## POLITECNICO DI TORINO

## Repository ISTITUZIONALE

## Graph-Informed Neural Networks for Regressions on Graph-Structured Data

Original
Graph-Informed Neural Networks for Regressions on Graph-Structured Data / Berrone, Stefano; Della Santa, Francesco; Mastropietro, Antonio; Pieraccini, Sandra; Vaccarino, Francesco. - In: MATHEMATICS. - ISSN 2227-7390.
ELETTRONICO. - (2022)

Availability:
This version is available at: 11583/2956960 since: 2022-03-20T13:14:00Z

Publisher:
MDPI

Published
DOI:

Terms of use.

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
(Article begins on next page)

Article

# Graph-Informed Neural Networks for Regressions on Graph-Structured Data 

Stefano Berrone ${ }^{1,2,3}{ }^{(\mathbb{D}}$, Francesco Della Santa ${ }^{1,2,3, *(\mathbb{C}}$, Antonio Mastropietro ${ }^{1,2,4}$ © , Sandra Pieraccini ${ }^{3,5}$ and Francesco Vaccarino ${ }^{1,2}$ (D)<br>1 Dipartimento di Scienze Matematiche (DISMA), Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy; stefano.berrone@polito.it (S.B.); antonio.mastropietro@polito.it (A.M.); francesco.vaccarino@polito.it (F.V.)<br>2 SmartData@PoliTO Center , Politecnico di Torino, 10129 Turin, Italy<br>3 Member of the INdAM-GNCS Research Group, 00100 Rome, Italy; sandra.pieraccini@polito.it<br>4 Addfor Industriale s.r.l., Via Giuseppe Giocosa 36/38, 10125 Turin, Italy<br>5 Dipartimento di Ingegneria Meccanica e Aerospaziale (DIMEAS), Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy<br>* Correspondence: francesco.dellasanta@polito.it

Citation: Berrone, S.; Della Santa, F.; Mastropietro, A.; Pieraccini, S.; Vaccarino, F. Graph-Informed Neural Networks for Regressions on Graph-Structured Data. Mathematics 2022, 10, 786. https:/ / doi.org/ 10.3390/math10050786

Academic Editor: Konstantin Kozlov

Received: 1 February 2022
Accepted: 25 February 2022
Published: 1 March 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/).


#### Abstract

In this work, we extend the formulation of the spatial-based graph convolutional networks with a new architecture, called the graph-informed neural network (GINN). This new architecture is specifically designed for regression tasks on graph-structured data that are not suitable for the wellknown graph neural networks, such as the regression of functions with the domain and codomain defined on two sets of values for the vertices of a graph. In particular, we formulate a new graphinformed (GI) layer that exploits the adjacent matrix of a given graph to define the unit connections in the neural network architecture, describing a new convolution operation for inputs associated with the vertices of the graph. We study the new GINN models with respect to two maximum-flow test problems of stochastic flow networks. GINNs show very good regression abilities and interesting potentialities. Moreover, we conclude by describing a real-world application of the GINNs to a flux regression problem in underground networks of fractures.


Keywords: graph neural networks; deep learning; regression on graphs
MSC: 05C21; 65D15; 68T07; 90C35

## 1. Introduction

Graphs are frequently used to describe and study many different phenomena, such as transportation systems, epidemic- or economic-default spread, electrical circuits, and social interactions; the literature typically refers to the use of graph theory to analyze such phenomena with the term "network analysis" [1].

Recently, new key contributions to network analyses have been proposed by the neural network (NN) community; in particular, deep learning (DL) approaches can be extended to graph-structured data via the so-called graph neural networks (GNNs). The origin of GNNs dates back to the late 2000s [2-4] when their processing was still too computationally expensive [5]. Nonetheless, the huge success of the convolutional neural networks (CNNs) inspired a new family of GNNs, re-defining the notation of convolutions for graph-structured data and developing the graph convolutional networks (GCNs). According to the taxonomy defined in [5], two main families of GCNs can be observed: the spectral-based GCNs [6-8], which are based on the spectral graph theory, and the spatial-based GCNs [9-12], which are based on the aggregation of the neighbor nodes' information. In particular, the spatial-based GCNs are, nowadays, preferred in many applications, thanks to their flexibility and efficiency [5].

Typically, GCNs are used to perform the following tasks on graph data [5]: (i) semisupervised node regression or classification; (ii) edge classification or link prediction;
and (iii) graph classification. Nonetheless, even if GCNs have been proven to be good instruments to learn graph data, some challenges still exist. The two main challenges for GCNs are [5]: (i) to build deep architectures with good performances; and (ii) to be scalable for large graphs. The first issue is the most problematic one; indeed, the success of DL lies in its depth, but the literature suggests that going deeper into a GCN is not usually beneficial [5]. Moreover, experimental results for the spectral-based GCNs showed that performances dropped considerably as the number of graph convolutional layers increased [13].

In this work, we present a new type of spatial-based graph convolutional layer designed for regression tasks on graph-structured data, a framework for which previous GCNs are not well suited. Given a graph $G$ with $n$ nodes, a regression task on graphstructured data based on $G$ consists of approximating a function $F: \Omega \subseteq \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}, m \leq n$, depending on the adjacency matrix of $G$, and that returns the $m$ values related to a fixed subset of $m$ nodes for each set of values assigned to the nodes of $G$. This type of regression task has applications in many interesting fields, such as circulation with demand (CwD) problems (see [chap. 7.7] in [14]), network interdiction models (NIMs) [15], and flux regression problems in underground fractured media $[16,17]$. A classic multi-layer perceptron (MLP), or its suitable variants, can perform this regression task on the graph data with a good performance [16,17], implicitly learning the node relationships during the training (see $[18,19]$ ). On the other hand, the current GCNs in the literature are not comparable to MLPs for such a regression task; indeed, as mentioned above, they are designed mainly for other kinds of tasks and, in practice, they cannot exploit deep architectures. Then, the idea is to define a new graph convolutional layer that exploits the graph structure to improve the training of the NN (compared to an MLP), and that makes it possible to build deep NN architectures. The new convolution operation for graph data that we define is closer to the convolution of CNNs (see [chap. 9] in [20]) than the convolution of all the other GCNs. Nonetheless, similarities with the classic NN4G layers in [3] and the diffusion-convolutional neural networks (DCNNs) in [21] exist.

Put simply, the simplest version of our graph layer is characterized by a filter with one weight $w_{i}$ associated with each graph node $v_{i}$. Then, the output feature of a node is computed by summing up the input features of the node itself and of its neighbors, where each one is multiplied by the corresponding node weights. We call this new type of graph layer a graph-informed (GI) layer. Indeed, given a scalar value $w_{j}$ associated with each graph vertex $v_{j}$, a GI layer looks like a fully-connected (FC) layer where, for each unit $v_{i}$ connected with a unit $v_{j}$ of the previous layer, the weight $w_{j i}$ is equal to 0 if $\left(v_{j}, v_{i}\right)$ is not an edge of the graph and $i \neq j$, otherwise $w_{j i}=w_{j}$ (see Equation (4) in Section 2).

Numerical experiments have shown the potentiality of the GI layers, which involves training deep NNs made up of a sequence of GI layers. We define these NNs as graphinformed neural networks (GINNs). In particular, the numerical experiments showed that GINNs were characterized by improved regression abilities with respect to MLPs, thanks also to their ability to overcome the depth problem typical of the other GCNs.

The work is organized as follows: in Section 2, the GI layers are formally introduced and defined, explaining their fundamental operations, properties, similarities, and differences with respect to other spatial-based graph convolutional layers. Section 3 is dedicated to the numerical experiments; in particular, we analyze the regression abilities of the GINNs on a maximum-flow regression problem and we compare the results with the performances obtained on the same problem with MLPs. We conclude the section with a real-world application, studying the application of GINNs to a flux regression problem in underground networks of fractures. In Section 4, we summarize the results and draw some conclusions.

## 2. Mathematical Formulation of the Graph-Informed Layers

In this section, we describe the mathematical formulation of the new GI layers, based on the adjacency matrix of a graph. In particular, we describe the mathematical details that define the function $\mathcal{L}^{G I}$, describing the action of a GI layer $L^{G I}$. From now on, we will call
the function describing the action of a generic NN layer as the characterizing function of the layer.

Definition 1 (Graph-Informed layer: basic form). Let $A \in \mathbb{R}^{n \times n}$ be the adjacency matrix characterizing a given graph $G=(V, E)$ without self-loops, and let $\widehat{A}$ be the matrix $\widehat{A}:=A+\mathbb{I}_{n}$, where $\mathbb{I}_{n} \in \mathbb{R}^{n \times n}$ is the identity matrix. Then, a graph-informed (GI) layer $L^{G I}$, with respect to the graph $G$, is an NN layer with $n$ units connected to a layer with outputs in $\mathbb{R}^{n}$ with a characterizing function $\mathcal{L}^{G I}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ defined by

$$
\begin{equation*}
\mathcal{L}^{G I}(\boldsymbol{x})=f\left(\widehat{W}^{\top} \boldsymbol{x}+\boldsymbol{b}\right) \tag{1}
\end{equation*}
$$

where:

- Given a vector $\boldsymbol{w} \in \mathbb{R}^{n}$ of weights associated with the vertices $V$, the defined filter of $L^{G I}$ the matrix $\widehat{W}$ is obtained by multiplying the $i$-th row of $\widehat{A}$ by the weight $w_{i}$, i.e.,

$$
\begin{equation*}
\widehat{W}:=\operatorname{diag}(\boldsymbol{w}) \hat{A} \tag{2}
\end{equation*}
$$

where $\operatorname{diag}(\boldsymbol{w})$ is the diagonal matrix with a diagonal that corresponds to vector $\boldsymbol{w}$;

- Given the layer activation function $f: \mathbb{R} \rightarrow \mathbb{R}$, we denote by $f$ the element-wise application of $f$;
- $\boldsymbol{b} \in \mathbb{R}^{n}$ is the vector of biases.

Broadly speaking, given a directed graph $G=(V, E)$, with $n$ nodes and an adjacency matrix $A \in \mathbb{R}^{n \times n}$, the main idea behind a GI layer is to generalize the convolutional layer filters to the graph-structured features. Indeed, the objective is to endow the layer with the implicit relationship between the features of the adjacent graph nodes, and also to take advantage of the sparse interaction- and parameter-sharing properties typical of convolutional NNs (see [chap. 9.2] in [20]).

Convolutional layers rely on the identification of images as lattices of pixels. The main idea for the GI layer formulation is to adapt convolutional layer concepts to graphs that are not characterized by a lattice structure. We generalize the filter mechanisms of the convolutional layers to the graph-structured data, introducing the concept of graph-based filters. In practice, for each node of the graph $G$, we consider a weight $w_{j}$ which is associated to node $v_{j} \in V$ and we re-define the convolution operation as

$$
\begin{equation*}
x_{i}^{\prime}=\sum_{j \in \mathrm{~N}_{\text {in }}(i) \cup\{i\}} x_{j} w_{j}+b_{i}, \tag{3}
\end{equation*}
$$

where

- $\quad x_{j}$ denotes the input feature of node $v_{j} \in V$, for each $j=1, \ldots, n$;
- $\quad \mathrm{N}_{\text {in }}(i)$ is the set of indices $j$, such that there exists an incoming edge $\left(v_{j}, v_{i}\right) \in E$;
- $\quad b_{i}$ is the bias corresponding to node $v_{i}$;
- $\quad x_{i}^{\prime}$ is the output feature associated to $v_{i}$, computed by the filter (see Figure 1).

For a non-directed graph $G$, Equation (3) does not change, since a non-directed edge $\left\{v_{j}, v_{i}\right\}$ is equivalent to two directed edges, $\left(v_{j}, v_{i}\right)$ and $\left(v_{i}, v_{j}\right)$. Indeed, Definition 1 holds for both directed and non-directed graphs.

Similar to the convolutional layers, which act on the current pixel and on all its neighbors for computing the output feature, in (3) the layer acts on $x_{i}$ and on the values associated with the incoming neighbors of $v_{i}$ for the computation of $x_{i}^{\prime}$ (see Figure 1). Nonetheless, despite the inspiration received from convolutional layers, a GI layer $L^{G I}$ described by (3) can be seen also as a constrained FC layer, where the weights are such that

$$
w_{j i}= \begin{cases}w_{j}, & \text { if }\left(v_{j}, v_{i}\right) \in E  \tag{4}\\ w_{i}, & \text { if } j=i \\ 0, & \text { otherwise }\end{cases}
$$

for each $i, j=1, \ldots, n$.


Figure 1. Case of non-directed graph with $n=4$ nodes. Example of the action of a filter $\boldsymbol{w} \in \mathbb{R}^{4}$ (grey "layer" of the plot) of a GI layer, applied to feature $x_{1}$ of the first graph-node $v_{1}$, for the computation of $x_{1}^{\prime}$; for simplicity, the bias is not illustrated. The orange edges describe the multiplication of feature $x_{i}$, of node $v_{i}$, with the filter's weight $w_{i}$, for each $i=1, \ldots, 4$.

In the next sections, we generalize the action of these kinds of layers to make them able to: (i) receive any arbitrary number $K \geq 1$ of input features for each node; and (ii) return any arbitrary number $F \geq 1$ of output features for each node.

### 2.1. Generalization to K Input Node Features

Equation (1) describes the simplest case of GI layers, where just one feature is considered for each vertex of the graph for both the inputs and the outputs. We start generalizing the previous definition, taking into account a larger number of features tackled by $L^{G I}$.

