the same topology as the input graph of the related pooling layer, see Fig. 5. The operation is defined via the previously saved subgraphs G_j (with nodes S_j) as

$$d_{i,\beta}^{l+1} = d_{j,\beta}^l \quad \text{for } i \in S_j, \quad (9)$$

and it simply replicates the features of a node j to the nodes specified by S_j . As such, this operation is analogous to the related upsampling operation used in CNNs.

Graph unpooling + concatenation

Concatenations, also known as skipped connections, are characteristic to U-Net architectures. Thanks to them, the layers from the decoder part gain a direct access to features from the encoder part. Concatenations help to mitigate the issue of vanishing gradients, and add extra information that could have been lost due to the earlier downsampling (pooling).

In our case, the concatenations are always related to the respective pooling/unpooling operation pairs, see Fig. 2. It is done by stacking the output of an unpooling layer l , given by Eq. (9), with the input of a respective pooling layer l' :

$$d_{i,c'+\alpha}^{l+1} = d_{i,\alpha}^{l'} \quad (10)$$

In the formula above, c' is the number of channels in unpooling inputs. As the result, the total number of output channels of unpooling+concatenation is $c' + c^l$.

2.5. Information-passing interpretation of MAg and pooling layers

During a single forward pass of the MAg layer, the aggregation is performed locally for each individual node, i.e., each node of the graph will have an access to the aggregated feature information from its adjacent nodes only, specified by the adjacency matrix, A , see Eq (6). Therefore, the nodes that are not directly connected through A do not exchange information at a single MAg operation (see Fig. 6). Such long-distance exchange across the network is fundamental to allow the neural network model to express correlations between topologically distant input- and output nodes (e.g., how the output displacements at node C depend on the input loads at the node B, in Fig. 6).

One way to handle this issue would be to apply the MAg layer several times as shown in Fig. 6. However, in that case, the number of subsequent layers would be proportional to the diameter of the underlying graph, which could increase the number of training variables and the depth of the network, deteriorating its performance. A natural simple improvement, also utilized in the present paper, is to increase the support (neighborhood) of the MAg operations. In the proposed framework, this can be straightforwardly done by using higher powers of the adjacency matrix, e.g., A^2 or A^3 , instead of A . This improvement alone, however, would still require the number of MAg layers to be proportional to the graph diameter.

The above observations explain a natural motivation behind using the pooling/unpooling layers, and hence creating the U-Net architecture. The pooled graph can be seen as a reduced space representation of the parent graph, and each pooled node aggregates the feature information corresponding to multiple nodes of the parent graph, see Fig. 7. The pooled graph is of a coarsened topology when compared to the parent graph, and this allows for the feature information exchange




Fig. 6. This 2D mesh requires at least 4 subsequent local aggregation operations (orange areas with center nodes marked by dots) to propagate the feature information from node B to the distant node C.

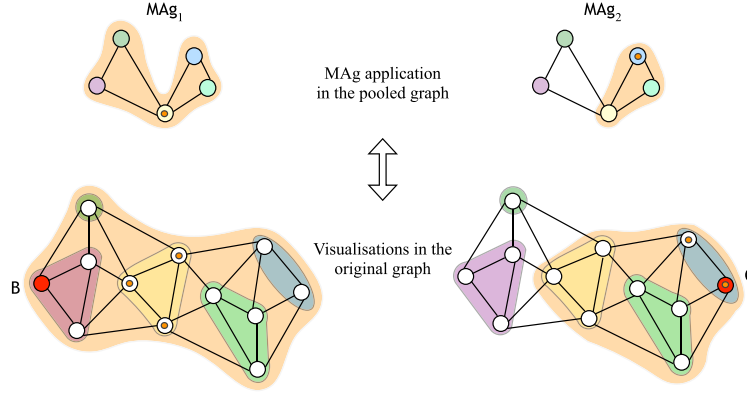


Fig. 7. Visualization of feature information exchange between nodes in the pooled graph. In the pooled space, only 2 MAG operations are sufficient to exchange feature information between spatially further located nodes in the original graph. The orange region shows the window of MAG operation.

with a lower number of MAG layers. The pooling/unpooling layers can be nested, which provides an exponential reduction rate of the graphs' diameters.

Remark. Note that the pooling layer proposed in this work represents a clique-pooling approach in which the cliques are non-overlapping. This allows us to achieve a very good level of graph coarsening (contraction). It is unlike a similar clique-based strategy that has been recently proposed, [Luzhnica et al. \(2019\)](#), in which the pooling cliques overlap, providing a much lower level of coarsening.

2.6. Application to FEM-based datasets

We will now focus on a particular graph structure of inputs/outputs that will be in a form of finite element mesh. Specifically, the MAGNET framework will be applied as a surrogate model to a finite element model in large-deformation elasticity. The finite element model will be shortly introduced below.

We consider a boundary value problem expressed in the weak form over the domain Ω :

$$\int_{\Omega} \mathbf{P}(\mathbf{F}(\mathbf{u})) \cdot \nabla \delta \mathbf{u} dV - \int_{\Omega} \rho \bar{\mathbf{b}} \cdot \delta \mathbf{u} dV - \int_{\Gamma_t} \bar{\mathbf{t}} \cdot \delta \mathbf{u} dS = 0 \quad \forall \delta \mathbf{u}, \quad (11)$$

where $\mathbf{P}(\bullet)$ is the first Piola–Kirchhoff stress tensor, $\bar{\mathbf{b}}$ are prescribed body forces, $\bar{\mathbf{t}}$ are prescribed tractions on the Neumann's boundary Γ_N , while the solution \mathbf{u} and the variation $\delta \mathbf{u}$ belong to appropriate functional spaces, with $\mathbf{u} = \bar{\mathbf{u}}$ and $\delta \mathbf{u} = \mathbf{0}$ on the Dirichlet boundary Γ_u . The hyperelastic constitutive relationship is expressed through the strain–energy density potential $W(\mathbf{F})$ as

$$\mathbf{P}(\mathbf{F}) = \frac{\partial W(\mathbf{F})}{\partial \mathbf{F}}, \quad (12)$$

where the deformation gradient tensor $\mathbf{F} = \mathbf{I} + \nabla \mathbf{u}$.

For all the cases considered in the present work, the Neo-Hookean hyperelastic law with the following strain energy potential is used, see ([Simo and Taylor, 1982](#)),

$$W(\mathbf{F}) = \frac{\mu}{2} (I_c - 3 - 2 \ln J) + \frac{\lambda}{4} (J^2 - 1 - 2 \ln J), \quad (13)$$

where the invariants J and I_c are given in terms of deformation gradient \mathbf{F} as

$$J = \det(\mathbf{F}), \quad I_c = \text{tr}(\mathbf{F}^T \mathbf{F}), \quad (14)$$

with μ and λ being Lamé's parameters, which can be expressed in terms of Young's modulus, E , and Poisson's ratio, ν , as

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}, \quad \mu = \frac{E}{2(1+\nu)}. \quad (15)$$

Note that one can use other forms of the volumetric part of the above potential, see [Doll and Schweizerhof \(2000\)](#), or other hyperelastic models, such as the Mooney–Rivlin and a more general class of Ogden models, see [Ogden \(2005\)](#).

Finite element discretization transforms the weak form expressed by Eq. (11) into the system of non-linear equations

$$\mathbf{R}(\mathbf{u}; \mathbf{f}_{\text{ext}}) = \mathbf{f}_{\text{int}}(\mathbf{u}) - \mathbf{f}_{\text{ext}} = \mathbf{0}, \quad (16)$$

that expresses the balance between external and internal nodal forces. In this work, the vector of external forces, \mathbf{f}_{ext} , represents boundary conditions, which can be surface tractions or body forces. Given $\mathbf{f}_{\text{ext}} = \mathbf{f}_m$, the system is solved for an unknown vector \mathbf{u} with the Newton–Raphson scheme, giving as a result the solution \mathbf{u}_m . A pair $(\mathbf{f}_m, \mathbf{u}_m)$ makes an element of the dataset \mathcal{D} introduced in Eq. (3), and the FE mesh that results from the FE discretization produces the adjacency matrix \mathbf{A} introduced in Section 2.2.

3. Results

In this section, we study the performance of the proposed framework in application to surrogate modeling in mechanics of solids. For that purpose, we use four benchmark problems. In Section 3.1, we give a detailed specification of the benchmark problems and the procedure for obtaining FEM-based datasets. In Section 3.2, we provide details of neural network architectures for each of the studied cases and will describe the training procedure. In Section 3.3, we study the predictions of neural network models by cross-validating results from MAGNET and CNN models, and by comparing them with the FEM results. Finally, in Section 3.4 we demonstrate the capabilities of the MAGNET framework to provide a surrogate model for the unstructured mesh cases.

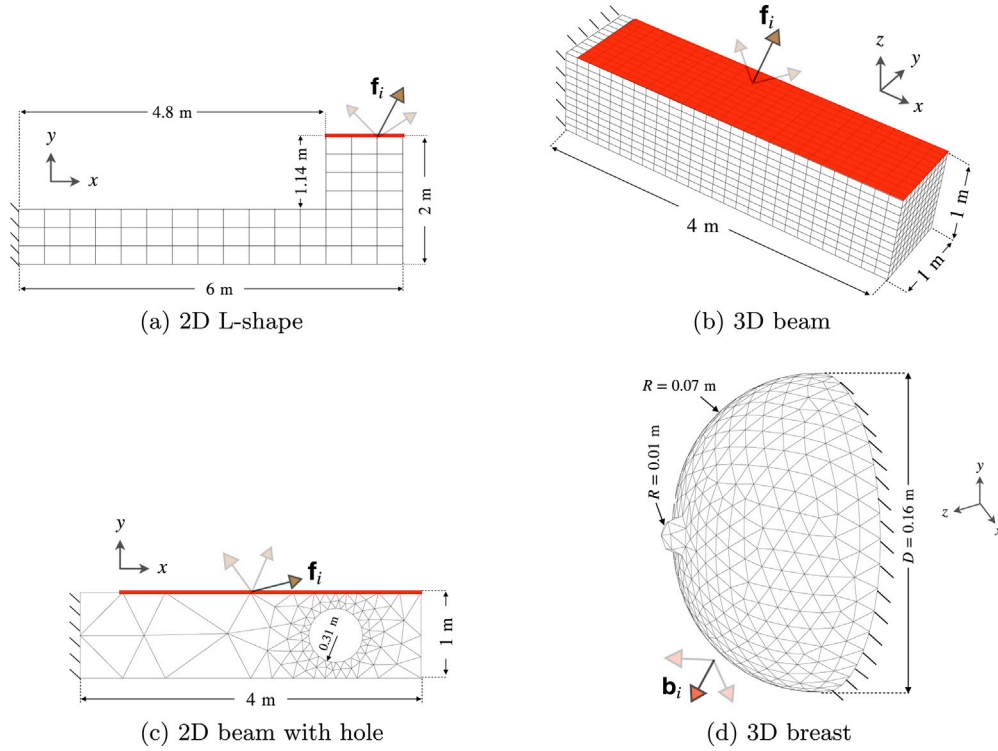


Fig. 8. Schematics of four benchmark problems. (a) 2D L-shape geometry (quad mesh), (b) 3D beam geometry (hexahedron mesh), (c) 2D beam with hole geometry (triangular mesh), and (d) 3D breast geometry (tetrahedron mesh). In examples (a)–(c), single nodal loads are applied on the region of boundary indicated with red color. In example (d), only body forces are considered.

