

Graph Neural Networks for Graph Drawing

Matteo Tiezzi¹, Gabriele Ciravegna², and Marco Gori², *Fellow, IEEE*

Abstract—Graph drawing techniques have been developed in the last few years with the purpose of producing esthetically pleasing node-link layouts. Recently, the employment of differentiable loss functions has paved the road to the massive usage of gradient descent and related optimization algorithms. In this article, we propose a novel framework for the development of *Graph Neural Drawers* (GNDs), machines that rely on neural computation for constructing efficient and complex maps. GND is Graph Neural Networks (GNNs) whose learning process can be driven by any provided loss function, such as the ones commonly employed in Graph Drawing. Moreover, we prove that this mechanism can be guided by loss functions computed by means of feedforward neural networks, on the basis of supervision hints that express beauty properties, like the minimization of crossing edges. In this context, we show that GNNs can nicely be enriched by positional features to deal also with unlabeled vertices. We provide a proof-of-concept by constructing a loss function for the edge crossing and provide quantitative and qualitative comparisons among different GNN models working under the proposed framework.

Index Terms—Graph Drawing, Graph Neural Drawers (GNDs), Graph Neural Networks (GNNs), graph representation learning.

I. INTRODUCTION

VISUALIZING complex relations and interaction patterns among entities is a crucial task, given the increasing interest in structured data representations [1]. The Graph Drawing [2] literature aims at developing algorithmic techniques to construct drawings of graphs—i.e., mathematical structures capable of efficiently representing the aforementioned relational concepts with nodes and edges connecting them—for example via the node-link paradigm [3]–[5]. The readability of graph layouts can be evaluated following some esthetic criteria, such as the number of crossing edges, minimum crossing angles, community preservation, edge length variance, and so on [6]. The final goal is to find suitable coordinates for the node positions, and this often requires explicitly expressing and combining these criteria through complicated mathematical formulations [7]. Moreover, effective approaches, such as energy-based models [8], [9] or

spring-embedders [10], [11], require hands-on expertise and trial and error processes to achieve certain desired visual properties. Additionally, such methods define loss or energy functions that must be optimized for each new graph to be drawn, often requiring adapting algorithm-specific parameters. Lately, two interesting directions have emerged in the Graph Drawing community. The former leverages the power of gradient descent to explore the manifold given by predefined loss functions or combinations of them. Stochastic gradient descent (SGD) can be used to move subsamples of vertices couples in the direction of the gradient of spring-embedder losses [12] substituting complicated techniques, such as Majorization [13]. This approach has been extended to arbitrary optimization goals, or combinations of them, which can be optimized via gradient descent if the corresponding criterion can be expressed via smooth functions [6]. The latter novel direction consists of the exploitation of deep learning models. Indeed, the flexibility of neural networks and their approximation capability can come in handy also when dealing with the Graph Drawing scenario. For instance, neural networks are capable to learn the layout characteristics from plots produced by other Graph Drawing techniques [14], [15], as well as the underlying distribution of data [16]. Very recently, the node positions produced by Graph Drawing frameworks [14] have been used as an input to Graph Neural Networks (GNNs) [17], [18] to produce a pleasing layout that minimizes combinations of esthetic losses [19].

We propose a framework, *Graph Neural Drawers* (GNDs), which embraces both the aforementioned directions. We borrow the representational capability and computational efficiency of neural networks to prove that: 1) differentiable loss functions guiding the common Graph Drawing pipeline can be provided directly by a neural network, a Neural Aesthete, even when the required esthetic criteria cannot be directly optimized. In particular, we propose a proof-of-concept where we focus on the criteria of edge crossing, proving that a neural network can learn to identify if two arcs are crossing or not and provide a differentiable loss function toward nonintersection. Otherwise, in fact, this simple esthetic criterion cannot be achieved through direct optimization, because it is nondifferentiable. Instead, the Neural Aesthete provides a useful and flexible gradient direction that can be exploited by (stochastic) gradient descent methods. 2) Moreover, we prove that GNNs, even in the nonattributed graph scenario if enriched with appropriate node positional features, can be used to process the topology of the input graph with the purpose of mapping the obtained node representation in a 2-D layout. We compare various commonly used GNN models [20]–[22], proving that the proposed framework is flexible enough to give these models the ability to learn a wide variety of solutions. In particular, GND

Manuscript received 20 September 2021; revised 30 April 2022; accepted 12 June 2022. This work was supported in part by the EU Horizon 2020 Project AI4Media, under Contract 951911 (<https://ai4media.eu/>). (Corresponding author: Matteo Tiezzi.)