Definition 2 (Graph-Informed layer with $K$ input features per node). Let $G, A$, and $\widehat{A}$ be as in Definition 1. Then, a GI layer with $K \in \mathbb{N}$ input features is an NN layer with $n$ units connected to a layer with outputs in $\mathbb{R}^{n \times K}$ with a characterizing function $\mathcal{L}^{G I}: \mathbb{R}^{n \times K} \rightarrow \mathbb{R}^{n}$ defined by

$$
\begin{equation*}
\mathcal{L}^{G I}(X)=f\left(\widetilde{W}^{\top} \operatorname{vertcat}(X)+\boldsymbol{b}\right) \tag{5}
\end{equation*}
$$

where:

- $\quad X \in \mathbb{R}^{n \times K}$ is the input matrix (i.e., the output of the previous layer) and vertcat $(X)$ denotes the vector in $\mathbb{R}^{n K}$ obtained by concatenating the columns of $X$;
- Given the matrix $W \in \mathbb{R}^{n \times K}$, the defined filter of $L^{G I}$, whose columns $\boldsymbol{w}_{\cdot 1}, \ldots, \boldsymbol{w}_{\cdot K} \in \mathbb{R}^{n}$ are the vectors of weights associated with the $k$-th input feature of the graph's vertices, the matrix $\widetilde{W} \in \mathbb{R}^{n K \times n}$ is defined as

$$
\widetilde{W}:=\left[\begin{array}{c}
\widehat{W}^{(1)}  \tag{6}\\
\vdots \\
\widehat{W}^{(K)}
\end{array}\right]=\left[\begin{array}{c}
\operatorname{diag}\left(\boldsymbol{w}_{\cdot 1}\right) \widehat{A} \\
\vdots \\
\operatorname{diag}\left(\boldsymbol{w}_{\cdot K}\right) \widehat{A}
\end{array}\right] \in \mathbb{R}^{n K \times n} .
$$

The idea behind the generalization from Definitions 1 to 2 is rather simple. Let $L$ be an NN layer with outputs in $\mathbb{R}^{n \times K}, K \geq 1$. Therefore, a generic output of $L$ is a matrix $X \in \mathbb{R}^{n \times K}$, whose row $i \in\{1, \ldots, n\}$ describes the $K$ features $x_{i 1}, \ldots, x_{i K}$ of node $v_{i}$; on the
other hand, each column $x_{\cdot 1}, \ldots, x_{\cdot K}$ of $X$ is equivalent to the output of an NN layer with outputs in $\mathbb{R}^{n}$.

Therefore, the generalization consists of summing up the action of the $K$ "basic" single-
 this sum, we add the bias vector and we apply the activation function. However, this approach is equivalent to (5), i.e., defining one filter $W$ obtained from the concatenation of the basic filters. Indeed:

$$
\begin{equation*}
\sum_{k=1}^{K} \widehat{W}^{(k) \top} \boldsymbol{x}_{\cdot k}=\widetilde{W}^{\top} \operatorname{vertcat}(X) \tag{7}
\end{equation*}
$$

Remark 1 (Parallelism with convolutional layers). It is worth noting that the operations summarized in (5) are an adaptation of the convolutional layer operations to the graph-based inputs. Indeed, the input $X \in \mathbb{R}^{n \times K}$ is equivalent to an $n \times 1$ image with $K$ channels, while $\boldsymbol{w}_{\cdot k}$ is equivalent to the part of the convolutional filter corresponding to the $k$-th channel of the input image. Then, the output $\mathcal{L}^{G I}(X) \in \mathbb{R}^{n}$ is equivalent to the so-called activation map of the convolutional layers.

### 2.2. Generalization to F Output Node Features

We can further generalize (5) by increasing the number of output features per node returned by the GI layer. This operation is equivalent to building a GI layer characterized by a number $F \geq 1$ of matricial filters, where each one used to compute one of the output features. In a nutshell, the output of these general GI layers is a matrix $Y \in \mathbb{R}^{n \times F}$ whose $l$-th column $y_{\cdot l} \in \mathbb{R}^{n}, l=1 \ldots, F$, describes the $l$-th feature of the nodes of $G$.

Definition 3 (Graph-Informed layer: general form). Let $G, A, \widehat{A}$ be as in Definition 1. Then, a GI layer with $K \in \mathbb{N}$ input features and $F \in \mathbb{N}$ output features is an NN layer with nF units connected to a layer with outputs in $\mathbb{R}^{n \times K}$ with a characterizing function $\mathcal{L}_{G I}: \mathbb{R}^{n \times K} \rightarrow \mathbb{R}^{n \times F}$ defined by

$$
\begin{equation*}
\mathcal{L}^{G I}(X)=f\left(\widetilde{\mathbf{W}}^{\top} \operatorname{vertcat}(X)+B\right) \tag{8}
\end{equation*}
$$

where:

- We define the filter of $L^{G I}$, the tensor $\mathbf{W} \in \mathbb{R}^{n \times K \times F}$, given by the concatenation along the third dimension of the weight matrices $W^{(1)}, \ldots W^{(F)} \in \mathbb{R}^{n \times K}$, corresponding to the $F$ output features of the nodes. Each column $\boldsymbol{w}_{\cdot k}^{(l)} \in \mathbb{R}^{n}$ of $W^{(l)}$ is the basic filter describing the contribution of the $k$-th input feature to the computation of the l-th output feature of the nodes, for each $k=1, \ldots, K$, and $l=1, \ldots, F$;
- The tensor $\widetilde{\mathbf{W}} \in \mathbb{R}^{n K \times F \times n}$ is defined as the concatenation along the second dimension (i.e., the column dimension) of the matrices $\widetilde{W}^{(1)}, \ldots, \widetilde{W}^{(F)}$, such that

$$
\widetilde{W}^{(l)}:=\left[\begin{array}{c}
\widehat{W}^{(l, 1)}  \tag{9}\\
\vdots \\
\widehat{W}^{(l, K)}
\end{array}\right]=\left[\begin{array}{c}
\operatorname{diag}\left(\boldsymbol{w}_{\cdot 1}^{(l)}\right) \widehat{A} \\
\vdots \\
\operatorname{diag}\left(\boldsymbol{w}_{\cdot K}^{(l)}\right) \widehat{A}
\end{array}\right] \in \mathbb{R}^{n K \times n},
$$

for each $l=1, \ldots, F$. Before the concatenation, the matrices $\widetilde{W}^{(1)}, \ldots, \widetilde{W}^{(F)}$ are reshaped as tensors in $\mathbb{R}^{n k \times 1 \times n}$ (see Figure 2);

- the operation $\widetilde{\mathbf{W}}^{\top}$ vertcat $(X)$ is a tensor-vector product (see Remark 2);
- $\quad B \in \mathbb{R}^{n \times F}$ is the matrix of the biases, i.e., each column $\boldsymbol{b}_{. l}$ is the bias vector corresponding to the l-th output feature of the nodes.

Notation 1. From now on, for the sake of simplicity, for each matrix $X \in \mathbb{R}^{n \times K}$, we denote by $x$ the vector vertcat $(X) \in \mathbb{R}^{n K}$.


Figure 2. Tensor $\widetilde{\mathbf{W}}$ obtained concatenating along the second dimension of the matrices $\widetilde{W}^{(1)}, \ldots, \widetilde{W}^{(F)} \in \mathbb{R}^{n K \times n}$. Before the concatenation, the matrices are reshaped as tensors in $\mathbb{R}^{n k \times 1 \times n}$.

The generalization of (5) to the case of $F$ output features is built as a function that, for each $X \in \mathbb{R}^{n \times K}$, a matrix is returned $Y \in \mathbb{R}^{n \times F}$ whose $l$-th column $y_{\cdot l}$, for $l=1, \ldots, F$, is defined as the application of (5) with respect to a proper filter $W^{(l)} \in \mathbb{R}^{n \times K}$. Indeed, given

$$
\begin{equation*}
\boldsymbol{y}_{\cdot l}=f\left(\widetilde{W}^{(l) \top} \boldsymbol{x}+\boldsymbol{b}_{\cdot l}\right) \tag{10}
\end{equation*}
$$

where $\boldsymbol{b}_{\cdot l} \in \mathbb{R}^{n}$ is the bias vector associated to the $l$-th filter, we have

$$
Y=\left[\begin{array}{ccc}
\boldsymbol{y}_{\cdot 1} & \cdots & \boldsymbol{y}_{\cdot F}
\end{array}\right]=\left[\boldsymbol{f}\left(\widetilde{W}^{(1) \top} \boldsymbol{x}+\boldsymbol{b}_{\cdot 1}\right) \quad \cdots \quad \boldsymbol{f}\left(\widetilde{W}^{(F) \top} \boldsymbol{x}+\boldsymbol{b}_{\cdot F}\right)\right]=\boldsymbol{f}\left(\widetilde{\mathbf{W}}^{\top} \boldsymbol{x}+B\right)
$$

Put simply, the generalization to the $F$ output features can be interpreted as a repetition of (5), with respect to $F$ different filters and biases, grouping the results in a matrix $Y$.

Remark 2. We recall that the matrix-tensor product of a matrix $M \in \mathbb{R}^{p \times q}$ by a tensor $\mathbf{T} \in \mathbb{R}^{q \times r \times s}$ is given by

$$
M \cdot \mathbf{T}=\mathbf{P} \in \mathbb{R}^{p \times r \times s},
$$

where the $(i, j, k)$-th component $p_{i j k}$ of the tensor $\mathbf{P}$ is defined as

$$
p_{i j k}=\sum_{h=1}^{q} m_{i h} t_{h j k}
$$

where $m_{i h}$ and $t_{h j k}$ are components of $M$ and $T$, respectively. Analogously, we can extend this product to tensor-matrix or tensor-tensor pairs.

Moreover, we recall that, for a three-way tensor as $\widetilde{\mathbf{W}}$, the transpose is defined such that the $(i, j, k)$-th element of $\widetilde{\mathbf{W}}^{\top}$ is equal to the $(k, j, i)$-th element of $\widetilde{\mathbf{W}}$.

Remark 3 (Total number of parameters). The total number of parameters in a GI layer, with a characterizing function (8), is $n K F+n F$, i.e., the number of weights plus the number of biases. Let us recall that the number of parameters of a fully-connected layer, with an input shape $n$ and an output shape $M$ is $n M+M$; then, in the case of $M=n$ and $(K F+F)<(n+1)$, we see that the

GI layers have a smaller number of parameters to be trained. This observation is important in the case of very large graphs $G$ (i.e., $n \gg 1$ ).

Remark 4 (GI layer contextualization). To the best of the authors' knowledge, the GI layers introduced above define a novel typology of spatial GCNs. Equation (1) partially recalls the NN4G layer (see [sec. V.B.] in [3,5]), if one removes both the so-called residual and skip connections (i.e., the extra connections used to directly transfer information between non-consecutive layers). However, the formulation given by the authors in Equation (1) is generalizable to the tensor form (8), i.e., to multiple input/output features, unlike the NN4G layers. It is worth noting that a tensor form, such as (8), is very useful to manage graph-structured regression problems with more than one feature per node.

Analogously, there are few similarities between the simple $L^{\text {GI }}$ layer of Equation (1) and the diffusion-convolutional NNs (DCNNs) of [21] which can be observed, but these GCNs are still different from the GINNs. Indeed, DCNNs are made for different types of tasks, such as node classification or graph classification tasks, inferred from the known features of a subset of nodes. In addition, DCNN layers are based on a degree-normalized transition matrix, computed from the adjacency matrix, "that gives the probability of jumping from node i to node j in one step" [21].

Other similarities between the GINNs and other models can be observed in [22,23], where the adjacency matrix is used to describe the flow of information. Nonetheless, in [22], the NN is built connecting a set of simpler NNs, according to the adjacency matrix. In [23], the interconnected NNs are trained similar to a physics-informed NN (see [24,25]).

In the end, we point the attention of the reader to the fact that, from a theoretical point of view, nothing prevents us from adding a softmax layer at the end of a GINN to extend the new model architecture to cover graph classification tasks with respect to vertex labels (like CNNs for image classification); however, we defer the study of this possibility to future work.

### 2.3. Additional Properties for GI Layers

The GI layers, in their general formulation (8), can be endowed with additional operations. As is commonly done for the convolutional layers, we add the possibility to endow the GI layers with a pooling operation. However, this operation is different from the one typically used in convolutional layers. Indeed, we define a pooling for GI layers that aggregates the information in the columns of the output matrix, i.e., the values of the $F$ output features of each graph vertex. Given a "reducing" operation (e.g., the mean, the max, the sum, etc.), labeled as rdc, and applied to each row of the matrix returned by (8), the pooling operation for GI layers modifies (8) in the following way:

$$
\begin{equation*}
\mathcal{L}^{(G I ; \mathrm{rdc})}(X)=\operatorname{rdc}\left(f\left(\widetilde{\mathbf{W}}^{\top} x+B\right)\right) \tag{11}
\end{equation*}
$$

where rdc is applied row-wise. For example, let $Y \in \mathbb{R}^{n \times F}$ denote the argument of the pooling operation in (11), namely $Y=f\left(\widetilde{\mathbf{W}}^{\top} \boldsymbol{x}+B\right)$; the max-pooling operation for a GI layer is such that:

$$
\mathcal{L}^{(G I ; \max )}(X)=\left[\begin{array}{c}
\max \left\{y_{11}, \ldots, y_{1 F}\right\}  \tag{12}\\
\vdots \\
\max \left\{y_{n 1}, \ldots, y_{n F}\right\}
\end{array}\right] \in \mathbb{R}^{n}
$$

Note that the pooling operation can be generalized to the application of subgroups to filters, instead of to all the filters. In this case, the pooling operation returns a matrix $Y \in \mathbb{R}^{n \times F^{\prime}}$, with $F^{\prime}<F$.