Table 2
Material properties used for the benchmark cases.

	Problem (element topology)	Is structured?	Young's modulus, E [Pa]	Poisson's ratio, ν	Density, ρ [kg/m^3]
(a)	2D L-shape (quad)	No (Yes)	500	0.4	–
(b)	3D beam (hexahedron)	Yes	500	0.4	–
(c)	2D beam with hole (triangular)	No	500	0.3	–
(d)	3D breast (tetrahedron)	No	800	0.4	1000

3.1. Generation of FEM based datasets

We consider four benchmark problems, see Fig. 8. Two of them, Fig. 8(a–b), utilize simple meshes, which makes it possible to assure structured (grid) inputs. They will be used to cross-validate between our MAGNET architecture and the standard CNN U-Net architecture. The other two examples, Fig. 8(c–d), are more complex and will serve us to demonstrate the applicability of MAGNET for general (unstructured) meshes. Each of those two groups consists of a 2D and a 3D problem, thanks to which the framework can be tested for four different finite element topologies: triangular, quadrilateral, tetrahedral, and hexahedral.

For all considered cases, we use the neo-Hookean material model, see Section 2.6, with material parameters provided in the Table 2. In order to generate training/testing datasets, for each discretized problem we individually specify a family of boundary conditions, as described below, see also schematics in Fig. 8. For the cases shown in Fig. 8(a–c), nodes on one side are fixed (Dirichlet boundary conditions), and only a single random nodal force is applied at a selected node in a prescribed region of interest (denoted by red line/surface in Fig. 8(a–c)). For the remaining nodes, the external forces are set to $\mathbf{0}$. For the case shown in Fig. 8(d), the uniform body force density is prescribed (force per unit mass, with density $\rho = 1000 \text{ kg}/\text{m}^3$). The body force field is integrated through element shape functions to obtain respective nodal forces that are used in datasets.

All the finite element computations were implemented and performed within the AceGen/AceFem framework (Korelc, 2002). For a

given problem, for each loading case, i , the entire vector $\mathbf{f}_{(i)}$ of external nodal forces and the vector $\mathbf{u}_{(i)}$ of computed nodal displacements were saved, which allowed to generate the final training/testing dataset $D = \{(\mathbf{f}_{(1)}, \mathbf{u}_{(1)}), \dots, (\mathbf{f}_{(M_{\text{tr}}+M_{\text{te}})}, \mathbf{u}_{(M_{\text{tr}}+M_{\text{te}})})\}$. The datasets were randomly split into training sets, M_{tr} (95%), and testing sets, M_{te} (5%). The sizes of datasets and the distribution of force magnitudes are provided in Table 3.

3.2. Design, implementation and training of neural network models

The implementation of the layers and mechanisms of the MAGNET framework described in Section 2 and of CNN U-Net framework introduced in Deshpande et al. (2022) has been performed within the TensorFlow libraries. We use them to build and train deep neural network models for the cases described in Section 3.1. Table 4 outlines individual properties of the network architectures implemented in this work. The codes and datasets are publicly available open source (Deshpande et al., 2023a), which makes it possible for other researchers to reproduce the present results and also to extend our frameworks to new cases/problems.

To provide a complete understanding of the neural network architectures listed in Table 4, we will now delve into the details of the MAGNET architecture used for the 2D L-shape example. Its schematics is shown in Fig. 9. As indicated in the third column in Table 4, it is a three-level graph U-Net architecture with two MAG operations at each level. The fourth column specifies the number of channels utilized for the MAG operations at each level of the graph U-Net. The forward pass

Table 3


Specification of FE-based datasets. For cases (a–c), the external force is applied at a selected node, and for case (d), external body forces are applied. The magnitudes of forces are randomly sampled from the multivariate uniform distribution, with ranges specified in the table. For cases (a–c), multiple samples per node are generated, for all nodes in the prescribed area of interest.

	Problem	N. of FEM DOFs (F)	Range (External forces/body force density)	Dataset size $M_{tr} + M_{te}$	Samples per node
(a)	2D L-shape	160	$f_x, f_y = -1$ to 1 N	3800 + 200	1000
(b)	3D beam	12096	$f_x, f_y, f_z = -2$ to 2 N	33858 + 1782	110
(c)	2D beam (hole)	198	$f_x, f_y = -5$ to 5 N	4560 + 240	400
(d)	3D breast	3105	$b_x, b_y = -6$ to 6 N/kg, $b_z = -3$ to 3 N/kg	7600 + 400	–

Table 4

Neural network architectures implemented in this work. The leaky ReLU activation function is used in all MAGNET cases, while ReLU activation is used for CNN cases. For the last layers, the linear activation function is always applied.

Example	Network type	(N. of poolings, N. of MAg/conv. layers per level, window size)	N. of channels per level	N. of parameters
2D L-shape	MAGNET	(3, 2, A^2)	16, 32, 64, 128	~ 4 E+6
	CNN U-Net	(2, 2, 3×3)	64, 128, 512	
3D beam	MAGNET	(5, 1, A^2)	3, 3, 3, 12, 24, 48	~ 75 E+6
	CNN U-Net	(4, 2, $3 \times 3 \times 3$)	256, 256, 256, 512, 512	
2D beam (hole)	MAGNET	(3, 2, A^2)	8, 16, 32, 64	~ 2 E+6
3D breast	MAGNET	(4, 1, A^2)	6, 12, 12, 24, 48	~ 19 E+6

**Fig. 9.** MAGNET architecture used for the 2D L-shape example.

starts with the input mesh (2D), to which the MAg layer is applied twice (with 16 output channels). This is followed by the graph pooling layer, which coarsens the mesh and transitions to the next level of the U-Net (from zeroth to the first level). This process repeats twice, with the first and second levels of the graph U-Net having MAg layers with 32 and 64 output channels, respectively, leading to the coarsest third level of the U-Net. At this level, two MAg layers (with 128 output channels) are applied. In the subsequent decoding phase, the graph unpooling layer is employed with the concurrent concatenation operation, and followed by two MAg operations (with 64 output channels). This upsampling sequence repeats twice with the use of 32- and 16-channel MAg layers. Finally, a single MAg layer (with 2 output channels) is applied, using a linear activation to produce the desired output mesh (the displacement mesh must have the same structure as the input mesh of forces). It is worth noting that analogous architectures of CNN U-Net networks are similar, with the only distinction being the use of convolution layers in place of MAg layers and CNN U-Net max poolings instead of graph poolings.

As demonstrated in Table 4, both 2D L-shape and 3D beam examples have been modeled utilizing both MAGNET and CNN U-Net architectures. These networks were designed to have a similar number of trainable parameters, thus facilitating a fair comparison of their fitting capabilities. The number of parameters in both types of networks is controlled by having a higher number of channels in the CNN U-Net architecture compared to its corresponding MAGNET architecture. This difference in the number of channels is attributed to the convolution

operators in the CNN architecture sharing parameters across a layer, which may necessitate a larger number of channels to ensure an optimal fit, while the aggregation operators in the MAg layer use individual weights per aggregation window, allowing for more flexible fitting across the mesh with a smaller number of channels. However, caution must be exercised when selecting the number of channels, as setting it too low can result in increased prediction errors (as seen in Fig. 15).

The number of neural network levels (pooling operations) and the size of convolution/aggregation windows have been adjusted on a case-by-case basis to obtain the desired fitting capabilities while keeping the number of trainable parameters low and comparable between the respective CNN U-Net and MAGNET models. The fitting capabilities heavily depend on the successful propagation of information from the input throughout the network. This can be compromised when the number of poolings or the window size is too small, as explained in Section 2.5. For this reason, a larger number of pooling operations is used for mesh graphs with larger diameters (e.g., the 3D cases in Table 4). Additionally, in the case of MAGNET models, a global optimization of graph pooling operations is performed to reduce the number of nodes at the coarsest level. In this optimization, Algorithm 1 is run 1000 times with different random seeds, and the case with the least number of nodes at the lowest level is selected.

Remark. Note that the example presented in Fig. 8(a) utilizes a non-structured mesh. As such, it cannot be directly used by the CNN U-Net model, and an additional preprocessing step needs to be done




Fig. 10. Zero-padding is applied to make the L-shape topology compatible with the CNN framework. The additional nodal values for inputs (forces) and outputs (displacements) are set as zero vectors.




Fig. 11. Training loss curve for the 3D breast MAgNET model.

to make the input and output meshes structured. In this case, we apply zero padding to convert the L-shape mesh into a structured mesh, see Fig. 10, which is then used for training with the CNN U-Net architecture. We do not need to do this preprocessing step for the MAgNET architecture.

The models presented in Table 4 were trained by minimizing the loss function, as described in Eq. (4), using the datasets introduced in Section 3.1. The Adam optimizer, an extension of the stochastic gradient descent algorithm, was used for this purpose. A mini-batch size of 4 and an initial learning rate of 1×10^{-4} , with a linear decay to 1×10^{-6} during training, were employed. The number of epochs (i.e., iterations of the Adam optimizer) was manually tailored on a case-by-case basis to achieve low values of the loss function. An example of the training loss is provided in Fig. 11. The network trainings were conducted using TensorFlow on a Tesla V100-SXM2 GPU at the HPC facilities of the University of Luxembourg, see Varrette et al. (2014).