Matteo Tiezzi is with the Department of Information Engineering and Mathematics, University of Siena, 53100 Siena, Italy (e-mail: mtiezzi@diism.unisi.it).

Gabriele Ciravegna is with the Models and Algorithms for Artificial Intelligence (MAASAI), Inria, I3S, CNRS, Université Côte d’Azur, 06000 Nice, France.

Marco Gori is with the Department of Information Engineering and Mathematics, University of Siena, 53100 Siena, Italy, and also with the MAASAI, Inria, I3S, CNRS, Université Côte d’Azur, 06000 Nice, France.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNNLS.2022.3184967>.

Digital Object Identifier 10.1109/TNNLS.2022.3184967

is capable to draw graphs: 1) from supervised coordinates, i.e., emulating Graph Drawing packages; 2) minimizing common esthetic loss functions and, additionally; 3) by descending toward the gradient direction provided by the Neural Aesthete.

This article is organized as follows. Section II introduces some basics on the Graph Drawing scenario as well as references on gradient descent approaches. Section III introduces the Neural Aesthete and provides a proof-of-concept on the edge crossing task. Section IV describes the requirements of using GNNs to draw graphs in the nonattributed scenario, as well as the problem definition and the experimental evaluation. Finally, conclusions are drawn in Section V.

II. RELATED WORK

There exists a large variety of methods in the literature to improve graph readability. A straightforward approach, which has been proven to be effective in improving the human understanding of the graph topology, consists in minimizing the number of crossing edges [23]. However, the computational complexity of the problem is NP-hard, and several authors proposed complex solutions and algorithms to address this problem [24]. Shabbeer *et al.* [25] employ an expectation–maximization algorithm based on the decision boundary surface built by a support vector machine. The underlying idea is that two edges do not cross if there is a line separating the node coordinates. Further esthetic metrics have been explored, such as the minimization of node occlusions [26], neighborhood preservation, the maximization of crossing edges angular width [27], and many more [6], [28]. Given the Graph Drawing categorization depicted in surveys [28] (i.e., force-directed, dimension reduction, and multi-level techniques), interesting, and esthetically pleasing layouts are produced by methods regarding a graph as a physical system, with forces acting on nodes with attraction and repulsion dynamics up to a stable equilibrium state [9]. Force-directed techniques inspired many subsequent works, from spring-embedders [29] to energy-based approaches [8]. The main idea is to obtain the final layout of the graph minimizing the *stress* function [see (1)]. The forces characterizing this formulation can be thought of as springs connecting pairs of nodes. This very popular formulation, exploited for graph layout in the seminal work by Kamada and Kawai [9], was optimized with the localized 2-D Newton–Raphson method. Further studies employed various complicated optimization techniques, such as the *stress majorization* approach which produces graph layout through an iterative resolution of simpler functions, as proposed by Gasner *et al.* [30]. In this particular context, some recent contributions highlighted the advantages of using gradient-based methods to solve Graph Drawing tasks. The SGD method was successfully applied to efficiently minimize the Stress function in Zheng *et al.* [12], displacing pairs of vertices following the direction of the gradient, computed in closed form. A recent framework, (GD)², leverages gradient descent to optimize several readability criteria at once [6], as long as the criterion can be expressed by smooth functions. Indeed, thanks to the powerful autodifferentiation tools available in modern machine learning frameworks [31], several criteria, such as ideal edge lengths, stress majorization, node occlusion, angular resolution, and many others, can be easily

optimized. We build our first contribution upon these ideas, proving that neural networks can be used to learn decomposed single criteria (i.e., edge crossing) approximating smooth functions, with the purpose of providing a useful descent direction to optimize the graph layout.