Another operation that is defined for GI layers is the application of a mask on the graph, such that the layer returns values only for a subset $\left\{v_{i_{1}}, \ldots, v_{i_{m}}\right\}$ of the chosen nodes. Let $I=\left\{i_{1}, \ldots, i_{m}\right\} \subseteq\{1, \ldots, n\}$ label the subset of nodes on which we want to focus the output of the GI layer. Then, a GI layer with a mask operation defined by the set I returns a
sub-matrix $Y^{\prime} \in \mathbb{R}^{m \times F}$ of the matrix $Y \in \mathbb{R}^{n \times F}$ defined by (8), obtaining extracting rows with index in $I$; namely,

$$
\mathcal{L}^{(G I ; I)}(X)=\left[\begin{array}{c}
\left(f\left(\widetilde{\mathbf{W}}^{\top} \boldsymbol{x}+B\right)\right)_{i_{1}}  \tag{13}\\
\vdots \\
\left(f\left(\widetilde{\mathbf{W}}^{\top} \boldsymbol{x}+B\right)\right)_{i_{m} .}
\end{array}\right] \in \mathbb{R}^{m \times F}
$$

We end this section with the following proposition, characterizing the relationship between the input and the output features of the graph nodes with respect to a GINN with a subset of $H$ consecutive GI layers. The proof of the statement is straightforward.

Proposition 1 (Number of consecutive GI layers and node interactions). Let $H \in \mathbb{N}, H \geq 1$, be fixed and let $A \in \mathbb{R}^{n \times n}$ be the adjacency matrix of a given graph $G$. Let us consider a GINN with a subset of $H$ consecutive GI layers $L_{1}^{G I}, \ldots, L_{H}^{G I}$, built according to $A$ and with $L_{h}^{G I}$ connected to $L_{h+1}^{G I}$, for $h=1, \ldots, H-1$. Let $d_{i j}:=\operatorname{dist}_{G}\left(v_{i}, v_{j}\right) \in \mathbb{N} \cup\{+\infty\}$ be the distance between node $v_{i}$ and node $v_{j}$ in $G$. Then, the input feature corresponding to node $v_{i}$ in $L_{1}^{G I}$ contributes to the computation of the output feature corresponding to $v_{j}$ in $L_{H}^{G I}$ if $H \geq d_{i j}$.

The proposition above introduces a dependency of a GINN's depth on the complexity of the graph $G=(V, E)$. Let $F: \Omega \subseteq \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ be a function defined on the $n$ vertices of $G$, and returning a vector of $m$ values associated with the vertices $v_{i_{1}}, \ldots, v_{i_{m}} \in V$. Let $\widehat{\boldsymbol{F}}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ be the characterizing function of a GINN that approximates $F$. If the output feature of vertex $v_{j} \in\left\{v_{i_{1}}, \ldots, v_{i_{m}}\right\}$ through $\boldsymbol{F}$ depends on the input feature of vertex $v_{i}$, then the GINN needs at least $d_{i j}=\operatorname{dist}_{G}\left(v_{i}, v_{j}\right)$ consecutive GI layers to guarantee that the input feature of $v_{i}$ contributes in making predictions for the output feature of $v_{j}$.

## 3. Numerical Tests

In this section, we study the potentialities of the GI layers, comparing the regression abilities of GINNs and MLPs for graph-structured data.

The main test problem we consider is the maximum-flow problem: given a flow network, that is, a graph with a source, a sink, and capacities defined on the edges, the goal is to find the maximum flow value that can reach the sink (see [ch. 7.1] in [14]). In particular, we are interested in the stochastic maximum-flow problem, i.e. a problem where the edge capacities are modeled as random variables and the target is to find the distribution of the maximum-flow (e.g., see [26]).

The stochastic maximum-flow problem is a sufficiently general problem for testing GINNs, and it has many interesting applications in network analyses, such as the circulation with demand (CwD) problems (see [ch. 7.7] in [14] and network interdiction models (NIMs) [15]. Put simply, a CwD problem should identify whether or not the maximum flow satisfies a given demand, varying the supply provided by the source and the capacities of the edges; a NIM describes a game in which one or more agents modify the edge capacities to minimize/maximize the maximum flow of the network. These models have many interesting real-world applications, such as the administration of city traffic, the optimization of goods distributions, and identifying vulnerabilities in an operational system.

The section is organized as follows: after a brief description of the maximum-flow problem (Section 3.1), we illustrate the maximum-flow regression problem (Section 3.2). Then, in Section 3.3, we report and compare the performances of the trained GINNs and MLPs. We conclude the section with an example that shows the potentialities of using GINNs in practical applications related to realistic underground flow simulations (Section 3.4).

### 3.1. The Maximum-Flow Problem

Flow networks are useful models to describe transportation networks, i.e., networks where some sort of traffic flows from a source to a sink along the edges, using the nodes as switches to let the traffic move from an edge to another one (see [ch. 7.1] in [14]). Here, we briefly recall the definition of a flow network.

Definition 4 (Flow Network). A flow network $\mathcal{G}=(G, s, t, c)$ is a directed graph $G=(V, E)$, of nodes $V$ and edges $E \subseteq V \times V$, such that:

- The two nodes $s, t \in V, s \neq t$, are defined as the source and the sink of the network, respectively;
- $\quad c$ is a real-valued non-negative function defined on the edges, $c: E \rightarrow \mathbb{R}_{\geq 0}$, assigning to each edge $e \in E$ a capacity $c_{e}:=c(e)$.

A flow network $\mathcal{G}$ can be endowed with a flow function.
Definition 5 (Flow). Let $\mathcal{G}$ be a flow network. An s-t flow (or just flow) on $\mathcal{G}$ is a function

$$
\varphi: E \rightarrow \mathbb{R}_{\geq 0}
$$

that satisfies the following properties:

- The capacity condition: for each $e \in E$, it holds $0 \leq \varphi(e) \leq c_{e}$;
- The conservation condition: for each $v \in V \backslash\{s, t\}$, the amount of flow entering $v$ must be equal to the amount of flow leaving $v$, i.e.,

$$
\sum_{e \in E_{\text {in }}(v)} \varphi(e)=\sum_{e \in E_{\text {out }}(v)} \varphi(e), \quad \forall v \in V \backslash\{s, t\},
$$

where $E_{\text {in }}(v) \subset V$ is the subset of the incoming edges of $v$, and $E_{\text {out }}(v) \subset V$ is the subset of outcoming edges of $v$;

- The amount of flow leaving the source s must be greater than, or equal to, the one entering s, i.e., $\sum_{e \in E_{\text {in }}(s)} \varphi(e) \leq \sum_{e \in E_{\text {out }}(s)} \varphi(e)$.

For the sake of notation, for each $v \in V$ we set

$$
\varphi_{\text {in }}(v)=\sum_{e \in E_{\text {in }}(v)} \varphi(e), \quad \varphi_{\text {out }}(v)=\sum_{e \in E_{\text {out }}(v)} \varphi(e), \quad \Delta \varphi_{v}=\varphi_{\text {out }}(v)-\varphi_{\text {in }}(v),
$$

and we call the flow value of a vertex $v$ the quantity $\Delta \varphi_{v}$.
Then, due to the conservation condition, it holds that $\Delta \varphi_{v}=0$, for each $v \in V \backslash\{s, t\}$, and $\Delta \varphi_{s} \geq 0$. Note that the flow value of the source $s$ is equal to the opposite of the flow value of the $\operatorname{sink} t$, i.e., $\Delta \varphi_{t}=-\Delta \varphi_{s}$; for this reason, we refer to $\Delta \varphi_{s}$ as the flow value of the network.

One of the most common issues concerning a flow network $\mathcal{G}$ is to find a flow that maximizes the effective total flow of the $\operatorname{sink} t$, i.e., to find $\varphi^{*}$, such that

$$
\varphi^{*}=\arg \max _{\varphi}\left|\Delta \varphi_{t}\right| .
$$

Such a kind of problem is called maximum-flow problem, and it can be solved through the linear programming or many other algorithms (e.g., [27-31]). From a practical point of view, the relationship between the maximum-flow problem and the minimum-cut problem on a flow network $\mathcal{G}$ is particularily important (see [ch. 7.2] in [14] for more details).

Remark 5 (Flow networks and undirected graphs). Definitions 4 and 5 can be extended to the more complicated case of undirected graphs. Indeed, as observed in Section 2, the non-directed edges $\left\{v_{j}, v_{i}\right\}$ of a graph are equivalent to the two directed edges $\left(v_{j}, v_{i}\right)$ and $\left(v_{i}, v_{j}\right)$. Then, a flow network based on an undirected graph $G=(V, E)$ can be defined as a flow network $\mathcal{G}=\left(G^{\prime}, s, t, c\right)$, where
$G^{\prime}=\left(V, E^{\prime}\right)$ is a directed graph, such that $(u, v),(v, u) \in E^{\prime}$ if $\{u, v\} \in E$, whose capacity is defined on the edges of $G$,i.e., $c: E \rightarrow \mathbb{R}_{\geq 0}$. As a result, a flow $\varphi$ defined on such a flow network is a function $\varphi: E^{\prime} \rightarrow \mathbb{R}_{\geq 0}$ characterized by a slightly different capacity condition; namely, $\varphi$ is such that

$$
0 \leq \varphi((u, v))+\varphi((v, u)) \leq c(\{u, v\}), \quad \forall\{u, v\} \in E .
$$

Another approach is to introduce an arbitrary ordering, denoted by " $<$ ", on the graph nodes and define a directed graph $G^{\prime}=\left(V, E^{\prime}\right)$, such that $(u, v) \in E^{\prime}$ if $\{u, v\} \in E$ and $u<v$. In this case, a flow $\varphi$ on the flow network $\mathcal{G}=\left(G^{\prime}, s, t, c\right)$, with $c$ defined on the edges $E^{\prime}$ is a function $\varphi: E^{\prime} \rightarrow \mathbb{R}$, such that the capacity condition is

$$
0 \leq|\varphi(e)| \leq c(e), \quad \forall e \in E^{\prime}
$$

and where the entering/exiting behavior of the flows is described by the sign of $\varphi(e)$ and not by the edge direction; i.e., for $(u, v) \in E^{\prime}$, if $\varphi((u, v))>0$ the flow $\varphi((u, v))$ enters in $v$, whereas if $\varphi((u, v))<0$ then $\varphi((u, v))$ enters in $u$. This latter approach is mainly adopted by software implementations.

### 3.1.1. The Stochastic Maximum-Flow Problem

The idea of the flow network, flow, and the maximum-flow problem can be easily extended to a stochastic framework, in which edge capacities are modeled as random variables.

Definition 6 (Stochastic flow network). A stochastic flow network $\mathcal{G}=(G, s, t, p)$ is a directed graph $G=(V, E)$ of nodes $V$ and edges $E \subseteq V \times V$, such that:

- The two nodes $s, t \in V, s \neq t$, are defined as the source and the sink of the network, respectively;
- $\quad p$ is a real-valued non-negative probability distribution for the edge capacities of the network.

We let $\mathcal{G}(c)$ denote the flow network ( $G, s, t, c$ ) with edge capacities returned by $c$ sampled from $p$. More specifically, let $e_{1}, \ldots, e_{|E|}$ be all the edges of $G$; then $\mathcal{G}(c)=(G, s, t, c)$ if:

- $\quad c \in \mathbb{R}^{|E|}$ is a vector whose $c_{i}$ is sampled from $p ;$
- $\quad$ The function $c$ is such that $c\left(e_{i}\right)=c_{i}$, for each $i=1, \ldots,|E|$.

We denote by $\varphi^{(c)}$ a flow defined on the flow network $\mathcal{G}(\boldsymbol{c})$.
The stochastic maximum-flow problem consists of finding the flow

$$
\begin{equation*}
\varphi^{*(c)}=\arg \max _{\varphi^{(c)}}\left|\Delta \varphi_{t}^{(c)}\right|, \tag{14}
\end{equation*}
$$

for each fixed vector $\boldsymbol{c}$.
Alternatively, in stochastic maximum-flow problems, one may seek the flow distribution and/or its moments, or the maximum flow value entering the $\operatorname{sink} t$, i.e.,

$$
\begin{equation*}
\left|\Delta \varphi_{t}^{*(c)}\right|=\max _{\varphi^{(c)}}\left|\Delta \varphi_{t}^{(c)}\right|, \tag{15}
\end{equation*}
$$

### 3.2. The Maximum-Flow Regression Problem

A maximum-flow regression problem, with respect to a given stochastic flow network, consists of finding a function that, for each capacity vector $c$, returns an approximation of the maximum flow $\left|\Delta \varphi_{t}^{*(\boldsymbol{c})}\right|$ or an approximation of all the flows reaching the sink $t$.

Let $\mathcal{G}$ be a stochastic flow network of $n=|E|$ edges and, without a loss of generality, let $e_{1}, \ldots, e_{m} \in E, m \leq n$ be all the incoming edges of the $\operatorname{sink} t$. Let $F: \Omega \subseteq \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ be a function, such that

$$
\begin{equation*}
\boldsymbol{F}(\boldsymbol{c})=\left[\varphi^{*(\boldsymbol{c})}\left(e_{1}\right), \ldots, \varphi^{*(\boldsymbol{c})}\left(e_{m}\right)\right]^{\top}:=\boldsymbol{\varphi}, \tag{16}
\end{equation*}
$$

for each capacity vector $c \in \Omega \subseteq \mathbb{R}^{n}$ with the elements sampled from the distribution $p$ of the given network $\mathcal{G}$.