3.3. Cross validation of CNN U-Net and MAgNET predictions

We are going to compare the predictions of MAgNET and CNN U-Net models for two problems with structured inputs/outputs that were introduced in Fig. 8(a,b). Let us look at the individual examples with the highest nodal displacement magnitudes. In the 2D L-shape example, shown in Fig. 12a, MAgNET predictions visually coincide with the reference FEM solution very well. This is quantitatively shown in Fig. 12b and c, where the L_2 error field

$$\text{err}(\mathbf{X}) = \|u_{\text{FEM}}(\mathbf{X}) - u_{\text{pred}}(\mathbf{X})\|_2 \quad (17)$$

is presented for MAgNET and CNN U-Net, respectively, demonstrating low level of errors for both models. A similar tendency can be observed in the 3D beam case shown in Fig. 13. Here, although the level of errors is relatively a bit higher than in the 2D example, the MAgNET and CNN U-Net perform similarly, which proves good capabilities of the proposed MAgNET model as compared to the CNN U-Net model.

Table 5

Error metrics for the structured mesh examples. M_{te} stands for the number of test examples, and \bar{e} , $\sigma(e)$, e_{max} are error metrics defined by Eqs. (20)–(22).

Example	M_{te}	\bar{e} [m]	$\sigma(e)$ [m]	e_{max} [m]
2D L-shape (MAgNET)	200	0.5 E-3	0.2 E-3	1.1 E-2
2D L-shape (CNN U-Net)		0.7 E-3	0.6 E-3	1.8 E-2
3D beam (MAgNET)	1782	0.8 E-3	0.7 E-3	7.7 E-2
3D beam (CNN U-Net)		0.7 E-3	0.5 E-3	5.4 E-2

In the following, we will analyze and compare the performance of both models for all cases in the test datasets. For that purpose, we need aggregated error metrics. As an error metric for a single test example, we use the mean absolute error,

$$e_m = e(\mathbf{f}_m, \mathbf{u}_m) = \frac{1}{F} \sum_{i=1}^F |\mathcal{G}(\mathbf{f}_m)^i - \mathbf{u}_m^i|, \quad (18)$$

where the force–displacement pair $(\mathbf{f}_m, \mathbf{u}_m)$ is an element of the test dataset

$$D_{\text{te}} = \{(\mathbf{f}_{M_{\text{tr}}+1}, \mathbf{u}_{M_{\text{tr}}+1}), \dots, (\mathbf{f}_{M_{\text{tr}}+M_{\text{te}}}, \mathbf{u}_{M_{\text{tr}}+M_{\text{te}}})\}, \quad (19)$$

and F is the number of dofs of the mesh. The metric e_m gives us the notion of error between an expected finite element solution, \mathbf{u}_m , and the prediction of the neural network, $\mathcal{G}(\mathbf{f}_m)$. To analyze the overall quality of fitting, we define a single error metric over the entire test set as the average mean absolute error

$$\bar{e} = \frac{1}{M_{\text{te}}} \sum_{m=M_{\text{tr}}+1}^{M_{\text{tr}}+M_{\text{te}}} e_m, \quad (20)$$

with the corrected sample standard deviation (standard deviation of averaged errors) defined as

$$\sigma(e) = \sqrt{\frac{1}{M_{\text{te}} - 1} \sum_{m=M_{\text{tr}}+1}^{M_{\text{tr}}+M_{\text{te}}} (e_m - \bar{e})^2}. \quad (21)$$

Finally, in addition to that, we also use the maximum error per degree of freedom over the entire test set

$$e_{\text{max}} = \max_{m,i} |\mathcal{G}(\mathbf{f}_m)^i - \mathbf{u}_m^i|. \quad (22)$$

The aggregated error metrics for the entire test sets of structured mesh examples obtained using MAgNET and CNN approaches are summarized in Table 5. The first observation is that both MAgNET and CNN models exhibit similar prediction accuracy, demonstrating that the MAgNET architecture can achieve a comparable predictive capacity to the CNN U-Net architecture for a similar number of trainable parameters. The prediction errors, with respect to a characteristic length of 1 m, fall below 0.1% for the average mean absolute error (Eq. (20)), which is a promising result given the presence of geometric and constitutive nonlinearities. Additionally, we analyze the performance of the MAgNET model as a function of the maximum nodal displacement per test example. This dependency is visualized in Fig. 14 for both benchmark examples. Although there is a general trend of increased errors for larger maximum displacement magnitudes, the sensitivity is low, and the errors remain small (the regression lines are $e(d) \propto$

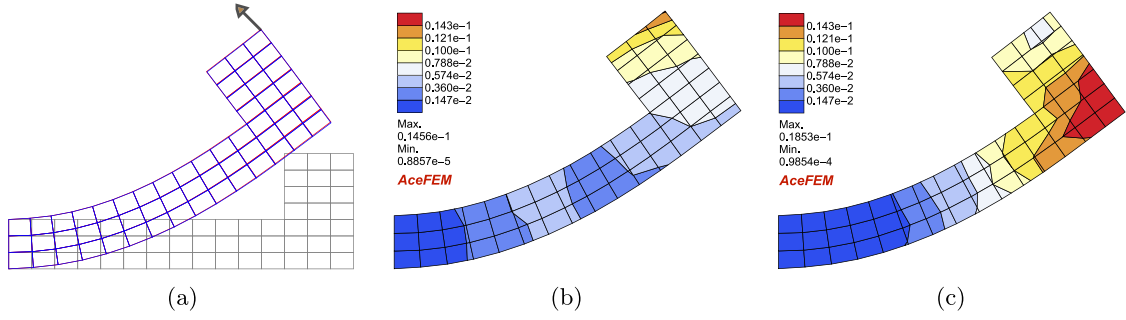


Fig. 12. Deformation of 2D L-shape under point load $(-0.93, 0.91)$ N on the corner node (a) Deformed mesh predicted using MAGNET (blue), for comparison FEM solution is presented (red) (b) L_2 error of nodal displacements between MAGNET and FEM solution. The relative error for the corner node displacement using MAGNET is 0.5% (c) L_2 error of nodal displacements between CNN U-Net and FEM solution. The relative error for the corner node displacement using CNN is 0.3%.

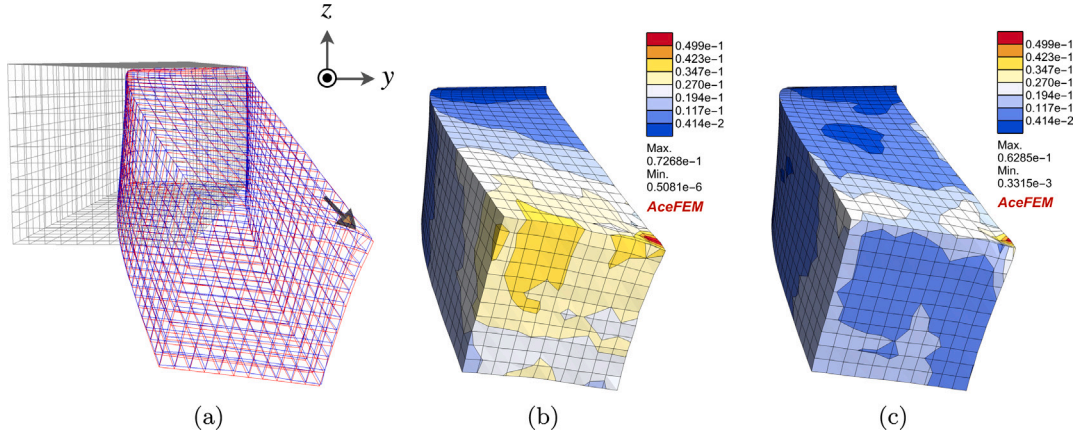


Fig. 13. Deformation of the 3D beam under point load $(-1.75, 1.31, -1.7)$ N on the second last node (a) Deformed mesh predicted using MAGNET (blue), for comparison FEM solution is presented (red) and undeformed mesh is represented by gray (b) L_2 error of nodal displacements between MAGNET and FEM solution. The relative error in predicting displacement of the node of application of load using MAGNET is 4.4% (c) L_2 error of nodal displacements between CNN U-Net and FEM solution. The relative error in predicting displacement of the node of application of load using CNN is 3.0%.

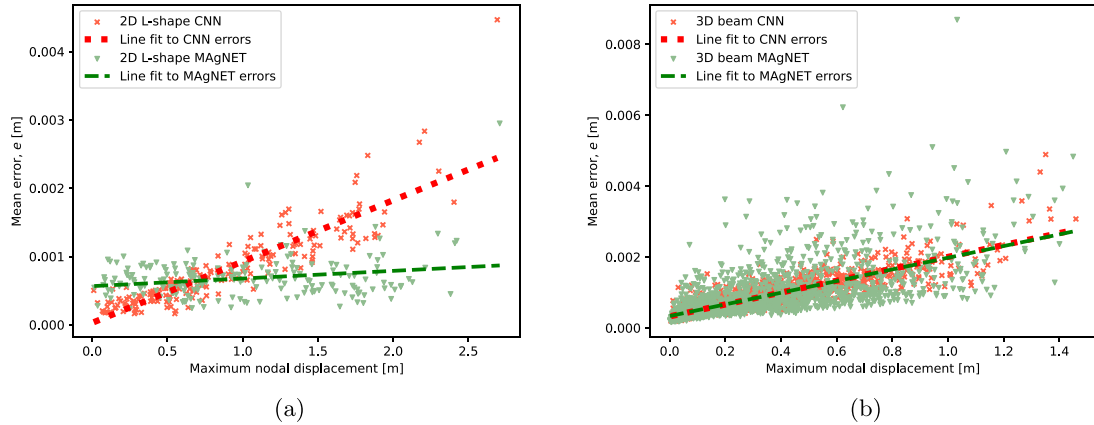


Fig. 14. Mean absolute errors (see Eq. (18)) as a function of maximum nodal displacements for all test examples for 2D L-shape and 3D beam cases for CNN U-NET and MAGNET networks.