Deep Learning has been successfully applied to data belonging to the non-Euclidean domain, e.g., graphs, in particular, thanks to GNNs [18], [32]. The seminal work by Scarselli *et al.* [17] proposes a model based on an information diffusion process involving the whole graph, fostered by the iterative application of an *aggregation function* among neighboring nodes up to an equilibrium point. The simplification of this computationally expensive mechanism was the goal of several works which leverage alternative recurrent neural models [33] or constrained fixed-point formulations. This problem was solved via reinforcement learning algorithms as done in stochastic steady-state embedding (SSE) [34] or cast to the Lagrangian framework and optimized by a gradient descent-ascent approach, like in Lagrangian-Propagation GNNs (LP-GNNs) [35], even with the advantage of multiple layers of feature extraction [36]. The iterative nature of the aforementioned models inspired their classification under the umbrella of RecGNNs in recent surveys [18], [37].

In addition to RecGNNs, several other flavors of GNN models have been proposed, such as the ConvGNNs [38] or Attentional GNNs [21], [39], [40]. All such models fit into the powerful concept of message exchange, the foundation on which is built the very general framework of message passing neural networks (MPNNs) [41], [42].

Recent works analyze the expressive capabilities of GNNs and their aggregation functions, following the seminal work on graph isomorphism by Xu *et al.* [22]. The model proposed by the authors, Graph Isomorphism Network (GIN), leverages an injective aggregation function with the same representational power of the Weisfeiler–Leman (WL) test [43]. Subsequent works (sometimes denoted with the term WL-GNNs) try to capture higher order graph properties [44]–[47]. Bearing in mind that we deal with the nonattributed graph scenario, i.e., graphs lacking node features, we point out the importance of the nodal feature choice. Several recent works investigated this problem [48]–[50]. We borrow the highly expressive Laplacian eigenvector-based positional features described by Dwivedi *et al.* [51].

There have been some early attempts in applying deep learning models and GNNs to the Graph Drawing scenario. Wang *et al.* [15] proposed a graph-based LSTM model able to learn and generalize the coordinates patterns produced by other Graph Drawing techniques. However, this approach is limited by the fact that the model drawing ability is highly dependent on the training data, such that processing different graph classes or layout styles requires recollecting and retraining procedures. We prove that our approach is more general, given that we are able to learn both drawing styles from Graph Drawing techniques and to draw by minimizing esthetic losses. Another very recent work, DeepGD [19], consists of a message-passing GNN, which process starting positions produced by Graph Drawing frameworks [14], to construct pleasing layouts that minimize combinations of esthetic losses (stress loss combined with others). Both DeepDraw and

DeepGD share the common need of transforming the graph topology into a more complicated one: DeepDrawing [15] introduces skip connections (*fake edges*) among nodes in order to process the graph via a bidirectional-LSTM; DeepGD converts the input graph to a complete one, so that each node couples is directly connected, and requires to explicitly provide the shortest path between each node couple as an edge feature. The introduction of additional edges into the learning problem increases the computational complexity of the problem, hindering the model’s ability to scale to bigger graphs. More precisely, in the DeepGD framework, the computational complexity grows quadratically in the number of nodes $O(N^2)$. Conversely, we show that the GNN is capable of producing esthetically pleasing layouts without inserting additional edges, by simply leveraging powerful positional-structural features. Additionally, we introduce a novel neural-based mechanism, the *Neural Aesthete*, capable to express differentiable esthetic losses delivering flexible gradient direction also for nondifferentiable goals. We show that this mechanism can be exploited by gradient-descent-based Graph Drawing techniques and by the proposed GND framework. Finally, GNN-based encoder–decoder architectures can learn a generative model of the underlying distribution of data from a collection of graph layout examples [16].

III. LOSS FUNCTIONS AND NEURAL ESTHETES

A. Graph Drawing Algorithms

Graph drawing algorithms typically optimize functions that somehow express a sort of beauty index, leveraging information visualization techniques, graph theory concepts, topology, and geometry to derive a graphical visualization of the graph at hand in a bidimensional or tridimensional space [2], [52], [53]. Amongst others, typical beauty indexes are those of measuring the degree of edge crossings [23], the measurements to avoid small angles between adjacent or crossing edges and measurements to express a degree of uniform allocation of the vertexes [6], [28]. All these requirements inherently assume that the Graph Drawing only consists of the allocation of the vertexes in the layout space, since the adjacent matrix of the graph can drive the drawing of the arcs as segments. However, we can also choose to link pairs of vertexes through a spline by involving some associated variables in the optimization process [54].