From now on, for the sake of simplicity, we drop the dependency from $c$ and the star symbol from the elements of $\boldsymbol{\varphi}$, denoted by $\varphi_{1}, \ldots, \varphi_{m}$ the $m$ elements of the vector $\boldsymbol{\varphi}=\boldsymbol{F}(\boldsymbol{c})$. Moreover, assuming the convention of the non-negative flow functions on the graph (see Section 3.1), we denote by $\varphi$ the $\ell_{1}$-norm of $\boldsymbol{\varphi}$; specifically:

$$
\begin{equation*}
\sum_{j=1}^{m} \varphi_{i}=\sum_{j=1}^{m}\left|\varphi_{i}\right|=\|\boldsymbol{\varphi}\|_{1}=: \varphi \tag{17}
\end{equation*}
$$

Then, the target maximum flow with respect to $\boldsymbol{c}$ coincides with $\varphi$; indeed, $\left|\Delta \varphi_{t}{ }^{*(\boldsymbol{c})}\right|=$ $\sum_{j=1}^{m} \varphi^{*(c)}\left(e_{j}\right)=\sum_{j=1}^{m} \varphi_{i}=\varphi$.

Given the target function $\boldsymbol{F}$ defined by (16), we consider the maximum-flow regression problem with respect to $\mathcal{G}$ looking for an NN with a characterizing function $\widehat{\boldsymbol{F}}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$, such that $\widehat{\boldsymbol{F}}(\boldsymbol{c})$ approximates $\boldsymbol{F}(\boldsymbol{c})$ for each capacity vector $\boldsymbol{c}$. Namely, setting $\widehat{\boldsymbol{\varphi}}=\widehat{\boldsymbol{F}}(\boldsymbol{c})$ and $\boldsymbol{\varphi}=\boldsymbol{F}(\boldsymbol{c})$, we seek $\widehat{\boldsymbol{\varphi}} \approx \boldsymbol{\varphi}$. To train an NN with respect to $\boldsymbol{F}$, we build a dataset (i.e., a multi-set) $\mathcal{D}$ of pairs $(\boldsymbol{c}, \boldsymbol{\varphi}) \in \mathbb{R}^{n} \times \mathbb{R}^{m}$, with $\boldsymbol{\varphi}=\boldsymbol{F}(\boldsymbol{c})$, where the capacity vectors are sampled with respect to the distribution $p$ of $\mathcal{G}$; then, $\mathcal{D}$ is split into a training set $\mathcal{T}$, a validation set $\mathcal{V}$, and a test set $\mathcal{P}$ of arbitrary cardinalities. In particular, denoting $\Theta$ as the multi-set $\mathcal{T}+\mathcal{V}$, we denote $\vartheta$ as the total number of pairs involved in the training operations of the NN, i.e., the pairs in the multi-set $\Theta(\vartheta:=|\Theta|=|\mathcal{T}|+|\mathcal{V}|)$. For more details about multi-sets, see Appendix A.

Once an NN is trained, we evaluate its regression performances by computing two performance measures on the test set $\mathcal{P}$ : the edge-wise average mean relative error $\left(\mathrm{MRE}_{a v}\right)$, and the mean relative error on the predicted maxflow $\left(\mathrm{MRE}_{\varphi}\right)$. These two errors represent the mean relative error (weighted with respect to the true maxflow) of the predicted flows of the $m$ edges $e_{1}, \ldots, e_{m}$ and the mean relative error of the predicted maxflow $\widehat{\varphi}:=\sum_{i=1}^{m} \widehat{\varphi}_{j}$ (i.e., the sum of the elements of $\widehat{\boldsymbol{\varphi}}=\widehat{\boldsymbol{F}}(\boldsymbol{c})$ ), respectively. For each prediction $\widehat{\boldsymbol{\varphi}}$, let us denote

$$
\operatorname{err}(\widehat{\boldsymbol{\varphi}}, \boldsymbol{\varphi})=\left[\operatorname{err}_{1}(\widehat{\boldsymbol{\varphi}}, \boldsymbol{\varphi}), \ldots, \operatorname{err}_{m}(\widehat{\boldsymbol{\varphi}}, \boldsymbol{\varphi})\right]^{\top}:=\left[\frac{\left|\widehat{\varphi}_{1}-\varphi_{1}\right|}{\varphi}, \ldots, \frac{\left|\widehat{\varphi}_{m}-\varphi_{m}\right|}{\varphi}\right]^{\top}
$$

as the vector of relative errors computed with respect to the true maxflow $\varphi=\sum_{j=1}^{m} \varphi_{j}$ (see (17)). Then, the performance measures $\mathrm{MRE}_{a v}$ and $\mathrm{MRE}_{\varphi}$ are defined as

$$
\begin{equation*}
\operatorname{MRE}_{a v}(\mathcal{P}):=\frac{1}{m} \sum_{j=1}^{m}\left(\frac{1}{|\mathcal{P}|} \sum_{(c, \boldsymbol{\varphi}) \in \mathcal{P}} \operatorname{err}_{j}(\widehat{\boldsymbol{\varphi}}, \boldsymbol{\varphi})\right) \tag{18}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{MRE}_{\varphi}(\mathcal{P}):=\frac{1}{|\mathcal{P}|} \sum_{(c, \varphi) \in \mathcal{P}} \frac{|\widehat{\varphi}-\varphi|}{\varphi}, \tag{19}
\end{equation*}
$$

respectively.
The smaller both the $\mathrm{MRE}_{a v}$ and the $\operatorname{MRE}_{\varphi}$ values are on the test set, the better the performances of the NN, with respect to the maximum-flow regression task, are.

Remark 6 (Interpretation of $\mathrm{MRE}_{a v}$ and $\mathrm{MRE}_{\varphi}$ ). It is worth highlighting the different meanings of the errors (18) and (19): MRE $E_{a v}$ describes the average quality of the NN in predicting the single elements $\varphi_{1}, \ldots, \varphi_{m}$ of the target vector $\varphi$, while $M R E_{\varphi}$ describes the ability of the NN to predict a vector $\widehat{\boldsymbol{\varphi}}$, such that the corresponding maxflow $\widehat{\varphi}=\sum_{j=1}^{m} \widehat{\varphi}_{j}$ approximates the true maxflow $\varphi=\sum_{j=1}^{m} \varphi_{j}$. Therefore, a small $M R E_{a v}$ corresponds to a good approximation of the flow vectors (i.e., $\widehat{\boldsymbol{\varphi}} \approx \boldsymbol{\varphi}$ ) and a small $M R E_{\varphi}$ corresponds to a good approximation of the maximum-flows (i.e., $\widehat{\varphi} \approx \varphi$ ). Nonetheless, it is important to point out that a small $M R E_{a v}$ does not necessarily imply a small $M R E_{\varphi}$, and vice-versa. For example, an $N N$ with large $M R E_{a v}$, characterized by the underestimation of the flows $\varphi_{j_{1}}$ and by the overestimation of the flows $\varphi_{j_{2}}$, may return a small $M R E_{\varphi}$ because the sum of the flows is not so far from the true maximum-flow; similarly, a
large $M R E_{\varphi}$ can be obtained from a sufficiently small $M R E_{a v}$ if, e.g., the $N N$ underestimates or overestimates all the flows $\varphi_{1}, \ldots, \varphi_{m}$ equivalently, such that $\widehat{\boldsymbol{\varphi}} \approx \varphi$ but $\widehat{\varphi} \not \approx \varphi$.

### 3.2.1. Line Graphs for the Exploitation of GINN Models

Since the inputs of the target function $\boldsymbol{F}$ are the capacity vectors $c$, which are defined on the edges of the graph $G$ and not on the nodes, we need to compute the line graph $L$ of $G$ in order to exploit the GINN models for the maximum-flow regression problem. We recall, here, the definition of line graph (see [32,33]).

Definition 7 (Line Graph). Let $G=(V, E)$ be a graph (either directed or not). The line graph of $G$ is a graph $L=\left(E, E^{\prime}\right)$, such that:

- $\quad$ The vertices of $L$ are the edges of $G$;
- Two vertices in $L$ are adjacent if the corresponding edges in $G$ share at least one vertex.

Given the line graph $L$ of the graph $G$ of a stochastic flow network $\mathcal{G}$, we can use the adjacency matrix $A_{L}$ of $L$ to define NN models characterized by GI layers to perform the maximum-flow regression task. See the next section for more details about the GINN architectures that are built.

### 3.3. Maximum-Flow Numerical Experiments

For the experiments related to the maximum-flow regression problem, we take into account two stochastic flow networks:

- $\mathcal{G}_{1}=\left(G_{\mathrm{BA}}, s, t, p\right)$. The graph $G_{\mathrm{BA}}$ of $\mathcal{G}_{1}$ characterizes a flow network built on an extended Barabási-Albert (BA) model graph [34,35]. Put simply, an extended BA model graph is a random graph generated using a preferential attachment criterion. This family of graphs describes a very common behavior in many natural and human systems, where few nodes are characterized by a higher degree if they are compared to the other nodes of the network.
In particular, we generate an extended BA undirected graph using the NetworkX Python module [36] (function extended_barabasi_albert_graph, input arguments $\mathrm{n}=50, \mathrm{~m}=2, \mathrm{p}=0.15$, and $\mathrm{q}=0.35$ ); then, we denote $t$ (the sink of the network) as the node with the highest betweenness centrality [37] and we add a new node $s$ (the source of the network) connected to the 10 nodes with smallest closeness centrality $[38,39]$. With these operations, we obtain a graph $G_{B A}$ of 51 nodes and $n=|E|=126$ edges, where the source $s$ is connected to the 10 nodes and the $\operatorname{sink} t$ is connected to the $m=15$ nodes (see Figure 3-left).
In the end, since, in real-world applications, truncated normal distributions seem to be very common (see Remark 7), in order to simulate a rather general maximum-flow regression problem, we chose a truncated normal distribution between 0 and 10, with a mean of 0 , and a standard deviation of $5 / 3$, as a probability distribution $p$ for the edge capacities (see Section 3.1.1); i.e.,

$$
\begin{equation*}
c_{i} \sim p \equiv \mathcal{N}_{[0,10]}(5,5 / 3), \quad \forall i=1, \ldots, n \tag{20}
\end{equation*}
$$

- $\mathcal{G}_{2}=\left(G_{\mathrm{ER}}, s, t, p\right)$. The graph $G_{\mathrm{ER}}$ of $\mathcal{G}_{2}$ characterizes a flow network built on an Erdős-Rényi (ER) model graph [40,41]. Put simply, an ER model graph is a random graph generated with a fixed number of nodes, where the edge $e_{i j}=\left(v_{i}, v_{j}\right)$ has a fixed probability of being created. This family of graphs is typically used to prove and/or find new properties that hold for almost all the graphs; for this reason, we consider a stochastic flow network based on an ER graph in our experiments.
In particular, we generate an ER undirected graph using the NetworkX Python module [36] (function fast_gnp_random_graph, input arguments $n=200, p=0.01$ ) and we select its largest connected component $G_{0}$ (in terms of the number of vertices). Then, we add to $G_{0}$ two new nodes: a node $s$ (the source of the network) connected to
all the nodes with degree equal to 1 , and a node $t$ (the sink of the network) connected to the 15 most distant nodes from $s$. With these operations, we obtain a graph $G_{E R}$ of 171 nodes and $n=|E|=269$ edges, where the source $s$ is connected to 37 nodes and the $\sin k t$ is connected to $m=15$ nodes (see Figure 3-right).
In the end, we chose the truncated normal distribution (20) as the probability distribution $p$ for the edge capacities.


Figure 3. Graph $G_{\mathrm{BA}}$ of $\mathcal{G}_{1}$ (left) and graph $G_{\mathrm{ER}}$ of $\mathcal{G}_{2}$ (right). In cyan, and with a circle around the source $s$, in magenta and with a circle around the $\operatorname{sink} t$, in green the nodes connected to $s$, and in red the nodes connected to $t$. All the other nodes are in blue.

Remark 7 (Regarding the truncated normal distribution for capacities). In a network describing a system of highroads, the capacity of a road is defined as $c=k \ell / S$ [42], where $k \in \mathbb{R}^{+}$is a value depending on the type of the road, $\ell$ is the road length, and $S$ is the average distance between two vehicles, typically chosen as a constant value. Then, assuming a network with all roads of the same type (i.e., $k$ constant) and a truncated normal distribution for the length $\ell$ of the roads, the capacity can be modeled as a random variable with a truncated normal distribution. Therefore, generalizing the concept of the highroad capacity to other similar problems (e.g., a network of pipes, a communication network, etc.) the distribution (20) can be considered sufficiently generic for the numerical experiments of this section.