$1.0 \cdot d \cdot 10^{-4}$ (2D L-shape) and $e(d) \propto 1.6 \cdot d \cdot 10^{-3}$ (3D beam) for MAGNET and $e(d) \propto 7.0 \cdot d \cdot 10^{-4}$ (2D L-shape) and $e(d) \propto 1.6 \cdot d \cdot 10^{-3}$ (3D beam) for the CNN U-NET case.

The increasing sizes of input/output intensify the challenge in non-convex optimization, necessitating more training data and additional epochs for convergence. This trend is universally observed across all deep learning-based surrogate techniques and relates to the well-known concept of the ‘curse of dimensionality’ in machine learning applications. As discussed in Section 3.2, the choice of hyperparameters for the MAGNET and CNN U-Net architectures has been tailored to specific

mesh structures and sizes, ensuring effective predictive capabilities while maintaining a comparable number of training parameters. In general, finding the most suitable set of hyperparameters would require fine-tuning, which goes beyond the scope of this work. To shed more light on that subject, below, we conduct a simple ablation study to determine the extent to which accuracy declines when simplifying the DNN architectures.

In the first ablation test we reduce the number of channels in MAG and CNN layers. As explained in Section 2.3, the number of channels

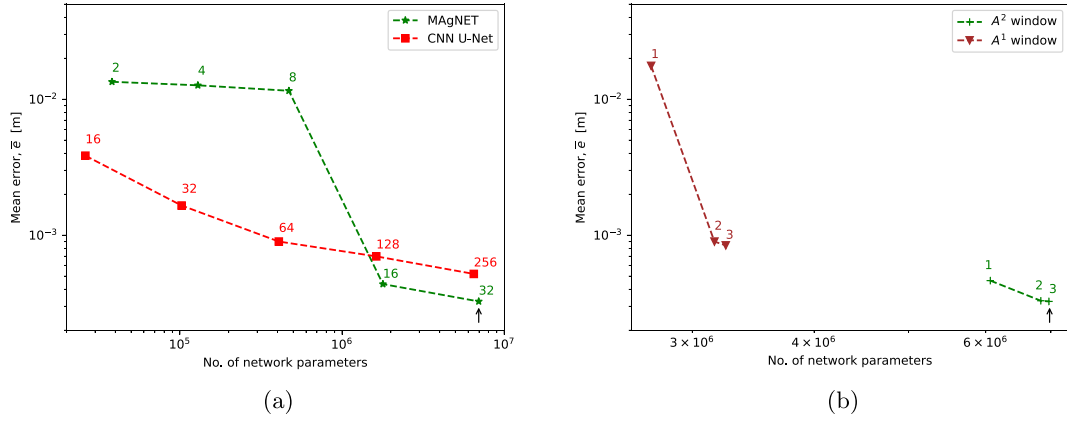


Fig. 15. Average mean error over the test set for the L-shape case for different network architectures. The reference case for the MAGNET is marked with black arrows in both plots, and represents the architecture of A^2 window size, three pooling layers (four levels), and 32 channels. (a) The effect of changing the number of channels used for MAGNET and CNN U-Net architectures (number of channels are marked as numbers in the plot). (b) The effect of reducing the MAG window size and number of pooling layers for MAGNET architecture with 32 channels at each level (number of pooling layers are marked as numbers in the plot).

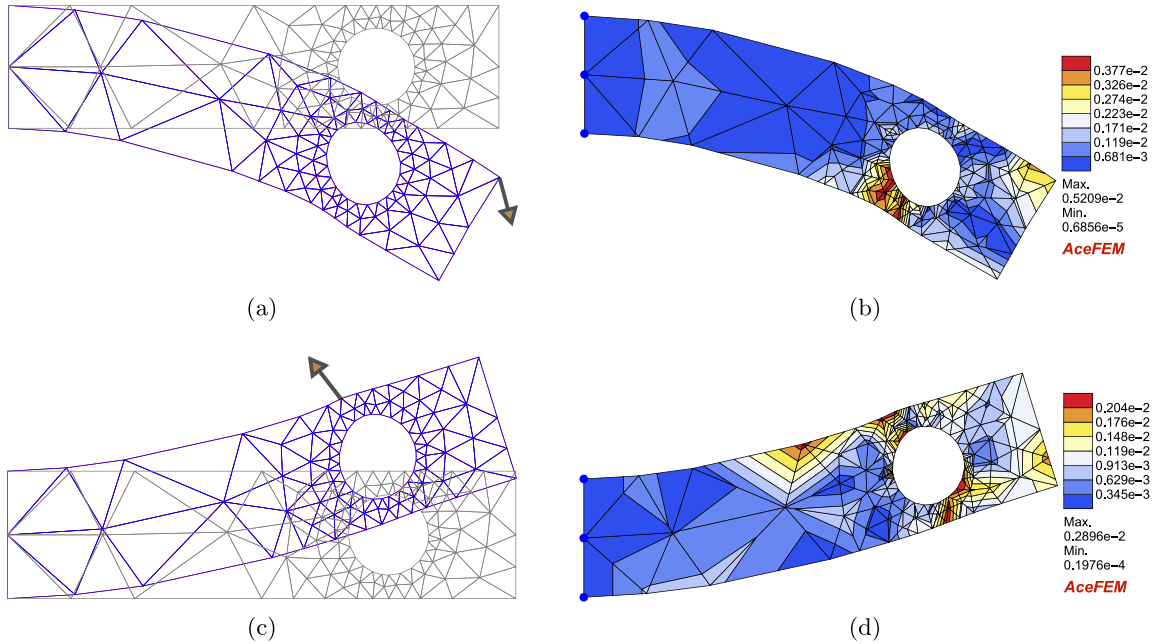


Fig. 16. Deformation of the 2D beam under two different point loads (upper case: (1.28, -4.43) N, lower case: (-3.38, 4.04) N). (a)&(c) Deformed meshes computed using MAGNET (blue) and FEM (red), with the undeformed configuration (gray). (b)&(d) L_2 error of nodal displacements between MAGNET and FEM solutions.

modulates the model capacity to capture non-linearities in the underlying data. Importantly, the convolution windows are shareable in CNN architectures, whereas the aggregation windows in MAG architectures are independent. To this end, we expect CNN networks to require more channels than their respective MAGNET networks to achieve the same level of accuracy. We used this fact when designing CNN and MAGNET architectures in Section 3.2. To verify this hypothesis and demonstrate this effect, we trained five MAGNET and five CNN U-Net networks on the L-shape dataset with different numbers of channels. In all analyzed cases, we used 4-level MAGNET and 3-level CNN U-Net architectures, with two MAG/Conv layers per level, and with a constant number of channels at all levels. Fig. 15(a) shows that there is indeed a strong dependency of accuracy on the number of channels for both analyzed network architectures. For a comparable number of trainable parameters, CNN U-Nets can use more channels than their respective MAGNETs, providing them with comparable predictive accuracy. We can also observe that too few channels significantly reduce the fitting capabilities of both networks, with a step jump between the 8- and 16-channel case for MAGNET. For the two largest cases (16 and 32

channels for MAGNET and 128 and 256 channels for CNN U-Net), the accuracy of both architectures is comparable.

In the second ablation test, we modulate two other important hyper-parameters of MAGNET architectures: the MAG layer window size and the number of pooling layers. Again, we utilize the L-shape dataset, with the reference MAGNET architecture characterized by a window size of A^2 , four levels (three pooling operations), and 32 channels at each level. Fig. 15(b) demonstrates a general trend where the model's simplification leads to a drop in accuracy. This is observable both when reducing the window size from A^2 to A^1 and when decreasing the number of levels (pooling operations). This reduced capacity of the simplified models can be explained by the overall reduced number of parameters, but also by the less intensive exchange of information across the graph. In particular, a significant drop in accuracy is observed for the single-pooling A^1 case. This can be attributed to the inability of this oversimplified network to fully exchange the nodal input information, confirming the effect discussed in more detail in Section 2.5.

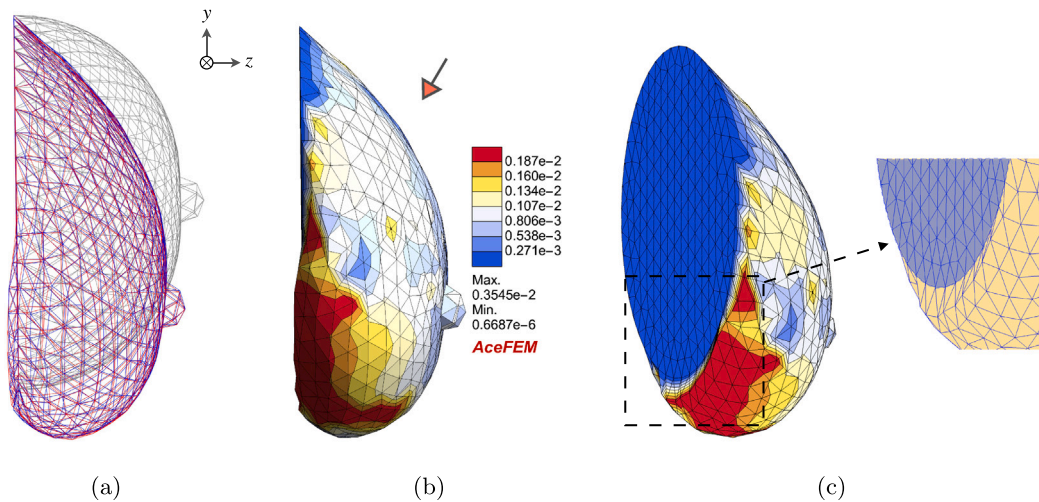


Fig. 17. Deformation of the 3D breast geometry with force density of $(-5.94, -5.23, -2.56)$ N/kg. (a) Deformed meshes computed using MAGNET (blue) and FEM (red), with the undeformed configuration (gray). (b) L_2 error of nodal displacements between MAGNET and FEM solutions. (c) Titled view of the figure (b), MAGNET efficiently captures fixed boundary and nearby high non-linear deformations by learning implicitly from the data.

Remark. The high accuracy of the proposed surrogate models is accompanied by very rapid prediction times compared to the respective high-fidelity FEM predictions. For example, in the 3D beam case, the nonlinear FEM solution required over 3 seconds for some test examples, while MAGNET could make predictions in just 0.18 seconds. Moreover, while MAGNET maintains consistent prediction times regardless of the input force, the time for the FEM solution increases with larger load cases. This increase is due to the use of an iterative solver and an adaptive load-stepping scheme, which are necessary to address convergence issues in highly nonlinear problems. A more comprehensive study on how MAGNET's computational efficiency compares with nonlinear FEM, as well as with CNN-based and attention-based neural networks, is provided in Deshpande et al. (2023b).