Without loss of generality, in this work, we restrict our objective to the vertex coordinates optimization, but the basic ideas can be extended also to the case of appropriate arc drawing.

As usual, we denote a graph by $G = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, \dots, v_N\}$ is a finite set of N nodes and $E \subseteq V \times V$ collects the arcs connecting them. The neighborhood of node v_i is denoted by \mathcal{N}_i . We denote the coordinates of each vertex with $p_i : \mathcal{V} \mapsto \mathbb{R}^2$, for a node i mapped to a bidimensional space. We denote with $P \in \mathbb{R}^{N \times 2}$ the matrix of the node coordinates.

One of the techniques that empirically proved to be very effective for an esthetically pleasing node coordinates selection

is the *stress* function [9]

$$\text{STRESS}(P) = \sum_{i < j} w_{ij} (\|p_i - p_j\| - d_{ij})^2 \quad (1)$$

where p_i, p_j are the coordinates of vertexes i and j , respectively, d_{ij} is the graph theoretic distance (or shortest path) between node i and j , and w_{ij} is a weighting factor leveraged to balance the influence of certain pairs given their theoretical distance. Usually, it is defined as $w_{ij} = d_{ij}^{-\alpha}$ with $\alpha \in [0, 1, 2]$. The optimization of this function is generally carried out leveraging complicated resolution methods (i.e., 2-D Newton–Raphson method, Stress Majorization, and so on) that hinder its efficiency.

Recently, gradient descent methods were employed to produce graph layouts [12] by minimizing the stress function, and noticeably, Ahmed *et al.* [6] proposed a similar approach employing autodifferentiation tools. The advantage of this solution is that as long as esthetic criteria are characterized by smooth differentiable functions, it is possible to undergo an iterative optimization process¹ following, at each variable update step, the gradient of the criteria.

Clearly, the definition of esthetic criteria as smooth functions could be hard to express. For instance, while we can easily count the number of arc intersections, devising a smooth function that may drive a continuous optimization of this problem is not trivial [6], [25]. Indeed, finding the intersection of two lines,² is as simple as solving the following equation system:

$$\begin{cases} a_1x + b_1y + c_1 = 0 \\ a_2x + b_2y + c_2 = 0. \end{cases} \quad (2)$$

By employing the classic Cramer’s rule we can see that there is an intersection only in the case of a nonnegative determinant of the coefficient matrix A , $\text{Det}(A) = a_1b_2 - a_2b_1 \neq 0$. Clearly, the previous formula cannot be employed as a loss function in an optimization problem since it does not provide gradients. To tackle this issue and provide a scoring function optimizable via gradient descent, we propose the *Neural Aesthete*.

B. Neural Esthete

A major contribution of this article is that of introducing the notion of *Neural Aesthete*, which is in fact a neural network that learns beauty from examples with the perspective of generalizing to unseen data. The obtained modeled function that is expressed by the Neural Aesthete is smooth and differentiable by definition and offers a fundamental heuristic for Graph Drawing. As a proof-of-concept, we focus on edge crossing. In this case, we define the Neural Aesthete as a machine, which processes two arcs as inputs and returns the information on whether or not they intersect each other. Each arc is identified by the coordinates of the corresponding pair of vertexes, $e_u = (p_i, p_j)$ for $e_u \in \mathcal{E}$. Hence, the Neural Aesthete $v(\cdot, \cdot, \cdot) : \mathcal{E}^2 \times \mathbb{R}^m \rightarrow \mathbb{R}$ operates on the concatenation of two arcs,

¹The variables of the optimization process are the node coordinates P .