Given the stochastic flow networks $\mathcal{G}_{1}$ and $\mathcal{G}_{2}$, the corresponding maximum-flow regression problems consist of the approximation of the functions $F_{1}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}, n=126$, $m=15$, and $F_{2}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}, n=269, m=15$, respectively, where $F_{1}$ and $\boldsymbol{F}_{2}$ are defined as in (16). For each $i=1,2$, we build the dataset $\mathcal{D}_{i}$ of $\mathcal{G}_{i}$ made of 10,000 pairs $\left(\boldsymbol{c}, \boldsymbol{\varphi}=\boldsymbol{F}_{i}(\boldsymbol{c})\right) \in \mathbb{R}^{n} \times \mathbb{R}^{m}$, where 3000 of them are used as the test set $\left(\mathcal{P}_{i}\right)$ and the remaining 7000 pairs are used to generate the multi-set $\Theta_{i}$ sampling $\vartheta \in\{1, \ldots, 7000\}$ pairs. In particular, $80 \%$ of the pairs in $\Theta_{i}$ are used as the training set $\left(\mathcal{T}_{i}\right)$ and the remaining $20 \%$ are used as the validation set $\left(\mathcal{V}_{i}\right)$. An important aspect of our numerical experiments consists of analyzing the performance of the trained NNs, varying the quantity of available data for the training process (i.e., $\vartheta$ ), and not only varying the hyper-parameters related to the architecture and optimization method; in particular, we study the NN performances when the number of training and validation data is $\vartheta=7000,1000,500$. Indeed, in real-world problems, the amount of available data can be limited for many reasons (e.g., limited computational resources for simulations, limited time for measurements, etc.). Then, studying the performances of a regression model while decreasing $\vartheta$ is important to understand the sample efficiency of the model.

The fluxes $F_{i}(\boldsymbol{c})=\boldsymbol{\varphi}$ for the dataset creation are computed using the maximum_flow NetworkX function that, specifically, allows the computatio of the flows for all the edges of the network (given the capacities $c$ ). Then, considering all the 10,000 simulations executed
to build the dataset $\mathcal{D}_{i}$, and denoting $\ell_{\max }^{(i)}(c)$ the length of the longest source-sink path in $\mathcal{G}_{i}(\boldsymbol{c})$, we observe that:

1. $\quad \ell_{\text {min }}\left(\mathcal{G}_{1}\right)=4, \ell_{\mathrm{av}}\left(\mathcal{G}_{1}\right) \approx 5.5$, and $\ell_{\max }\left(\mathcal{G}_{1}\right)=9$;
2. $\quad \ell_{\text {min }}\left(\mathcal{G}_{2}\right)=7, \ell_{\mathrm{av}}\left(\mathcal{G}_{2}\right) \approx 10.7$, and $\ell_{\text {max }}\left(\mathcal{G}_{2}\right)=17$;
where $\ell_{\min }\left(\mathcal{G}_{i}\right), \ell_{\mathrm{av}}\left(\mathcal{G}_{i}\right)$, and $\ell_{\max }\left(\mathcal{G}_{i}\right)$ are the minimum, the average, and the maximum lengths, respectively, of the longest source-sink path of the flow for $\mathcal{G}_{i}$ with respect to $\mathcal{D}_{i}$, i.e.,

$$
\begin{gather*}
\ell_{\min }\left(\mathcal{G}_{i}\right):=\min _{(\boldsymbol{c}, \boldsymbol{\varphi}) \in \mathcal{D}_{i}} \ell_{\max }^{(i)}(\boldsymbol{c}),  \tag{21}\\
\ell_{\mathrm{av}}\left(\mathcal{G}_{i}\right):=\frac{1}{\left|\mathcal{D}_{i}\right|} \sum_{(\boldsymbol{c}, \boldsymbol{\varphi}) \in \mathcal{D}_{i}} \ell_{\max }^{(i)}(\boldsymbol{c}), \tag{22}
\end{gather*}
$$

and

$$
\begin{equation*}
\ell_{\max }\left(\mathcal{G}_{i}\right):=\max _{(\boldsymbol{c}, \boldsymbol{\varphi}) \in \mathcal{D}_{i}} \ell_{\max }^{(i)}(\boldsymbol{c}) \tag{23}
\end{equation*}
$$

The values reported in items 1 and 2 show that, on the average, in a radius of the length $\ell_{\mathrm{av}}\left(\mathcal{G}_{i}\right)$ from the $\operatorname{sink} t$, it is likely to find almost all the nodes characterizing the maximumflow of the network $\mathcal{G}_{i}$. This information is then taken into account while choosing the depth values for the construction of the GINNs in the following Section 3.3.1. Indeed, we recall that the number of consecutive GI layers in an NN tells us if the input feature of node $v_{i}$ contributes to the computation of the output feature of node $v_{j}$ (see Proposition 1). Therefore, it is interesting to verify if the regression performance of a GINN improves or not, when the number of GI layers is related to one of the quantities (21)-(23).

### 3.3.1. NN Architectures, Hyper-Parameters, and Training

In the numerical experiments of this section, we study and compare the performances of MLPs and GINNs concerning the maximum-flow regression problems related to $F_{1}$ and $\boldsymbol{F}_{2}$. Then, we consider the two archetypes of NN architectures: an MLP archetype and a GINN archetype.

- MLP Archetype: The NN architecture is characterized by one input layer $L_{0}, H \in \mathbb{N}$, hidden layers $L_{1}, \ldots, L_{H}$ with a nonlinear activation function $f$, and one output layer $L_{H+1}$ with a linear activation function. The output layer is characterized by $m$ units, while all the other layers are characterized by $n$ units. Finally, we apply a batch normalization [43] before the activation function for each hidden layer $L_{1}, \ldots, L_{H}$. See Figure 4.
- GINN Archetype: The NN architecture is characterized by one input layer $L_{0}$ of $n$ units, $H \in \mathbb{N}$ hidden $G I$ layers $L_{1}^{G I}, \ldots, L_{H}^{G I}$ with a nonlinear activation function $f$, and one output layer $L_{H+1}^{G I}$ with a linear activation function. All the GI layers are built with respect to the adjacency matrix $A_{L} \in \mathbb{R}^{n \times n}$ of the line graph of the network (see Section 3.2.1) and they are characterized by $F \in \mathbb{N}$ filters (i.e., output features). Then, the number of input features $K$ of the GI layer $L_{h}^{G I}$ is $K=F$, if $h>1$, and $K=1$, if $h=1$. As for the MLP archetype, we apply a batch normalization before the activation function of each hidden layer. Finally, the output layer is characterized by a pooling operation and by the application of a mask (see Section 2.3) to focus on the $m$ units corresponding to the $m$ target flows. See Figure 5 .

For our experiments, given the two NN archetypes above, we built a set of untrained NN models, varying the main hyper-parameters of the architectures. In particular, for the MLPs, we varied the hyper-parameters $H$ and $f$ (i.e., the depth and activation functions of the hidden layers), while for the GINNs, we also varied $F$ (i.e., the number of filters of the GI layers) and the pooling operation. Specifically, the hyper-parameters vary among these values:

- MLP archetype. $f \in\{$ relu, elu, swish, softplus $\}$ and $H \in\{2,3,4,5\}$. We do not use deeper MLPs to avoid the so-called degradation problem [44], i.e., the problem in
which increasing the number of hidden layers causes the performance of an NN to saturate and degrade rapidly.
- GINN archetype. $f \in\{$ relu, elu, swish, softplus $\}, F \in\{1,5,10\}$, and pooling operations in $\{$ max, mean $\}$ (only if $F=5,10$ ); $H \in\{3,5,7,9\}$ for $\mathcal{G}_{1}$ and $H \in\{4,9,14,19\}$ for $\mathcal{G}_{2}$. In particular, we select these values of $H$ because they are a discrete interval around the value $\ell_{\mathrm{av}}\left(\mathcal{G}_{i}\right)$, also including cases near, or equal to, the minimum and maximum values $\ell_{\min }\left(\mathcal{G}_{i}\right)$ and $\ell_{\max }\left(\mathcal{G}_{i}\right)$, respectively.


Figure 4. MLP archetype. The units of the input layer $L_{0}$ are in green, the units of the hidden layers $L_{1}, \ldots, L_{H}$ are in purple, and the units of the output layer $L_{H+1}$ are in red.


Figure 5. Example of a GINN archetype with depth $H=2$ and max-pooling operation for the output layer. The output matrices $Y$ of the NN layers are in orange, the weight tensors $\mathbf{W}$ of the hidden GI layers are in red (see Definition 3), and the weight tensor $\mathbf{W}$ of the output GI layer with max-pooling and masking operations (see Section 2.3) are in purple.

Then, these models are all trained on $\vartheta=7000,1000,500$ input-output pairs sampled from $\mathcal{D}_{i}-\mathcal{P}_{i}$, using a mini-batch size $\beta=128,64,32$; the weight initialization is a Glorot normal distribution [45] for the MLPs and it varies among a Glorot normal and a normal distribution $\mathcal{N}(0,0.5)$ for the GINNs. All the biases are initialized as zeroes.

The remaining training options are fixed and shared by all the models during the training. In particular, these options are:

- Mean square error (MSE) loss, i.e.,

$$
\begin{equation*}
\operatorname{loss}(\mathcal{B}):=\frac{1}{m} \sum_{j=1}^{m}\left(\frac{1}{|\mathcal{B}|} \sum_{(c, \varphi) \in \mathcal{B}}\left(\widehat{\varphi}_{j}-\varphi_{j}\right)^{2}\right) \tag{24}
\end{equation*}
$$

where $\mathcal{B}$ is any generic batch of input-output pairs;

- The Adam optimizer [46] (learning rate $\epsilon=0.002$, moment decay rates $\rho_{1}=0.9$, $\rho_{2}=0.999$ );
- Early stopping regularization $[20,47]$ (200 epochs of patience, restore best-weights), to avoid overfitting;
- Learning rate reduction on plateau [47] (reduction factor $\alpha=0.5,100$ epochs of patience, minimum learning rate $\epsilon=10^{-6}$ ).
The training of all the NN models, with respect to all the different training configurations, returns 3168 trained NNs; in particular, we have 144 MLPs and 1140 GINNs, for each stochastic flow network $\mathcal{G}_{1}$ and $\mathcal{G}_{2}$.


### 3.3.2. Performance Analysis of Maximum-Flow Regression

To evaluate the performance of an NN trained with respect to the maximum-flow regression task, we consider the errors $\mathrm{MRE}_{a v}$ and $\operatorname{MRE}_{\varphi}$ (see Section 3.2 and Equations (18) and (19), respectively) measured on the test set. In particular, to better analyze the performances, we visualize the NNs as points in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane (see Figures 6 and 7). Then, the nearer a point is to the origin (i.e., the ideal zero-error NN ), the better the regression performances of the corresponding NNs are. We decide to use this representation because of the characteristics of the errors reported in Remark 6. Indeed, it is important to analyze the behavior of the NNs with respect to $\mathrm{MRE}_{a v}$ and $\mathrm{MRE}_{\varphi}$ together.

We start the analysis with the first stochastic flow network $\mathcal{G}_{1}$. In general, looking at Figure 6, we clearly see that the GINNs have better regression performances than the MLPs. In particular:

1. The $\operatorname{MRE}_{\varphi}$ of the GINNs is generally smaller than the MLPs, and this effect increases with $\vartheta$;
2. The $\mathrm{MRE}_{a v}$ of the GINNs is almost always smaller than the MLPs, and this effect seems to be almost stable while varying $\vartheta$;
3. Looking at the hyper-parameter $F$, we observe that the cases with $F=1,10$ generally perform better with fewer training samples (i.e., $\vartheta=1000,500$ ) while the cases with $F=5$ generally perform better with $\vartheta=7000$. This phenomenon suggests that increasing the number of filters can improve the quality of the training, even if a clear rule for the best choice of $F$ is not apparent.


Figure 6. Network $\mathcal{G}_{1}$. Scatter plots in the $\left(\operatorname{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane. Left to right: NNs trained with $\vartheta=7000,1000,500$ samples. Red circles: MLPs; green stars, blue crosses and purple " $x$ ": GINNs with $F=1,5,10$, respectively.

We continue the analysis with the second stochastic flow network $\mathcal{G}_{2}$, increasing the size and complexity of the flow network. Indeed, the graph $G_{\mathrm{BA}}$ of $\mathcal{G}_{1}$ is characterized by a reduced complexity of the maximum-flow problem, because the BA graphs are generated using a preferential attachment criterion that keeps the average length of the maximum source-sink path $\ell_{\mathrm{av}}\left(\mathcal{G}_{1}\right)$ small, even when increasing the nodes of the graph (this phenomenon was observed during some preliminary experiments).

Looking at Figure 7, we notice the same characteristics observed for $\mathcal{G}_{1}$, but much more emphasized. In particular, for each $\vartheta=7000,1000,500$, we clearly see that the GINNs
generally outperform the MLPs, especially with respect to the MRE ${ }_{a v}$. The reason for these similarities probably lies in the nature of the graphs $G_{\mathrm{BA}}$ and $G_{\mathrm{ER}}$ of $\mathcal{G}_{1}$ and $\mathcal{G}_{2}$, respectively; indeed, the ER graphs are used to represent generic graphs and are typically used to show properties that hold for almost all the graphs. On the other hand, the graph $G_{B A}$ is simpler than $G_{E R}$. Then, it is reasonable that the observations made for $\mathcal{G}_{1}$ are confirmed looking at $\mathcal{G}_{2}$ and it is reasonable that the performance differences observed in $\mathcal{G}_{2}$ are less emphasized in $\mathcal{G}_{1}$, since the maximum-flow problem on $G_{B A}$ is less complex than on $G_{E R}$.


Figure 7. Network $\mathcal{G}_{2}$. Scatter plots in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane. Left to right: NNs trained with $\vartheta=7000,1000,500$ samples. Red circles: MLPs; green stars, blue crosses and purple " $x$ ": GINNs with $F=1,5,10$, respectively.