3.4. Predictions of MAGNET for general (unstructured) meshes

In Section 3.3, we demonstrated that the MAGNET architectures can achieve very good predictive capabilities for structured mesh cases, which was also cross-validated against respective CNN U-Net architectures. In this section, we aim to show that the high prediction accuracy of MAGNET can also be expected for unstructured mesh cases, which is the central point of the results section. We consider two cases: the first one is deformation under the application of point loads (the 2D beam with hole case), similar to the case of structured examples, and the second is deformation under body forces (the 3D breast case, inspired from Lavigne et al. (2023)).

Let us first analyze individual examples. In Fig. 16 we present two particular loading cases of the 2D beam with hole. One of them is loaded at the tip, featuring the highest nodal displacement magnitude of all test cases, and the other one is loaded close to the hole, representing high local distortions. Similarly, for the three-dimensional problem, in Fig. 17 we show the case featuring the highest nodal displacement magnitude of all test cases. In all mentioned examples we can observe overall good accuracy when visually comparing MAGNET predictions with the respective FEM solutions. This can also be checked quantitatively by analyzing maximum displacement errors. In the cases of the 2D beam with hole, those errors are 1.4% and below when related to the characteristic length of 1 m. In the case of 3D breast geometry, such relative maximum error is higher, reaching almost 3.1% (related to the breast diameter of 0.16 m). Despite this fact, we can observe that high local shape distortions are very well recovered. This property is more emphasized in Fig. 17c where one can additionally observe that also the Dirichlet boundary conditions are very well predicted, even though they were only introduced implicitly by training data.

Table 6

Error metrics for the unstructured mesh examples. M_{te} stands for the number of test examples, and \bar{e} , $\sigma(e)$, e_{max} are error metrics defined in Section 3.3.

Example	M_{te}	\bar{e} [m]	$\sigma(e)$ [m]	e_{max} [m]
2D beam (hole)	240	0.7 E-3	0.4 E-3	1.4 E-2
3D breast	400	8.9 E-5	3.1 E-5	5.1 E-3

The aggregated error metrics for the entire test sets are provided in Table 6. The maximum displacement errors over all test cases, e_{max} , are at the levels observed for particular cases in Figs. 16 and 17. At the same time, the average mean errors, \bar{e} , are at least an order of magnitude lower, which suggests that the errors close to maximum levels are not that often. The average mean errors are further analyzed in a case-by-case manner in Fig. 18, which is analogous to the analysis done for the structured cases in Fig. 14. Again, we plot the mean error e , of each test example as a function of the maximum nodal displacement. The regression lines $e(d) \propto 5.0 \cdot d \cdot 10^{-4}$ (2D-beam) and $e(d) \propto 8.0 \cdot d \cdot 10^{-4}$ (3D Breast) show low sensitivity of the MAGNET predictions to displacement magnitudes.

Above, we have demonstrated a good prediction accuracy of MAGNET within the test dataset (which is located in the interpolated domain). However, it is well known that this accuracy can gradually deteriorate when moving to the extrapolated region, see, e.g., Deshpande et al. (2022). We are going to study this effect for MAGNET for a particular case that is based on the 3D breast geometry. As described in the Table 3, during the training, the b_z component of body force density is varied from -3 to 3 N/kg only. At the inference time, we applied b_z from -7 to 7 N/kg (keeping other components 0) to see how the predictions perform within and outside the training magnitudes. Fig. 19(a) shows that the error is fairly low and is not increasing within the training region and it increases rapidly outside, which confirms this well-known effect. Figs. 19(b) and 19(c) show deformed meshes predicted for $b_z = 5$ and $b_z = 9$ N/kg, respectively, both outside the training data region. MAGNET is observed to give visually acceptable results although the accuracy of the framework decreases as we move away from the training data.

3.4.1. A note on physics-informed errors

The proposed MAGNET framework has only been trained by minimizing the loss function for displacement errors, with no additional explicit information about the underlying physics/mechanics. As demonstrated earlier in this work, such training can provide very good accuracy in terms of predicted displacements. However, this accuracy is not

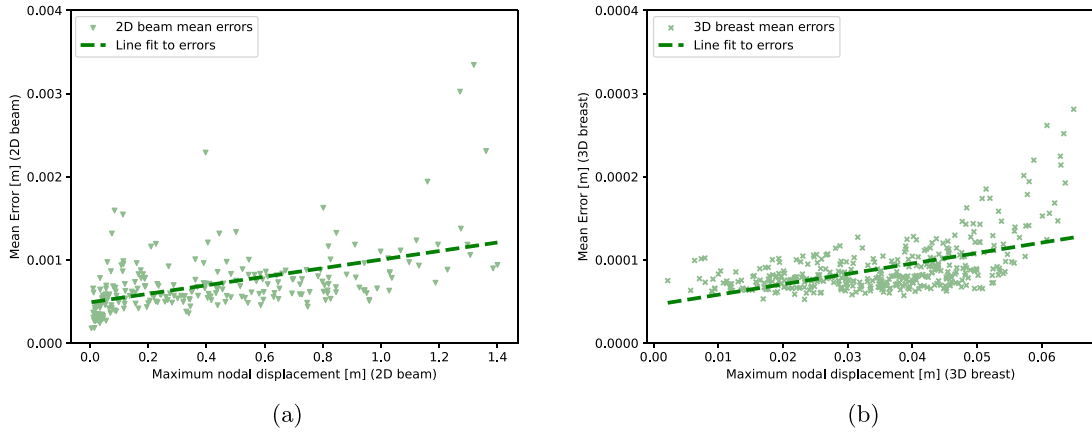


Fig. 18. Mean absolute errors (see Eq. (18)) as a function of maximum nodal displacements for all test examples (with unstructured meshes) predicted using MAgNET for (a) 2D beam with hole (b) 3D breast case.

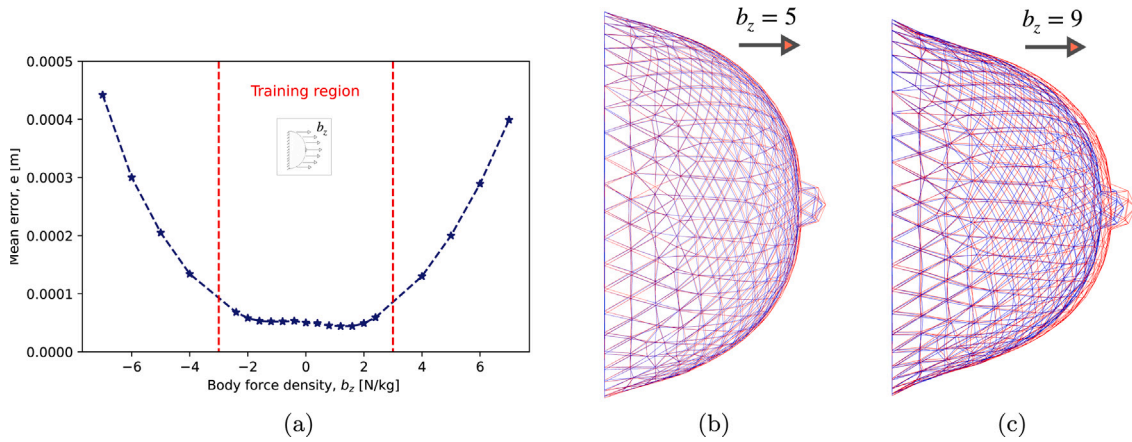


Fig. 19. 3D Breast deformation under horizontal body force densities, $(0, 0, b_z)$ N/kg. (a) Mean absolute error for testing cases in interpolated and extrapolated regions. The error increases rapidly in the extrapolated region while it remains low in the training (interpolated) region. (b)&(c) Visualization of deformed meshes for force densities outside the training region computed using MAgNET (blue) and FEM solution (red).

of machine precision. To this end, a natural question arises: how far the displacement errors can violate physics? To answer that question, we are going to analyze some problem-based quantities of interest, such as residuals (balance of forces) or stresses, in comparison to the expected ground-truth results.

In Fig. 20 we show nodal internal residual forces for the 2D beam with hole cases that we introduced earlier (compare Fig. 16 for respective displacement errors). Ideally, the residual forces should be zero (the balance of forces), except for the boundary condition areas in which they should be exactly opposite to the reaction at the support and the applied external force. However, due to inaccuracies in displacements obtained from the MAgNET model, differences with respect to the ground true residuals can be noted. In Fig. 20, we can observe the expected high residual forces in the areas where Dirichlet and Neumann boundary conditions are applied, however, also localized residual force spots are present in the fine mesh region around the hole. The magnitude of those errors in the localized spots can go up to 20% of the maximal magnitude of applied forces. Also, when more closely analyzing the residuals at boundary condition areas, it turns out that they do not fully match the respective FEM residuals. For instance, the relative error in residuals at the support in Fig. 20a is almost 5%. When analyzing the entire test set, we observed that the mean error in retrieving residuals at the Dirichlet boundary related to the maximum external force is 2.4%. A similar relative mean error value for retrieving the Neumann boundary residual is 14.6%. The higher error for Neumann boundary residual is attributed to high local non-linear

deformations at the vicinity of the point of application of force. Though the displacement errors provoked by these non-linearities are not so high, they can result in high residual errors through the relatively high magnitude of element stiffness.

The errors observed in Figs. 16 and 20 can have a direct impact on some application-dependent quantities of interest. As an example, in Fig. 21, we present the field of von Mises stresses, which is a commonly used measure of shear stresses. We can observe that the MAgNET solution provides similar profiles of stresses as compared to respective FEM solutions, however, high localized errors are present at the fine mesh region (up to 30% of the reference FEM maximal von Mises stresses).

The above mentioned localized errors in residual forces, von Mises stresses and other relevant quantities of interest can be reduced by enriching the loss function with physics-informed terms. For instance, in the context of mesh-based force–displacement data, in Odot et al. (2021) such enrichment has been introduced by scaling individual components of the loss function with the respective computed residual values, which proved to reduce residual errors. In As'ad et al. (2022), the authors introduced an energy-based approach that provided purely physics-informed training for the Gauss-point stress–strain relationship, which allowed them to satisfy the expected frame indifference. Similar concepts of physics-informed loss functions can be seamlessly integrated into the MAgNET framework, which would convert it into a Physics Informed MAgNET.