²Obviously, in the case of arcs, one should also check whether the intersection point, if exists, lies on one of the two segments.

e_u and e_v and returns

$$y_{e_u, e_v} = v(\theta, e_u, e_v) \quad (3)$$

where $\theta \in \mathbb{R}^m$ is the vector which represents the weights of the neural network. The Neural Aesthete is learned by optimizing a cross-entropy loss function $L(y_{e_u, e_v}, \hat{y}_{e_u, e_v})$ over the arcs $(e_u, e_v) \in \mathcal{E}$, which is defined as

$$L(y_{e_u, e_v}, \hat{y}_{e_u, e_v}) = -(\hat{y}_{e_u, e_v} \cdot \log(y_{e_u, e_v}) + (1 - \hat{y}_{e_u, e_v}) \cdot \log(1 - y_{e_u, e_v})) \quad (4)$$

where \hat{y}_{e_u, e_v} is the target

$$\hat{y}_{e_u, e_v} = \begin{cases} 0, & \text{if } (e_u, e_v) \text{ do not intersect} \\ 1, & \text{otherwise} \end{cases} \quad (5)$$

and the intersection of e_u and e_v is automatically computed, e.g., by solving (2).

Notice that the learning process from a finite set of supervised examples yields weights that allow us to estimate the probability of the intersection of any two arcs. Basically, the learned output of the neural network can be regarded as a degree of intersection between any arc couple. Once learned, this characteristic of the Neural Aesthete comes in handy for the computation of the gradient of a loss function for Graph Drawing. In general, we want to move the extreme nodes defining the two arcs toward the direction of *nonintersection*.

Hence, for the Graph Drawing task, the Neural Aesthete is able to process an unseen edge couple (e_u, e_v) randomly picked from the edge list \mathcal{E} , and to predict their degree of intersection y_{e_u, e_v} . We define the loss function $L(\cdot, \cdot)$ on this edge pair as the cross-entropy with respect to the target *no-intersection*, $\hat{y}_{e_u/e_v} = 0$

$$H_{u,v} = L(y_{e_u, e_v}, \hat{y}_{e_u/e_v}) = -\log(1 - y_{e_u, e_v}). \quad (6)$$

This smooth and differential loss function fosters the utilization of gradient descent methods to optimize the problem variables, i.e., the arc node coordinates (e_u, e_v) .

This same procedure can be replicated to all the graph edges

$$H(P) = \sum_{(e_u, e_v) \in \mathcal{E}} L(y_{e_u, e_v}, \hat{y}_{e_u/e_v}). \quad (7)$$

Overall, a possible Graph Drawing scheme is the one which returns

$$P^* = \arg \min_P H(P). \quad (8)$$

This can be carried out by classic optimization methods. For instance, a viable solution is by gradient descent as follows:

$$P \leftarrow P - \eta \nabla_P H(P)$$

where η specifies the learning rate.

It is worth mentioning that overall, this approach leverages the computational efficiency and parallelization capabilities of neural networks. Hence, the prediction of the edge-crossing degree can be carried out for many edge couples in parallel.

Moreover, this same approach can be conveniently combined with other esthetic criteria, for instance, coming from




Fig. 1. NEURAL ESTHETE FOR EDGE CROSSING. Left-to right, graph layouts with starting random node coordinates (START), optimized by minimizing stress function with gradient descent (STRESS), optimized by gradient descent applied on the Neural Aesthete for edge-crossing loss (NA-CROSS), optimized by alternating stress loss and Neural Aesthete loss in subsequent iterations (COMBINED). We report the graph layouts generated in three random sparse graphs, one for each row.

other Neural Aesthetes or from classical loss function (e.g., stress). For example, we could consider

$$E = H(P) + \lambda_A A(\cdot) + \lambda_B B(\cdot) \quad (9)$$

where $A(\cdot)$ and $B(\cdot)$ denotes other esthetic criteria characterized by smooth differentiable functions.

C. Example: Neural Esthete on Small-Sized Random Graphs

We provide a qualitative proof-of-concept example for the aforementioned Neural Aesthete for edge crossing in Fig. 1.

We built an artificial dataset composed of 100K entries to train the Neural Aesthete. Each entry of the dataset is formed by an input-target couple (x, \hat{y}) . The input pattern x corresponds to the Neural Aesthete arcs input positions³ as defined in Section III, whose node coordinates are randomly picked inside the interval $[0, 1]$. The corresponding target \hat{y} is defined as in (5).