Remark 8 (GINNs, small MRE $_{a v}$, and regressions on graphs). We have just observed that the GINNs genrally perform better than MLPs for regression tasks on graphs but, if we focus on the $M R E_{a v}$ values, the GINNs clearly show better performances (see Figures 6 and 7 and Tables 1 and 2). Specifically, Tables 1 and 2 show the three GINNs and MLPs with lowest MRE av value on the test sets of $\mathcal{G}_{1}$ and $\mathcal{G}_{2}$, respectively. The better performances of the GINNs, with respect to this error measure, are particularly important if we extend the regression problem, e.g., if we want to learn all the flow values $\varphi^{*(c)}\left(e_{1}\right), \ldots, \varphi^{*(c)}\left(e_{n}\right)$ on the edges of the graph and not only the ones characterizing the maximum-flow $\varphi$ reaching the sink. Indeed, in this case, an NN with small errors on the single elements of the target vector is fundamental, while an NN with small errors on the sum of the elements of the target vector is useless; for this reason, in vector-valued regression tasks, we use loss functions such as (24) (evidently similar to the performance measure MRE ${ }_{a v}$ ). In conclusion, we believe that the GINNs have great potential in the field of regression graphs.

We continue to conclude the performance analysis, and we analyze how the errors of the NN models change when varying the hyper-parameters.

The first observation is related to the activation functions and the mini-batch size. We observe that the trained NN models (both GINNs and MLPs) generally exhibit worse regression performances with a relu activation function and a mini-batch size equal to 128 ; in Figure 8, we report the same scatterplots of Figures 6 and 7 but without the points corresponding to the NNs with the relu activation function or a mini-batch size equal to 128 . It is worth noting that the general observations concerning the NN performances are even more evident when removing these models. Moreover, among the remaining models, we do not observe activation functions or mini-batch size values that are evidently better than the others; in general, as can be expected, we only observe a few advantages of using a mini-batch size of 32 samples instead of a mini-batch size of 64 samples, while decreasing $\vartheta$.

From now on, we do not include in our analyses the models characterized by a relu activation function or mini-batch size equal to 128. Moreover, we only focus on the GINN models and, in particular, on their performances with respect to the hyper-parameter $H$, characterizing the number of hidden layers. Indeed, the remaining hyper-parameters (pooling operations and weight initializations) do not seem to have a particular impact on the results.

The study of the GINN performances with respect to $H$ is particularly interesting if we consider Proposition 1. In fact, from this proposition, we expect the GINN models to
have a better performance if the depth $H$ is such that $H+1 \gtrsim \ell_{\mathrm{av}}\left(\mathcal{G}_{i}\right)$. This guess is indeed satisfied. In particular, for $\mathcal{G}_{1}$, we see that by increasing the depth $H$, we obtain GINNs with better performances in general (see Figure 9). On the other hand, for $\mathcal{G}_{2}$ we observe a slightly different behavior. The GINNs that are sufficiently deep (i.e., $H \geq 9$ ) show better performances than the GINNs with $H=4$, but their errors tend to increase with a small $\vartheta$; in particular, the more $H$ is greater than $\ell_{\mathrm{av}}\left(\mathcal{G}_{2}\right)$, the more the GINN performances seem to deteriorate (see Figure 10). To summarize, the depth in a GINN model is very important to obtaining good regression abilities, keeping in mind Proposition 1. Nonetheless, the practice of using as much of a GI layer as possible is not always the best choice, and this topic deserves attention in future work.

Table 1. Network $\mathcal{G}_{1}$. Top three GINNs and MLPs, for $\vartheta=7000,1000,500$. Models are sorted with respect to the $\mathrm{MRE}_{a v}$ error; the "rank" column describes their global position with respect to all the other models.

| $\mathcal{G}_{1}$ | GINNs |  |  |  |  |  |  |  | MLPs |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\vartheta$ | Rank | $\mathrm{MRE}_{\text {av }}$ | H | $F$ | $f$ | $\beta$ | Pool. | Init. | Rank | $\mathrm{MRE}_{a v}$ | H | $f$ | $\beta$ |
| 7000 | 1/528 | 0.00707 | 9 | 1 | elu | 64 | - | G.Norm. | 446/528 | 0.00914 | 3 | swish | 32 |
|  | 2/528 | 0.00712 | 9 | 1 | elu | 128 | - | Norm. | 453/528 | 0.00934 | 3 | swish | 64 |
|  | 3/528 | 0.00713 | 9 | 1 | elu | 32 | - | Norm | 455/528 | 0.00939 | 4 | swish | 32 |
| 1000 | 1/528 | 0.00933 | 9 | 1 | softplus | 32 | - | G.Norm. | 338/528 | 0.01272 | 5 | softplus | 32 |
|  | 2/528 | 0.00940 | 7 | 1 | elu | 32 | - | Norm. | 353/528 | 0.01289 | 4 | elu | 32 |
|  | 3/528 | 0.00952 | 7 | 1 | elu | 32 | - | G.Norm. | 357/528 | 0.01292 | 5 | elu | 32 |
| 500 | 1/528 | 0.01056 | 7 | 1 | softplus | 32 | - | Norm. | 263/528 | 0.01433 | 5 | softplus | 32 |
|  | 2/528 | 0.01077 | 5 | 1 | relu | 32 | - | G.Norm. | 279/528 | 0.01448 | 5 | softplus | 64 |
|  | 3/528 | 0.01092 | 7 | 1 | softplus | 32 | - | G.Norm. | 284/528 | 0.01452 | 4 | softplus | 32 |

Table 2. Network $\mathcal{G}_{2}$. Top three GINNs and MLPs, for $\vartheta=7000,1000,500$. Models are sorted with respect to the $\mathrm{MRE}_{a v}$ error; the "rank" column describes their global position with respect to all the other models.

| $\mathcal{G}_{1}$ | GINNs |  |  |  |  |  |  |  | MLPs |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\vartheta$ | Rank | $\mathrm{MRE}_{a v}$ | H | $F$ | $f$ | $\beta$ | Pool. | Init. | Rank | $\mathrm{MRE}_{a v}$ | H | $f$ | $\beta$ |
| 7000 | 1/528 | 0.00087 | 14 | 10 | softplus | 32 | max | G.Norm. | 442/528 | 0.00557 | 2 | elu | 32 |
|  | 2/528 | 0.00092 | 14 | 10 | softplus | 32 | mean | G.Norm. | 445/528 | 0.00571 | 5 | softplus | 64 |
|  | 3/528 | 0.00099 | 12 | 10 | softplus | 32 | mean | G.Norm | 446/528 | 0.00573 | 5 | elu | 32 |
| 1000 | 1/528 | 0.00465 | 9 | 1 | elu | 32 | - | G.Norm. | 263/528 | 0.01038 | 3 | softplus | 32 |
|  | 2/528 | 0.00478 | 19 | 1 | elu | 32 | - | G.Norm. | 264/528 | 0.01038 | 2 | swish | 32 |
|  | 3/528 | 0.00479 | 14 | 1 | softplus | 32 | - | Norm. | 267/528 | 0.01043 | 3 | softplus | 64 |
| 500 | 1/528 | 0.00593 | 14 | 1 | elu | 32 | - | G.Norm. | 114/528 | 0.01266 | 2 | swish | 32 |
|  | 2/528 | 0.00674 | 14 | 1 | softplus | 32 | - | G.Norm. | 118/528 | 0.01272 | 2 | swish | 64 |
|  | 3/528 | 0.00688 | 19 | 1 | softplus | 32 | - | G.Norm. | 119/528 | 0.01272 | 5 | softplus | 32 |



Figure 8. Scatter plots in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane for NNs trained with respect to $\mathcal{G}_{1}$ (top) and $\mathcal{G}_{2}$ (bottom). Left to right: NNs trained with $\vartheta=7000,1000,500$ samples. Red circles: MLPs; green stars, blue crosses and purple " $x$ ": GINNs with $F=1,5,10$, respectively. We do not plot results with NNs that have a relu activation function or a mini-batch size equal to 128 .

Remark 9 (Training time). Along the conducted experiments (with training that occurred more than 3000 times), we find that the average training time for the GINN models is approximately 20 min in total, and one second per epoch; on the other hand, the average training time for the MLP models is approximately 10 min in total and half a second per epoch. Nonetheless, we point out that the difference in training times between GINNs and MLPs can be reduced with code optimization. Indeed, the GINN layers are a custom class of TensorFlow NN layers developed on purpose by the authors for numerical experiments. The code of TensorFlow's FC layers is extremely optimized. Therefore, at the present time, the GINNs and MLPs cannot be compared in terms of computational costs. More details about the average training times per epoch of the models are reported in Table 3, and this quantity is indicative of the training computational cost of the NNs. However, we recall that the experiments take into account more than three thousand models, each one with a different training configuration that characterizes the training time per epoch. All the training was performed on a workstation with a CPU of 4 Core and 8 Threads, 32 GB of RAM, and a GPU Nvidia 1080 $8 G B$.

Table 3. Global statistics of the average training time per epoch for GINN and MLP models, expressed in seconds.

## Avg. Time per Epoch (s)

|  | GINNs | MLPs |
| :---: | :---: | :--- |
| Mean | 1.099 | 0.565 |
| Std | 1.359 | 0.292 |
| 25th perc. | 0.318 | 0.380 |
| 50th perc. | 0.567 | 0.431 |
| 75th perc. | 1.296 | 0.632 |



Figure 9. Scatter plots in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane for GINNs trained with respect to $\mathcal{G}_{1}$. Left to right: GINNs trained with $\vartheta=7000,1000,500$ samples; top to bottom: red markers highlight GINNs with $H=3,5,7,9$ (black markers for all the other models).


Figure 10. Scatter plots in the $\left(\operatorname{MRE}_{a v}, \operatorname{MRE}_{\varphi}\right)$ plane of the GINNs trained with respect to $\mathcal{G}_{2}$. From left to right, the GINNs are trained using $\vartheta=7000,1000,500$ samples; from top to bottom, the red markers highlight the GINNs with hyper-parameters of $H=4,9,14,19$ (black markers for all the other models).

### 3.4. GINNs for Flux Regression in Discrete Fracture Networks

In Section 3.3, we showed the regression abilities and the potentialities of the GINN models for the maximum-flow regression problem, i.e., for a problem representative of generic real-world applications.

In this section, we address a specific real-world application where GINNs can be useful; in particular, we consider an uncertainty quantification (UQ) problem related to underground flows in fractured media. Flow characterization in underground fractured media is an interesting problem for many applications, such as civil engineering, industrial engineering, and environmental analyses. A helpful model to describe the flow in an
underground network of fractures is represented by the discrete fracture network (DFN) models [48-50]. These models express the fractures as two-dimensional polygons into a three-dimensional domain, and each fracture is described by specific hydro-geological properties (e.g., fracture transmissivity) and geometrical properties (e.g., barycenter position and orientation). Intersections between fractures, denominated "traces", define flux exchange phenomenons, such that the flow model is typically ruled by the following assumptions: (i) the rock matrix is impenetrable; (ii) the Darcy law is used for the flux propagation on each fracture; and (iii) the head continuity and flux balance is imposed on all the traces. However, since the hydro-geological and geometrical characteristics of the network of fractures are usually not accessible in detail, the DFN models are typically generated by sampling their hydro-geological and geometrical features from known distributions [51-54]. Therefore, a statistical approach is required to study real fractured media, resorting to UQ analyses.

In recent literature, several new methods have been proposed to reduce the computational costs of the DFN flow simulations (e.g., see [55-57]); nonetheless, they are still computationally expensive in many situations and the UQ analyses can involve thousands of these simulations. Therefore, it is fundamental to take into account the techniques for the complexity reduction, such as machine learning-based techniques (e.g., see [16,17,58]). In particular, in [16,17], NN models are trained on datasets built using DFN simulations to provide surrogate models; finally, in a negligible amount of time, the NN models are used to generate a large set of approximated DFN flow simulation results, which are particularly useful to speed up the UQ analyses.

In this section, the idea is to take advantage of the DFN's graph structure to build GINN models and to analyze the advantages of using such models for the DFN flux regression tasks instead of the more classic NN architectures (e.g., MLPs or multi-task NNs).

### 3.4.1. The DFN Model and the Flux-Regression Task in DFNs

Here, for the reader's convenience, we briefly describe the problem of the flow simulations in DFNs. We point the interested reader to [55-57] for full details.

A DFN model is a discrete representation of an underground network of fractures in a fractured rock medium using a set of intersecting planar polygons (the fractures) in a three-dimensional domain $\mathbb{D} \subset \mathbb{R}^{3}$ (the rock medium); see Figure 11 for a DFN example. Each fracture (i.e., polygon) is labeled and identified by an index belonging to an arbitrary set $I$; i.e., for each $i \in I$, we denote each fracture by $\mathcal{F}_{i}$. Then, a DFN is defined as the union of all the fractures, i.e., $\cup_{i \in I} \mathcal{F}_{i}$. Since the traces, i.e., the segments obtained from the intersection of two or more fractures that connect the fractures define a network, a DFN can be represented as a graph where the fractures are the nodes and the traces are edges (see Figure 11).


Figure 11. External surface of a natural fractured medium (left), a 3D view of a DFN (center), and a graph representation of the same DFN (right). The DFN illustrated in this figure is DFN158, i.e., the one used for the numerical tests of Section 3.4.

Each fracture in a DFN is characterized by a transmissivity parameter $\kappa_{i}$ and by geometrical properties, such as the size and the orientation in $\mathbb{R}^{3}$. These characteristics are generally sampled from known probability distributions, and they characterize the flow simulations of the DFN. Indeed, a DFN flow simulation depends both on the geometry and the hydro-geological properties of the fractures, such as the transmissivities $\kappa_{i}$ of the fractures.