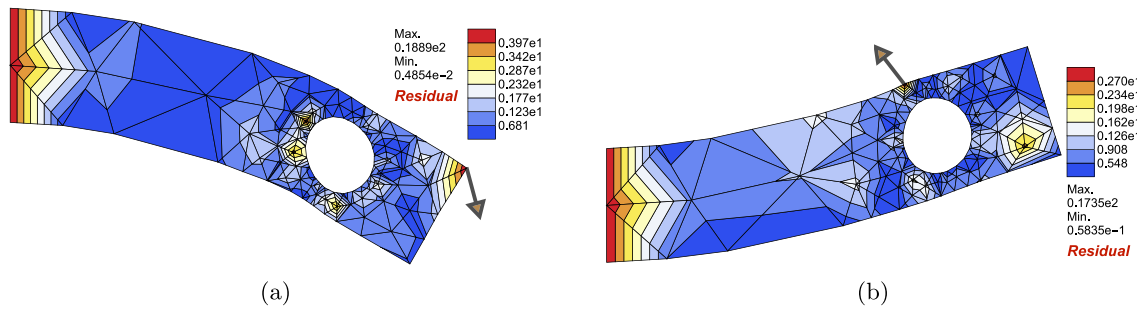


Fig. 20. Nodal residual forces obtained using MAGNET solutions for the examples in Fig. 16 (plotted on deformed meshes). The relative error for retrieving the total reaction force at the fixed interface is (a) 4.7% for the first example (b) 0.1% for the second example.

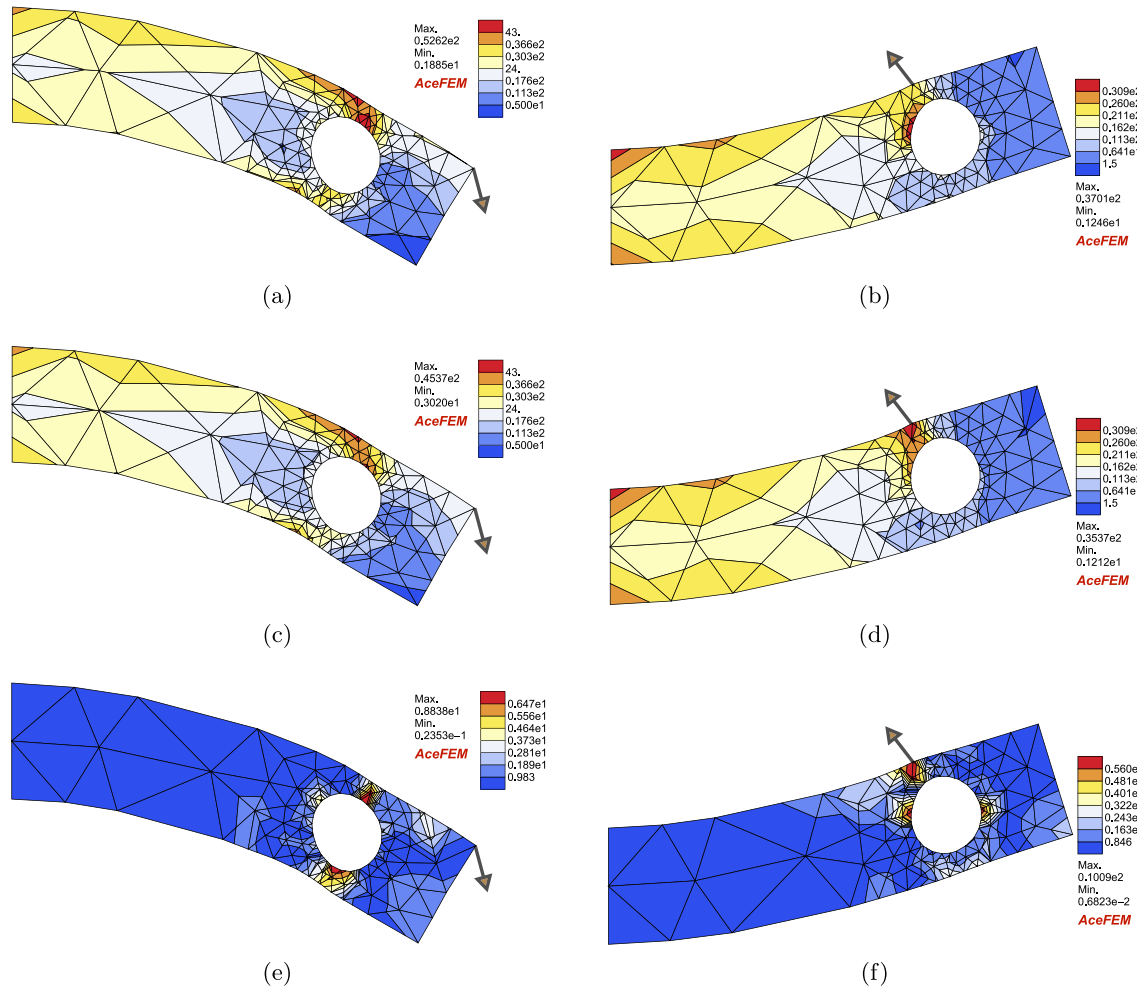


Fig. 21. Von Mises stresses obtained for the two examples as in Fig. 16 using (a)&(b) MAGNET solution (c)&(d) FEM solution. In (e)&(f) the absolute error between the MAGNET and FEM von Mises stresses is shown.

4. Conclusion and future directions

In this work we proposed MAGNET, a novel framework for efficient supervised learning on graph-structured data using geometric deep learning. The framework comprises two neural network operations: MAg and graph pooling/unpooling layers, which together form a graph U-Net architecture capable of learning on large-dimension inputs/outputs. Notably, the MAGNET framework is not restricted to any particular input→output relationship or any specific mesh- or discretization scheme, making it superior to existing convolutional neural network architectures. MAGNET allows for arbitrary non-grid inputs/outputs, meaning it can handle arbitrary meshes and support

complex geometries and local mesh refinements, making it suitable for a wide range of engineering applications.

We demonstrated and studied the capabilities of MAGNET in capturing nonlinear relationships in data. In particular, we showed that MAGNET can serve as an efficient surrogate framework for non-linear FEM simulations. For this purpose, we conducted quantitative cross-validation of predictions made by MAGNET and the well-known convolutional U-Net architecture, both of which have been verified against the ground-truth results obtained with FEM. The benchmarks have proven that MAGNET has similar predictive capabilities as CNN U-Net for structured meshes, and it can also be extended to arbitrary meshes while preserving similar accuracy of predictions.

There are multiple avenues for enhancing MAGNET's capabilities and applying it beyond the scope of the present work. One promising direction involves directly integrating the underlying physics into the training process. While the current purely data-based MAGNET accurately captures the displacement field, it exhibits increased errors for other physics-related quantities, as discussed in Section 3.4.1. Incorporating physics-based components into the loss function could significantly improve the performance of MAGNET in these areas. It could also result in reduction in the amount of necessary training data. A second potential research direction is applying MAGNET to unsteady problems. Such techniques already exist, for instance, in applications to dynamics (Pfaff et al., 2021; Meister et al., 2020) and in the context of elasto-plasticity (Mozaffar et al., 2019; Vijayaraghavan et al., 2021). We believe some of them could be seamlessly integrated with the MAGNET architecture. A third future research direction is to enhance the prediction accuracy of surrogate models for real-world applications involving online observational data. One approach is to combine any type of DNNs, such as MAGNET, with data assimilation techniques, like Kalman filtering (Liu et al., 2023). A fourth potential direction is to transform MAGNET into a probabilistic model. This could be achieved by performing local aggregations with probability distributions instead of discrete weights, similar to our approach with Bayesian CNNs in previous work (Deshpande et al., 2022). A Bayesian MAGNET would be capable of tracking uncertainties inherent in both the network architecture and real-world data. Finally, we envision applying MAGNET not only to forward problems but also to inverse problems. The MAGNET architecture's ability to handle non-structured meshes could be combined with existing ideas for data-driven calibration of model parameters, as seen, for example, in Raissi et al. (2019), allowing for a DNN extension of our earlier calibration schemes (Lavigne et al., 2023) to a broader class of systems and models.

We have made all the codes, datasets, and examples presented in this paper available open-access in the MAGNET repository at <https://github.com/saurabhdeshpande93/MAGNET>. Given the generality of MAGNET in supporting arbitrary non-linear relationships and arbitrary discretizations, we believe that the repository will provide a useful surrogate modeling framework for researchers and practitioners in various application areas across disciplines. We see it not only as a ready-to-use machine-learning library but also as a reference point and foundation for future developments and extensions in this emerging direction of research. The generality of MAGNET will enable the community to explore a range of new applications and modeling scenarios.

CRedit authorship contribution statement

Saurabh Deshpande: Conceptualization, Data curation, Formal analysis, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Stéphane P.A. Bordas:** Funding acquisition, Investigation, Methodology, Supervision, Writing – review & editing. **Jakub Lengiewicz:** Conceptualization, Data curation, Investigation, Methodology, Project administration, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

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.

Data availability

The source codes and supplementary dataset utilized in this study can be found at <https://github.com/saurabhdeshpande93/MAGNET> and <https://doi.org/10.5281/zenodo.7784804>.

Acknowledgments



This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 764644. Jakub Lengiewicz would like to acknowledge the support from

EU Horizon 2020 Marie Skłodowska Curie Individual Fellowship *MOR-PhEM* under Grant 800150. Stéphane Bordas and Jakub Lengiewicz are grateful for the support of the Fonds National de la Recherche Luxembourg FNR grant QuaC C20/MS/14782078. Stéphane Bordas received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 811099 TWINNING Project DRIVEN for the University of Luxembourg. This paper only contains the author's views and the Research Executive Agency and the Commission are not responsible for any use that may be made of the information it contains.