The DFN flux regression problem addressed in this section is characterized as follows. The DFN considered consists of $n=158$ fractures and we donote it by DFN158. The geometry of DFN158 is fixed, and the fractures are immersed into a domain $\mathbb{D}$ described by a cube of 1000 m along the long edge (see Figure 11). The domain boundary conditions are such that a fixed pressure difference $\Delta \mathrm{H}=10 \mathrm{~m}$ is imposed between two opposite faces of $\mathbb{D}$, therefore representing an inlet and outlet face, respectively. Moreover, those fractures cut by the inlet and outlet faces of the domain are defined as inflow and outflow fractures, respectively. The edges of all the other fractures are insulated (homogeneous Neumann condition). It is possible to observe that the fixed geometry and the given boundary conditions affect the flux directionality in DFN158; on the other hand, the fracture transmissivities characterize the flow intensity exiting from the outflow fractures. In particular, we assume that the fracture transmissivities of the fractures $\mathcal{F}_{1}, \ldots, \mathcal{F}_{n}$ are isotropic parameters $\kappa_{1}, \ldots, \kappa_{n}$, respectively, described by random variables with a lognormal distribution [51,52]:

$$
\log _{10} \kappa_{i} \sim \mathcal{N}(-5,1 / 3), \quad \forall i=1, \ldots, n
$$

In the test case considered, we use octagons to represent the fractures and they have been randomly created according to the distributions in [53,54]. For their geometrical properties: a truncated power law distribution for the fracuture radii, with an exponent $\gamma=2.5$ and an upper and lower cut-off $r_{u}=560$ and $r_{0}=50$, respectively; a Fischer distribution for fracture orientations, with a mean direction $\boldsymbol{\mu}=(0.0065,-0.0162,0.9998)$ and a dispersion parameter of 17.8; and a uniform distribution for the mass centers of the fractures. However, the first eight fractures $\mathcal{F}_{1}, \ldots, \mathcal{F}_{8}$ are not randomly sampled but are created to ensure an inlet-outlet path for the flow, where $\mathcal{F}_{1}$ is an inflow fracture and $\mathcal{F}_{8}$ is an outflow fracture.

Together with the boundary conditions described above, the geometry of DFN158 is such that we have $m=7$ outflow fractures in DFN158. Therefore, we can define a flux function representing the DFN flow simulations for DFN158; i.e., a function $F: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ such that, for each vector $\kappa=\left[\kappa_{1}, \ldots, \kappa_{n}\right]^{\top} \in \mathbb{R}^{n}$ of fracture transmissivities with $n=158$, it returns the vector $\boldsymbol{F}(\boldsymbol{\kappa})=\boldsymbol{\varphi}=\left[\varphi_{1}, \ldots, \varphi_{m}\right]^{\top} \in \mathbb{R}^{m}$ of fluxes exiting from the $m$ outflow fractures, $m=7$.

Now, assuming the need to perform UQ analyses of the fluxes $\varphi_{1}, \ldots, \varphi_{m}$ of DFN158 while varying the fracture transmissivities, it is worth considering the flux regression problem that looks for NN-based approximations $\widehat{\boldsymbol{F}}$ of $\boldsymbol{F}$.

### 3.4.2. Performance Analysis of DFN Flux Regression

In this subsection, we extend the numerical experiments and analyses of Section 3.3 to the example represented by DFN158. Then, we study and compare the performances of the MLPs and GINNs concerning the flux regression problem related to the function $F: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ described in Section 3.4.1. The two archetypes of MLP and GINN architectures are the same as in Section 3.3.1, as are most of the hyper-parameter values and the training options used; the only differences are the following:

- $\vartheta=1000,500$ (number of training and validation data);
- $\beta=64,32$ (mini-batch size);
- The relu activation function is not considered in the experiments;
- For the GINN models, we consider the depth parameter values $H \in\{4,7,9,14,19\}$. The rationale behind this choice is that it is a set of values around 8, which is the number of deterministic fractures that, on average, represent an inlet-outlet flow path
for DFN158 (in the absence of a value equivalent to $\ell_{\mathrm{av}}$ that cannot be easily computed for DFN158);
- The GI layers are built with respect to the adjacency matrix $A$ of DFN158; indeed, we do not need to introduce the line graph of the network since the features (i.e., the transmissivities) are assigned to the nodes of the graph and not to the edges.
As in Section 3.3.2, we evaluate the performance of the NNs trained with respect to the DFN flux regression task measuring the errors $\mathrm{MRE}_{a v}$ and $\mathrm{MRE}_{\varphi}$ on the test set (also, in this case, $\mathcal{P}$ is made of 3000 samples). Then, we visualize the results of the NNs as points in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane (see Figure 12). Analyzing the error values and looking at the scatter plots of Figure 12, we clearly observe that the GINN models outperform the MLPs, and that they are characterized by more regular error behaviors than the GINNs trained for the maximum-flow regression task, with respect to the filter hyper-parameter (see Section 3.3.2). In particular:

1. Both the $\operatorname{MRE}_{\varphi}$ and the $\mathrm{MRE}_{a v}$ of the GINNs are almost always smaller than the ones of the MLPs, independently of $\vartheta$;
2. Looking at the filter hyper-parameter $F$, we observe that the GINN performances are better as $F$ increases (from $F=1$, to $F=5$, to $F=10$ ).


Figure 12. Scatter plots in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane for NNs trained with respect to DFN158. NNs are trained using $\vartheta=1000$ (left) and $\vartheta=500$ (right) samples. Red circles: MLPs; green stars, blue crosses and purple " $x$ ": GINNs with $F=1,5,10$, respectively.

We continue the analysis of studying the relationships between the GINN errors and the other hyper-parameters of the models. With respect to the activation functions, we observe that the GINN models with the relu activation functions have, in general, slightly better performances than other models; on the other hand, all the GINN models with the worst performances (i.e., the points in the top-right corners of Figure 12) have softplus activation functions. Concerning the mini-batch size $\beta$, we observe that the GINNs with the best performances (corresponding to points in the bottom-left corners of Figure 12) are trained with $\beta=32$. Similar results hold for the weights initialization, where the best performing GINNs are initialized with a Glorot normal distribution. Concerning the pooling operations, we do not observe particular differences in the error values of GINN models using a max-pooling or a mean-pooling operation.

Analogous to Section 3.3.2, we conclude with a focus on the error behaviors with respect to the depth $H$ of the GINN models. Moreover, for the DFN flux regression task, we observe that the depth $H$ of the model can improve the regression quality. In particular, for each $\vartheta=1000,500$, we observe that the best performances are obtained by the GINNs with a depth of $H=7,9,14$, while both the shallowest and the deepest GINNs $(H=4,19)$ have higher errors (see Figures 13 and 14). In accordance with Proposition 1 and the observations of Section 3.3.2, this characteristic let us deduce that, on average, the maximum inlet-outlet flux path in DFN158 is probably made of 8 to 15 fractures; i.e., a value not so far from the length of the inlet-outlet path defined by the fractures $\mathcal{F}_{1}, \ldots, \mathcal{F}_{8}$.


Figure 13. Scatter plots in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane for GINNs trained with respect to DFN158, $\vartheta=1000$. Left to right, top to bottom: red markers highlight GINNs trained with $H=4,7,9,14,19$, respectively (black markers for all the other models).


Figure 14. Scatter plots in the $\left(\mathrm{MRE}_{a v}, \mathrm{MRE}_{\varphi}\right)$ plane for GINNs trained with respect to DFN158, $\vartheta=500$. Left to right, top to bottom: red markers highlight GINNs trained with $H=4,7,9,14,19$, respectively (black markers for all the other models).

## 4. Conclusions

In this work, we presented the graph-informed (GI) layers, a new type of spatialbased graph convolutional layer, designed for regression tasks on graph-structured data. With respect to the other types of GCN layers, our GI layers stand out for their tensor formulation (see (8)), managing multiple input/output features; moreover, these layers let users build deep NN architectures which are able to exploit the depth needed to improve
the regression performances (see Proposition 1 and Sections 3.3.2 and 3.4.2). The GI layers have been formally defined, from the simplest version to the most general and tensor version. Moreover, additional optional operations are introduced: the pooling operation and mask operations.

To study the regression abilities of graph-informed NNs (GINNs), i.e., NNs made from GI layers, we trained thousands of NN models (both GINNs and MLPs) on two maximum-flow regression problems, with networks based on a Barabási-Albert graph ( $\mathcal{G}_{1}$ ) and an Erdős-Rényi graph $\left(\mathcal{G}_{2}\right)$. We selected the maximum-flow regression problem as a representative test since it is a sufficiently general problem to demonstrate the applications in many topics of the network analysis. Analyzing the approximation errors of the NNs, we observed that the GINNs have better performances than the MLPs. in general. In particular, for $\mathcal{G}_{2}$, the GINNs in almost all the cases outperform the MLPs. The study of the regression performances also showed an interesting relationship between small errors and a depth greater than, or equal to, the average length of the maximum source-sink path in the stochastic network.

After the test on the maximum-flow regression task, we illustrated an example of a possible application of the GINNs to a real-world problem: a DFN flux regression problem, i.e., an uncertainty quantification problem for the characterization of the exiting flux distribution of an underground network of fractures. In this practical application, the GINN models completely outperform the MLPs; moreover, both the depth and the filter hyper-parameters of the GINNs proved to be significant enough to improve the approximation quality of the target function.

In conclusion, we believe that our work introduces a new, useful, contribution to the family of spatial-based graph convolutional networks; indeed, the numerical experiments illustrated here show that the new GI layers and the GINNs have great potentialities in the framework of regression tasks on graph-structured data.

Author Contributions: Conceptualization, S.B., F.D.S., A.M., S.P., F.V.; data curation, S.B., F.D.S., A.M., S.P., F.V.; formal analysis, S.B., F.D.S., A.M., S.P., F.V.; funding acquisition, S.B., S.P., F.V.; investigation, S.B., F.D.S., A.M., S.P., F.V.; methodology, S.B., F.D.S., A.M., S.P., F.V.; project administration, S.B., F.D.S., A.M., S.P., F.V.; resources, S.B., F.D.S., A.M., S.P., F.V.; software, F.D.S., A.M.; supervision, S.B., S.P., F.V.; validation, S.B., F.D.S., A.M., S.P., F.V.; visualization, F.D.S., A.M.; writing-original draft, S.B., F.D.S., A.M., S.P., F.V.; writing-review and editing, S.B., F.D.S., A.M., S.P., F.V. All authors have read and agreed to the published version of the manuscript.

Funding: Research performed in the framework of the Italian MIUR Award "Dipartimento di Eccellenza 2018-2022" to the Department of Mathematical Sciences, Politecnico di Torino, CUP: E11G18000350001. F.D.S. and S.P. also acknowledge support from Italian MIUR PRIN project 201752HKH8_003. A.M. gratefully acknowledges the support from Addfor Industriale. The research leading to these results was also partially funded by the SmartData@PoliTO center for Big Data and Machine Learning technologies.

Institutional Review Board Statement: Not applicable.
Informed Consent Statement: Not applicable.
Data Availability Statement: The code used to implement the Graph-Informed Neural Networks is available at https:/ / github.com/Fra0013To/GINN (accessed on 31 January 2022).

Acknowledgments: The authors acknowledge support from the GEOSCORE group (https:/ / areeweb. polito.it/geoscore/, accessed on 1 February 2022) of Politecnico di Torino (Department of Mathematical Sciences).

Conflicts of Interest: The authors declare no conflict of interest.

## Abbreviations

The following abbreviations and nomenclatures are used in this manuscript:

| BA | Barabási-Albert |
| :--- | :--- |
| CNN | Convolutional Neural Network |
| CwD | Circulation with Demand |
| DCNN | Diffusion-Convolutional Neural Networks |
| DFN | Discrete Fracture Network |
| DL | Deep Learning |
| ER | Erdős-Rényi |
| FC | Fully-Connected |
| GCN | Graph Convolutional Network |
| GI | Graph-Informed |
| GINN | Graph-Informed Neural Network |
| GNN | Graph Neural Network |
| ML | Machine Learning |
| MLP | Multi-Layer Perceptron |
| MRE | Mean Relative Error |
| MRE | Edge-Wise Average MRE |
| MRE | MRE on the Predicted Maxflow/Outflow |
| MSE | Mean Square Error |
| NIM | Network Interdiction Model |
| NN | Neural Network |
| UQ | Uncertainty Quantification |

## Appendix A. Multi-Sets

Definition A1 (Multi-set $[59,60])$. A multi-set $\mathcal{A}$ is a collection of objects, called elements, which may occur more than once. The number of times an element occurs in a multi-set is called its multiplicity. The cardinality of a multi-set (denoted by $|\mathcal{A}|$ ) is the sum of the multiplicities of its elements.

In other words, a multi-set may be formally defined as a pair $\mathcal{A}=(A, m)$, where $A$ is the underlying set of the multi-set, formed from its distinct elements, and $m: A \rightarrow \mathbb{Z}^{+}$is a function that, for each $a \in A$, returns the multiplicity $m(a) \geq 1$ of $a$ in the multi-set.

Definition A2 (Relations and Operations with Multi-sets [60]). The usual relations and operations on sets can be extended to multi-sets by considering the multiplicity function. Let $\mathcal{A}=\left(A, m_{A}\right)$ and $\mathcal{B}=\left(B, m_{B}\right)$ be two multi-sets; then, we can define the following relations and operations.