References

- As'ad, F., Avery, P., Farhat, C., 2022. A mechanics-informed artificial neural network approach in data-driven constitutive modeling. <http://dx.doi.org/10.2514/6.2022-0100>.
- Aydin, R.C., Braeu, F.A., Cyron, C.J., 2019. General multi-fidelity framework for training artificial neural networks with computational models. *Front. Mater.* 6, 61.
- Bianchi, F.M., Grattarola, D., Alippi, C., 2019. Spectral clustering with graph neural networks for graph pooling. <http://dx.doi.org/10.48550/ARXIV.1907.00481>, URL: <https://arxiv.org/abs/1907.00481>.
- Black, N., Najafi, A.R., 2022. Learning finite element convergence with the multi-fidelity graph neural network. *Comput. Methods Appl. Mech. Engrg.* 397, 115120, URL: <https://www.sciencedirect.com/science/article/pii/S004578252200305X>.
- Brenner, S.C., Scott, L.R., 2008. Finite element multigrid methods. In: *The Mathematical Theory of Finite Element Methods*. Springer New York, New York, NY, pp. 155–173. http://dx.doi.org/10.1007/978-0-387-75934-0_7.
- Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P., 2017. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Process. Mag.* 34 (4), 18–42. <http://dx.doi.org/10.1109/MSP.2017.2693418>.
- Brunet, J.-N., Mendizabal, A., Petit, A., Golse, N., Vibert, E., Cotin, S., 2019. Physics-based deep neural network for augmented reality during liver surgery. In: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.-T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. Springer International Publishing, Cham, pp. 137–145.
- Bui, H.P., Tomar, S., Courtecuisse, H., Cotin, S., Bordas, S.P.A., 2018. Real-time error control for surgical simulation. *IEEE Trans. Biomed. Eng.* 65 (3), 596–607. <http://dx.doi.org/10.1109/TBME.2017.2695587>.
- Cai, C., Wang, D., Wang, Y., 2021. Graph coarsening with neural networks. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=uxpzitPEooJ>.
- Chen, M., Wei, Z., Huang, Z., Ding, B., Li, Y., 2020. Simple and deep graph convolutional networks. In: Daumé, H., Singh, A. (Eds.), *Proceedings of the 37th International Conference on Machine Learning*. In: *Proceedings of Machine Learning Research*, vol. 119, PMLR, pp. 1725–1735, URL: <https://proceedings.mlr.press/v119/chen20v.html>.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3D U-net: learning dense volumetric segmentation from sparse annotation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 424–432.
- Daniel, T., Casenave, F., Akkari, N., Ryckelynck, D., 2020. Model order reduction assisted by deep neural networks (ROM-net). *Adv. Model. Simul. Eng. Sci.* 7 (1), 1–27.
- Deshpande, S., Bordas, S., Lengiewicz, J., 2023a. Magnet: A graph U-Net architecture for mesh-based simulations [dataset]. <http://dx.doi.org/10.5281/zenodo.000000>.
- Deshpande, S., Lengiewicz, J., Bordas, S.P., 2022. Probabilistic deep learning for real-time large deformation simulations. *Comput. Methods Appl. Mech. Engrg.* 398, 115307. <http://dx.doi.org/10.1016/j.cma.2022.115307>, URL: <https://www.sciencedirect.com/science/article/pii/S004578252200411X>.
- Deshpande, S., Sosa, R.I., Bordas, S.P.A., Lengiewicz, J., 2023b. Convolution, aggregation and attention based deep neural networks for accelerating simulations in mechanics. *Front. Mater.* 10, <http://dx.doi.org/10.3389/fmats.2023.1128954>.
- Doll, S., Schweizerhof, K., 2000. On the development of volumetric strain energy functions. *J. Appl. Mech.* 67 (1), 17–21.

- Fernández-León, J., Keramati, K., Miguel, C., González, C., Baumela, L., 2023. A deep encoder-decoder for surrogate modelling of liquid moulding of composites. *Eng. Appl. Artif. Intell.* 120, 105945. <http://dx.doi.org/10.1016/j.engappai.2023.105945>, URL: <https://www.sciencedirect.com/science/article/pii/S095219762300129X>.
- Gao, H., Ji, S., 2019. Graph U-net. URL: <https://openreview.net/forum?id=HJePrRoAct7>.
- Gao, H., Sun, L., Wang, J.-X., 2021. PhyGeoNet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state PDEs on irregular domain. *J. Comput. Phys.* 428, 110079, URL: <https://www.sciencedirect.com/science/article/pii/S0021999120308536>.
- Gao, H., Zahr, M.J., Wang, J.-X., 2022. Physics-informed graph neural Galerkin networks: A unified framework for solving PDE-governed forward and inverse problems. *Comput. Methods Appl. Mech. Engrg.* 390, 114502. <http://dx.doi.org/10.1016/j.cma.2021.114502>, URL: <https://www.sciencedirect.com/science/article/pii/S0045782521007076>.
- Goury, O., Duriez, C., 2018. Fast, generic, and reliable control and simulation of soft robots using model order reduction. *IEEE Trans. Robot.* 34 (6), 1565–1576.
- Henkes, A., Wessels, H., Mahnken, R., 2022. Physics informed neural networks for continuum micromechanics. *Comput. Methods Appl. Mech. Engrg.* 393, 114790. <http://dx.doi.org/10.1016/j.cma.2022.114790>, URL: <https://www.sciencedirect.com/science/article/pii/S0045782522001268>.
- Hennequin, R., Khelif, A., Voituret, F., Moussallam, M., 2020. Spleeter: a fast and efficient music source separation tool with pre-trained models. *J. Open Source Softw.* 5, 2154. <http://dx.doi.org/10.21105/joss.02154>.
- Hoq, E., Aljarrah, O., Li, J., Bi, J., Heryudono, A., Huang, W., 2023. Data-driven methods for stress field predictions in random heterogeneous materials. *Eng. Appl. Artif. Intell.* 123, 106267. <http://dx.doi.org/10.1016/j.engappai.2023.106267>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197623004517>.
- Huang, D.Z., Xu, K., Farhat, C., Darve, E., 2020. Learning constitutive relations from indirect observations using deep neural networks. *J. Comput. Phys.* 416, 109491. <http://dx.doi.org/10.1016/j.jcp.2020.109491>, URL: <https://www.sciencedirect.com/science/article/pii/S0021999120302655>.
- Jiang, C., Chen, N.-Z., 2023. Graph neural networks (GNNs) based accelerated numerical simulation. *Eng. Appl. Artif. Intell.* 123, 106370. <http://dx.doi.org/10.1016/j.engappai.2023.106370>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197623005547>.
- Johnsen, S.F., Taylor, Z.A., Clarkson, M.J., Hipwell, J., Modat, M., Eiben, B., Han, L., Hu, Y., Mertzaniidou, T., Hawkes, D.J., et al., 2015. NiftySim: A GPU-based nonlinear finite element package for simulation of soft tissue biomechanics. *Int. J. Comput. Assist. Radiol. Surg.* 10, 1077–1095.
- Kipf, T.N., Welling, M., 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P., Hoyer, S., 2021. Machine learning-accelerated computational fluid dynamics. *Proc. Natl. Acad. Sci.* 118 (21), e2101784118. <http://dx.doi.org/10.1073/pnas.2101784118>, URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2101784118>. *arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2101784118*.
- Korelc, J., 2002. Multi-language and multi-environment generation of nonlinear finite element codes. *Eng. Comput.* 18, 312–327. <http://dx.doi.org/10.1007/s003660200028>.
- Krokos, V., Bordas, S.P., Kerfriden, P., 2024. A graph-based probabilistic geometric deep learning framework with online enforcement of physical constraints to predict the criticality of defects in porous materials. *Int. J. Solids Struct.* 286–287, 112545. <http://dx.doi.org/10.1016/j.jisolsstr.2023.112545>, URL: <https://www.sciencedirect.com/science/article/pii/S0020768323004420>.
- Krokos, V., Bui Xuan, V., Bordas, S.P.A., Young, P., Kerfriden, P., 2022. A Bayesian multiscale CNN framework to predict local stress fields in structures with microscale features. *Comput. Mech.* 69, 733–766. <http://dx.doi.org/10.1007/s00466-021-02112-3>.
- Lavigne, T., Bordas, S., Lengiewicz, J., 2023. Identification of material parameters and traction field for soft bodies in contact. *Comput. Methods Appl. Mech. Engrg.* 406, 115889. <http://dx.doi.org/10.1016/j.cma.2023.115889>, URL: <https://www.sciencedirect.com/science/article/pii/S0045782523000129>.
- Lavigne, T., Mazier, A., Perney, A., Bordas, S., Hild, F., Lengiewicz, J., 2022. Digital volume correlation for large deformations of soft tissues: Pipeline and proof of concept for the application to breast ex vivo deformations. *J. Mech. Behav. Biomed. Mater.* 105490. <http://dx.doi.org/10.1016/j.jmbmb.2022.105490>, URL: <https://www.sciencedirect.com/science/article/pii/S1751616122003952>.
- Le, T.-H., Kang, H., Kim, H., 2022. Towards incompressible laminar flow estimation based on interpolated feature generation and deep learning. *Sustainability* 14 (19), <http://dx.doi.org/10.3390/su141911996>, URL: <https://www.mdpi.com/2071-1050/14/19/11996>.
- Lee, J., Lee, I., Kang, J., 2019. Self-attention graph pooling. <http://dx.doi.org/10.48550/ARXIV.1904.08082>, URL: <https://arxiv.org/abs/1904.08082>.
- Liu, W., Lai, Z., Bacsa, K., Chatzi, E., 2023. Neural extended Kalman filters for learning and predicting dynamics of structural systems. *Struct. Health Monit.* 14759217231179912. <http://dx.doi.org/10.1177/14759217231179912>, *arXiv:2210.04165*.
- Lorente, D., Martínez-Martínez, F., Rupérez, M., Lago, M., Martínez-Sober, M., Escandell-Montero, P., Martínez-Martínez, J., Martínez-Sanchis, S., Serrano-López, A., Monserrat, C., Martín-Guerrero, J., 2017. A framework for modelling the biomechanical behaviour of the human liver during breathing in real time using machine learning. *Expert Syst. Appl.* 71, 342–357. <http://dx.doi.org/10.1016/j.eswa.2016.11.037>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197617416306728>.
- Luzhnica, E., Day, B., Lio, P., 2019. Clique pooling for graph classification. <http://dx.doi.org/10.48550/ARXIV.1904.00374>, URL: <https://arxiv.org/abs/1904.00374>.
- Meister, F., Passerini, T., Mihalef, V., Tuysuzoglu, A., Maier, A., Mansi, T., 2020. Deep learning acceleration of total Lagrangian explicit dynamics for soft tissue mechanics. *Comput. Methods Appl. Mech. Engrg.* 358, 112628. <http://dx.doi.org/10.1016/j.cma.2019.112628>, URL: <https://www.sciencedirect.com/science/article/pii/S0045782519305109>.
- Mendizabal, A., Márquez-Neila, P., Cotin, S., 2019. Simulation of hyperelastic materials in real-time using deep learning. *Med. Image Anal.* 59, 101569. <http://dx.doi.org/10.1016/j.media.2019.101569>.
- Mozaffar, M., Bostanabad, R., Chen, W., Ehmman, K., Cao, J., Bessa, M.A., 2019. Deep learning predicts path-dependent plasticity. *Proc. Natl. Acad. Sci.* 116 (52), 26414–26420. <http://dx.doi.org/10.1073/pnas.1911815116>.
- Nguyen, T.N.K., Dairay, T., Meunier, R., Mougeot, M., 2022. Physics-informed neural networks for non-Newtonian fluid thermo-mechanical problems: An application to rubber calendaring process. *Eng. Appl. Artif. Intell.* 114, 105176. <http://dx.doi.org/10.1016/j.engappai.2022.105176>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197622002810>.
- Nikolopoulos, S., Kalogeris, I., Papadopoulos, V., 2022. Non-intrusive surrogate modeling for parametrized time-dependent partial differential equations using convolutional autoencoders. *Eng. Appl. Artif. Intell.* 109, 104652. <http://dx.doi.org/10.1016/j.engappai.2021.104652>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197621004541>.
- Obiols-Sales, O., Vishnu, A., Malaya, N., Chandramowlishwaran, A., 2020. CFNet: A Deep Learning-Based Accelerator for Fluid Simulations. *ICS '20*, Association for Computing Machinery, New York, NY, USA, <http://dx.doi.org/10.1145/3392717.3392772>.
- Odot, A., Haferssas, R., Cotin, S., 2021. DeepPhysics: a physics aware deep learning framework for real-time simulation. <http://dx.doi.org/10.48550/ARXIV.2109.09491>, URL: <https://arxiv.org/abs/2109.09491>.
- Ogden, R.W., 2005. *Non-Linear Elastic Deformations*. Dover Publications.
- Pant, P., Doshi, R., Bahl, P., Farimani, A.B., 2021. Deep learning for reduced order modelling and efficient temporal evolution of fluid simulations. *Phys. Fluids* 33 (10), 107101. <http://dx.doi.org/10.1063/5.0062546>.
- Pellicer-Valero, O.J., Rupérez, M.J., Martínez-Sanchis, S., Martín-Guerrero, J.D., 2020. Real-time biomechanical modeling of the liver using machine learning models trained on finite element method simulations. *Expert Syst. Appl.* 143, 113083. <http://dx.doi.org/10.1016/j.eswa.2019.113083>, URL: <https://www.sciencedirect.com/science/article/pii/S095219762109308000>.
- Pfaff, T., Fortunato, M., Gonzalez, A., Battaglia, P., 2021. Learning mesh-based simulation with graph networks. In: *International Conference on Learning Representations*. URL: https://openreview.net/forum?id=roNqYL0_XP.
- Raissi, M., Perdikaris, P., Karniadakis, G., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686–707. <http://dx.doi.org/10.1016/j.jcp.2018.10.045>, URL: <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- Rao, C., Liu, Y., 2020. Three-dimensional convolutional neural network (3D-CNN) for heterogeneous material homogenization. *Comput. Mater. Sci.* 184, 109850. <http://dx.doi.org/10.1016/j.commatsci.2020.109850>.
- Ren, X., Zhang, X., Chen, L., Zheng, X., Zhang, C., Guo, L., Yu, B., 2021. A causal U-net based neural beamforming network for real-time multi-channel speech enhancement. pp. 1832–1836. <http://dx.doi.org/10.21437/Interspeech.2021-1457>.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention. MICCAI*, In: LNCS, vol. 9351, Springer, pp. 234–241, URL: <http://lmb.informatik.uni-freiburg.de/Publications/2015/RFB15a>. (available on *arXiv:1505.04597 [cs.CV]*).
- Roy, A.M., Guha, S., 2023. A data-driven physics-constrained deep learning computational framework for solving von mises plasticity. *Eng. Appl. Artif. Intell.* 122, 106049. <http://dx.doi.org/10.1016/j.engappai.2023.106049>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197623002336>.
- Runge, G., Wiese, M., Raatz, A., 2017. FEM-based training of artificial neural networks for modular soft robots. In: *2017 IEEE International Conference on Robotics and Biomimetics. ROBIO*, pp. 385–392. <http://dx.doi.org/10.1109/ROBIO.2017.8324448>.
- Rus, D., Tolley, M.T., 2015. Design, fabrication and control of soft robots. *Nature* 521 (7553), 467–475.
- Samaniego, E., Anitescu, C., Goswami, S., Nguyen-Thanh, V., Guo, H., Hamdia, K., Zhuang, X., Rabczuk, T., 2020. An energy approach to the solution of partial differential equations in computational mechanics via machine learning: Concepts, implementation and applications. *Comput. Methods Appl. Mech. Engrg.* 362, 112790. <http://dx.doi.org/10.1016/j.cma.2019.112790>, URL: <https://www.sciencedirect.com/science/article/pii/S0045782519306826>.

- Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., Battaglia, P.W., 2020. Learning to simulate complex physics with graph networks. In: *Learning to Simulate Complex Physics with Graph Networks*. pp. 8459–8468, URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85094802982&partnerID=40&md5=57a515ee79e6915a2266fa3f2bc4870c>. Cited by: 53.
- Seo, M., Min, S., 2023. Graph neural networks and implicit neural representation for near-optimal topology prediction over irregular design domains. *Eng. Appl. Artif. Intell.* 123, 106284. <http://dx.doi.org/10.1016/j.engappai.2023.106284>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197623004682>.
- Shivadiya, M.V., Alves, J., Bugiotti, F., Magoules, F., 2022. Graph neural network-based surrogate models for finite element analysis. *arXiv:2211.09373*.
- Simo, J., Taylor, R., 1982. Penalty function formulations for incompressible nonlinear elastostatics. *Comput. Methods Appl. Mech. Engrg.* 35 (1), 107–118.
- Thakolkaran, P., Joshi, A., Zheng, Y., Flaschel, M., De Lorenzis, L., Kumar, S., 2022. NN-EUCLID: Deep-learning hyperelasticity without stress data. *J. Mech. Phys. Solids* 169, 105076. <http://dx.doi.org/10.1016/j.jmps.2022.105076>, URL: <https://www.sciencedirect.com/science/article/pii/S0022509622002538>.
- Varrette, S., Bouvry, P., Cartiaux, H., Georgatos, F., 2014. Management of an academic HPC cluster: The UL experience. <http://dx.doi.org/10.1109/HPCSim.2014.6903792>.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y., 2017. Graph attention networks. <http://dx.doi.org/10.48550/ARXIV.1710.10903>, URL: <https://arxiv.org/abs/1710.10903>.
- Vijayaraghavan, S., Wu, L., Noels, L., Bordas, S.P.A., Natarajan, S., Beex, L.A.A., 2021. Neural-network acceleration of projection-based model-order-reduction for finite plasticity: Application to RVEs. <http://dx.doi.org/10.48550/ARXIV.2109.07747>, URL: <https://arxiv.org/abs/2109.07747>.
- Vlassis, N.N., Ma, R., Sun, W., 2020. Geometric deep learning for computational mechanics part I: anisotropic hyperelasticity. *Comput. Methods Appl. Mech. Engrg.* 371, 113299. <http://dx.doi.org/10.1016/j.cma.2020.113299>, URL: <https://www.sciencedirect.com/science/article/pii/S0045782520304849>.
- Šarkić Glumac, A., Jadhav, O., Despotović, V., Blocken, B., Bordas, S.P., 2023. A multi-fidelity wind surface pressure assessment via machine learning: A high-rise building case. *Build. Environ.* 234, 110135. <http://dx.doi.org/10.1016/j.buildenv.2023.110135>, URL: <https://www.sciencedirect.com/science/article/pii/S0360132323001622>.
- Wang, F., Eljarrat, A., Müller, J., Henninen, T., Erni, R., Koch, C., 2020. Multi-resolution convolutional neural networks for inverse problems. *Sci. Rep.* 10, 5730. <http://dx.doi.org/10.1038/s41598-020-62484-z>.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Yu, P.S., 2021. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* 32 (1), 4–24. <http://dx.doi.org/10.1109/TNNLS.2020.2978386>.
- Yan, S., Xiong, Y., Lin, D., 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. <http://dx.doi.org/10.48550/ARXIV.1801.07455>, URL: <https://arxiv.org/abs/1801.07455>.
- Zhang, Y., Gong, Z., Zhou, W., Zhao, X., Zheng, X., Yao, W., 2023. Multi-fidelity surrogate modeling for temperature field prediction using deep convolution neural network. *Eng. Appl. Artif. Intell.* 123, 106354. <http://dx.doi.org/10.1016/j.engappai.2023.106354>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197623005389>.
- Zhao, X., Gong, Z., Zhang, Y., Yao, W., Chen, X., 2023. Physics-informed convolutional neural networks for temperature field prediction of heat source layout without labeled data. *Eng. Appl. Artif. Intell.* 117, 105516. <http://dx.doi.org/10.1016/j.engappai.2022.105516>, URL: <https://www.sciencedirect.com/science/article/pii/S0952197622005061>.
- Zhao, L., Peng, X., Tian, Y., Kapadia, M., Metaxas, D.N., 2019. Semantic graph convolutional networks for 3D human pose regression. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, IEEE, <http://dx.doi.org/10.1109/cvpr.2019.00354>.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., Sun, M., 2020. Graph neural networks: A review of methods and applications. *AI Open* 1, 57–81. <http://dx.doi.org/10.1016/j.aiopen.2021.01.001>, URL: <https://www.sciencedirect.com/science/article/pii/S2666651021000012>.