- Equality: $\mathcal{A}$ is equal to $\mathcal{B}(\mathcal{A}=\mathcal{B})$ if $A=B$ and $m_{A}(a)=m_{B}(a)$, for each $a \in A$.
- Inclusion: $\mathcal{A}$ is included in $\mathcal{B}(\mathcal{A} \subset \mathcal{B})$ if $A \subset B$ and $m_{A}(a)<m_{B}(a)$, for each $a \in A$. Analogously, $\mathcal{A}$ is included in, or equal to, $\mathcal{B}(\mathcal{A} \subseteq \mathcal{B})$ if $A \subseteq B$ and $m_{A}(a) \leq m_{B}(a)$, for each $a \in A$.
- Intersection: the intersection of $\mathcal{A}$ and $\mathcal{B}(\mathcal{A} \cap \mathcal{B})$ is a multi-set $\mathcal{C}=\left(C, m_{C}\right)$, such that $C=A \cap B$ and $m_{C}(c)=\min \left\{m_{A}(c), m_{B}(c)\right\}$, for each $c \in C$.
- Union: the union of $\mathcal{A}$ and $\mathcal{B}(\mathcal{A} \cup \mathcal{B})$ is a multi-set $\mathcal{C}=\left(C, m_{C}\right)$, such that $C=A \cup B$ and $m_{C}(c)=\max \left\{m_{A}(c), m_{B}(c)\right\}$, for each $c \in C$.
- Sum: the sum of $\mathcal{A}$ and $\mathcal{B}(\mathcal{A}+\mathcal{B})$ is a multi-set $\mathcal{C}=\left(C, m_{C}\right)$, such that $C=A \cup B$ and $m_{C}(c)=m_{A}(c)+m_{B}(c)$, for each $c \in C$.
- Difference: the difference of $\mathcal{A}$ and $\mathcal{B}(\mathcal{A}-\mathcal{B})$ is a multi-set $\mathcal{C}=\left(C, m_{C}\right)$, such that $\mathcal{A}=$ $\mathcal{C}+\mathcal{B}$.


## References

1. Brandes, U.; Erlebach, T. (Eds.) Network Analysis—Methodological Foundations; Theoretical Computer Science and General Issues; Springer: Berlin/Heidelberg, Germay, 2005; Volume 3418. [CrossRef]
2. Gori, M.; Monfardini, G.; Scarselli, F. A new model for learning in graph domains. In Proceedings of the 2005 IEEE International Joint Conference on Neural Networks, Montreal, QC, Canada, 31 July-4 August 2005; Volume 2, pp. 729-734. [CrossRef]
3. Micheli, A. Neural Network for Graphs: A Contextual Constructive Approach. IEEE Trans. Neural Netw. 2009, 20, 498-511. [CrossRef]
4. Scarselli, F.; Gori, M.; Tsoi, A.C.; Hagenbuchner, M.; Monfardini, G. The Graph Neural Network Model. IEEE Trans. Neural Netw. 2009, 20, 61-80. [CrossRef]
5. Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; Yu, P.S. A Comprehensive Survey on Graph Neural Networks. IEEE Trans. Neural Netw. Learn. Syst. 2021, 32, 4-24. [CrossRef]
6. Bruna, J.; Zaremba, W.; Szlam, A.; LeCun, Y. Spectral networks and locally connected networks on graphs. In Proceedings of the International Conference on Learning Representations, Banff, AB, Canada, 14-16 April 2014.
7. Defferrard, M.; Bresson, X.; Vandergheynst, P. Convolutional neural networks on graphs with fast localized spectral filtering. In Advances in Neural Information Processing Systems; Curran Associates, Inc.: Red Hook, NY, USA, 2016; Volume 29, pp. 3844-3852.
8. Kipf, T.; Welling, M. Semi-supervised classification with graph convolutional networks. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24-26 April 2017.
9. Hamilton, W.; Ying, Z.; Leskovec, J. Advances in Neural Information Processing Systems; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30.
10. Monti, F.; Boscaini, D.; Masci, J.; Rodolà, E.; Svoboda, J.; Bronstein, M.M. Geometric deep learning on graphs and manifolds using mixture model CNNs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21-26 July 2017; pp. 5115-5124.
11. Niepert, M.; Ahmed, M.; Kutzkov, K. Learning Convolutional Neural Networks for Graphs. In Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, $20-22$ June 2016; Volume 48, pp. 2014-2023.
12. Gao, H.; Wang, Z.; Ji, S. Large-Scale Learnable Graph Convolutional Networks. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery \& Data Mining, London, UK, 19-23 August 2018. [CrossRef]
13. Li, Q.; Han, Z.; Wu, X.M. Deeper Insights Into Graph Convolutional Networks for Semi-Supervised Learning. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2-7 February 2018; Volume 32.
14. Kleinberg, J.; Tardos, E. Algorithm Design; Addison-Wesley: Boston, MA, USA, 2005.
15. Dimitrov, N.B.; Morton, D.P. Interdiction Models and Applications. In Handbook of Operations Research for Homeland Security; Herrmann, J.W., Ed.; Springer: New York, NY, USA, 2013; pp. 73-103. [CrossRef]
16. Berrone, S.; Della Santa, F.; Pieraccini, S.; Vaccarino, F. Machine learning for flux regression in discrete fracture networks. GEM-Int. J. Geomath. 2021, 12, 9. [CrossRef]
17. Berrone, S.; Della Santa, F. Performance Analysis of Multi-Task Deep Learning Models for Flux Regression in Discrete Fracture Networks. Geosciences 2021, 11, 131. [CrossRef]
18. Berrone, S.; Della Santa, F.; Mastropietro, A.; Pieraccini, S.; Vaccarino, F. Discrete Fracture Network insights by eXplainable AI. Machine Learning and the Physical Sciences. In Proceedings of the Workshop at the 34th Conference on Neural Information Processing Systems (NeurIPS), Neural Information Processing Systems Foundation, Virtual, 6-7 December 2002. Available online: https:/ /ml4physicalsciences.github.io/2020/ (accessed on 1 February 2022).
19. Berrone, S.; Della Santa, F.; Mastropietro, A.; Pieraccini, S.; Vaccarino, F. Layer-wise relevance propagation for backbone identification in discrete fracture networks. J. Comput. Sci. 2021, 55, 101458. [CrossRef]
20. Goodfellow, I.; Bengio, Y.; Courville, A. Deep Learning; MIT Press: Cambridge, MA, USA, 2016.
21. Atwood, J.; Towsley, D. Diffusion-Convolutional Neural Networks. In Advances in Neural Information Processing Systems; Curran Associates, Inc.: Red Hook, NY, USA, 2016; Volume 29, pp. 1993-2001.
22. Donon, B.; Donnot, B.; Guyon, I.; Marot, A. Graph Neural Solver for Power Systems. In Proceedings of the International Joint Conference on Neural Networks, Budapest, Hungary, 14-19 July 2019; pp. 1-8. [CrossRef]
23. Donon, B.; Clément, R.; Donnot, B.; Marot, A.; Guyon, I.; Schoenauer, M. Neural networks for power flow: Graph neural solver. Electr. Power Syst. Res. 2020, 189, 106547. [CrossRef]
24. Raissi, M.; Karniadakis, G.E. Hidden physics models: Machine learning of nonlinear partial differential equations. J. Comput. Phys. 2018, 357, 125-141. [CrossRef]
25. Raissi, M.; Perdikaris, P.; Karniadakis, George, E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J. Comput. Phys. 2019, 378, 686-707. [CrossRef]
26. Ding, S. The $\alpha$-maximum flow model with uncertain capacities. Appl. Math. Model. 2015, 39, 2056-2063. [CrossRef]
27. Malhotra, V.; Kumar, M.; Maheshwari, S. An O(| V | 3) algorithm for finding maximum flows in networks. Inf. Process. Lett. 1978, 7, 277-278. [CrossRef]
28. Goldberg, A.V.; Tarjan, R.E. A New Approach to the Maximum-Flow Problem. J. ACM 1988, 35, 921-940. [CrossRef]
29. Cheriyan, J.; Maheshwari, S.N. Analysis of preflow push algorithms for maximum network flow. In Foundations of Software Technology and Theoretical Computer Science; Nori, K.V., Kumar, S., Eds.; Springer: Berlin/Heidelberg, Germay, 1988; pp. 30-48.
30. King, V.; Rao, S.; Tarjan, R. A Faster Deterministic Maximum Flow Algorithm. J. Algorithms 1994, 17, 447-474. [CrossRef]
31. Goldberg, A.V.; Rao, S. Beyond the Flow Decomposition Barrier. J. ACM 1998, 45, 783-797. [CrossRef]
32. Golumbic, M.C. (Ed.) Algorithmic Graph Theory and Perfect Graphs; Academic Press: Cambridge, MA, USA, 1980.
33. Degiorgi, D.G.; Simon, K. A dynamic algorithm for line graph recognition. In Graph-Theoretic Concepts in Computer Science; Nagl, M., Ed.; Springer: Berlin/Heidelberg, Germay, 1995; pp. 37-48. [CrossRef]
34. Barabási, A.L.; Albert, R. Emergence of Scaling in Random Networks. Science 1999, 286, 509-512. [CrossRef]
35. Albert, R.; Barabási, A.L. Topology of Evolving Networks: Local Events and Universality. Phys. Rev. Lett. 2000, 85, 5234-5237. [CrossRef]
36. Hagberg, A.A.; Schult, D.A.; Swart, P.J. Exploring Network Structure, Dynamics, and Function using NetworkX. In Proceedings of the 7th Python in Science Conference, Pasadena, CA, USA, 19-24 August 2008; Varoquaux, G., Vaught, T., Millman, J., Eds.; pp. 11-15.
37. Freeman, L.C. A Set of Measures of Centrality Based on Betweenness. Sociometry 1977, 40, 35-41. [CrossRef]
38. Bavelas, A. Communication Patterns in Task-Oriented Groups. J. Acoust. Soc. Am. 1950, 22, 725-730. [CrossRef]
39. Sabidussi, G. The centrality index of a graph. Psychometrika 1966, 31, 581-603. [CrossRef]
40. Erdős, P.; Rényi, A. On Random Graphs. Publ. Math. 1959, 6, 290-297.
41. Gilbert, E.N. Random Graphs. Ann. Math. Stat. 1959, 30, 1141-1144. [CrossRef]
42. Reilly, W. Highway Capacity Manual; Transport Research Board: Washington, DC, USA, 2000.
43. Ioffe, S.; Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In Proceedings of the 32nd International Conference on International Conference on Machine Learning, ICML'15, Lille, France, 7-9 July 2015; Volume 37, pp. 448-456.
44. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 26 June-1 July 2016; pp. 770-778. [CrossRef]
45. Glorot, X.; Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. J. Mach. Learn. Res. 2010, 9, 249-256.
46. Kingma, D.P.; Ba, J.L. Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015-Conference Track Proceedings, San Diego, CA, USA, 7-9 May 2015; pp. 1-15.
47. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. Software. Available online: tensorflow.org (accessed on 1 February 2022).
48. Adler, P. Fractures and Fracture Networks; Kluwer Academic: Dordrecht, The Netherlands, 1999.
49. Cammarata, G.; Fidelibus, C.; Cravero, M.; Barla, G. The Hydro-Mechanically Coupled Response of Rock Fractures. Rock Mech. Rock Eng. 2007, 40, 41-61. [CrossRef]
50. Fidelibus, C.; Cammarata, G.; Cravero, M. Hydraulic characterization of fractured rocks. In Rock Mechanics: New Research; Abbie, M., Bedford, J.S., Eds.; Nova Science Publishers Inc.: New York, NY, USA, 2009.
51. Hyman, J.D.; Hagberg, A.; Osthus, D.; Srinivasan, S.; Viswanathan, H.; Srinivasan, G. Identifying Backbones in Three-Dimensional Discrete Fracture Networks: A Bipartite Graph-Based Approach. Multiscale Model. Simul. 2018, 16, 1948-1968. [CrossRef]
52. Sanchez-Vila, X.; Guadagnini, A.; Carrera, J. Representative hydraulic conductivities in saturated grqundwater flow. Rev. Geophys. 2006, 44, 1-46. [CrossRef]
53. Svensk Kärnbränslehantering AB. Data Report for the Safety Assessment, SR-Site; Technical Report TR-10-52; SKB: Stockholm, Sweden, 2010.
54. Hyman, J.D.; Aldrich, G.; Viswanathan, H.; Makedonska, N.; Karra, S. Fracture size and transmissivity correlations: Implications for transport simulations in sparse three-dimensional discrete fracture networks following a truncated power law distribution of fracture size. Water Resour. Res. 2016, 52, 6489. [CrossRef]
55. Berrone, S.; Pieraccini, S.; Scialò, S. A PDE-constrained optimization formulation for discrete fracture network flows. SIAM J. Sci. Comput. 2013, 35, B487-B510. [CrossRef]
56. Berrone, S.; Pieraccini, S.; Scialò, S. On simulations of discrete fracture network flows with an optimization-based extended finite element method. SIAM J. Sci. Comput. 2013, 35, A908-A935. [CrossRef]
57. Berrone, S.; Pieraccini, S.; Scialò, S. An optimization approach for large scale simulations of discrete fracture network flows. J. Comput. Phys. 2014, 256, 838-853. [CrossRef]
58. Srinivasan, S.; Karra, S.; Hyman, J.; Viswanathan, H.; Srinivasan, G. Model reduction for fractured porous media: A machine learning approach for identifying main flow pathways. Comput. Geosci. 2019, 23, 617-629. [CrossRef]
59. Blizard, W.D. Multiset theory. Notre Dame J. Form. Log. 1989, 30, 36-66. [CrossRef]
60. Hein, J.L. Discrete Mathematics; Jones \& Bartlett Publishers: Burlington, MA, USA, 2003.