We balanced the dataset composition in order to have a comparable number of samples between the two classes (cross/no-cross). We trained a Neural Aesthete implemented as a multilayer perceptron (MLP) with two hidden layers of 100 nodes each and ReLu activation functions, minimizing the cross-entropy loss function with respect to the targets and leveraging the Adam optimizer [55]. We tested the generalization capabilities of the learned model on a test dataset composed of 50K entries, achieving a test accuracy of 97%.⁴ Hence, the learned model constitutes the Neural Aesthete for the task of edge crossing. Given an unseen input composed

³Which are defined as $x := [e_u, e_v] = [p_i, p_j, p_h, p_k]$, given two arcs $e_u = (p_i, p_j)$ and $e_v = (p_h, p_k)$.

⁴For comparison, a decision tree model, trained on the same dataset, only reaches a test accuracy of 78.7%.

of a couple of arcs, the learned model outputs a probability distribution representing a degree of intersection. Following the common pipeline of Graph Drawing methods with gradient descent, the Neural Aesthete output represents a differentiable function that provides an admissible descent direction for the problem parameters P .

To test the capability of the proposed solution, we leverage an artificial dataset of random graphs with a limited number of nodes ($N \in [20, 40]$). We generated Erdős-Rényi graphs with the method presented in [56] for efficiently creating sparse⁵ random networks, implemented in NetworkX [57]. We selected the connected component of the generated graph having the biggest size (max node number).

Fig. 1 reports a qualitative example of the proposed method in three graphs from the aforementioned dataset. To generate the graph layout, we carried out an optimization process on mini-batches of ten arc couples for an amount of 2K iterations (gradient steps). The first column depicts the starting random positions of the nodes; the second column reports the graph layout obtained with an in-house implementation of the stress function [see (1)], optimized via gradient descent as done in [6]; the third column contains the results obtained by optimizing the loss provided by the proposed Neural Aesthete for edge crossing; the fourth column reports the layouts obtained alternating the optimization of the stress function and edge crossing in subsequent update steps. It is noticeable to see how the solution provided by our approach is capable to avoid any arc intersection in these simple graphs. Moreover, the fact that the Neural Aesthete output represents a form of degree-of-intersection seems to provide a good gradient direction that easily moves the arcs into a recognizable angle pattern, even when combined with other criteria (fourth column). The proposed proof-of-concept proves that Neural Aesthetes represent a feasible, general and efficient solution for Graph Drawing. In the following, we prove that this same approach can be used to guide the training process of different kinds of deep neural models.

IV. GNNs FOR GRAPH DRAWING

The increasing adoption of GNNs in several research fields and practical tasks [58], [59] opens the road to an even wider spectrum of problems. Clearly, Graph Drawing and GNNs seem inherently linked, even if the formalization of this learning process under the GNN framework is not trivial. As pointed out in Section II, some recent works leveraged GNNs-inspired models for Graph Drawing. DeepDrawing [15] employs a graph-based LSTM model to learn node layout styles from Graph Drawing frameworks. DeepGD [19] is a concurrent work in which MPNN processes starting node positions to develop pleasing layouts that minimize combinations of esthetic losses (stress loss combined with others). Starting node positions, however, needs to be initialized by standard Graph Drawing frameworks [14]; in case a random initialization is employed, network performances deteriorate.

One of the drawbacks of both these approaches is the fact that they modify the graph topology, introducing additional connections that were not present in the original graph.

This fact entails an increased computational burden for the model. Indeed, a complete graph requires many more message exchanges than a sparse one, is the computational complexity of the GNN propagation linear in the number of edges [18]. Moreover, the complete graph processed by DeepGD is enriched with edge features being the shortest path between the nodes connected to the edge. This solution gives big advantages in tasks closely connected with stress minimization but could prevent the network from generalizing to other tasks.

We propose an approach to Graph Drawing, GNDs, that leverages the computational efficiency of GNNs and, thanks to informative nodal features (Laplacian eigenvectors, see Section IV-B), is general enough to be applied to several learning tasks.

A. Graph Neural Networks

First and foremost, let us introduce some notation. We denote with l_i all the input information (initial set of features) eventually attached to each node i in a graph G . The same holds for an arc connecting two nodes i and j , whose feature, if available, is denoted with $l_{(i,j)}$. Each node i has an associated hidden representation (or state) $x_i \in \mathbb{R}^s$, which in recent models is initialized with the initial features, $x_i = l_i$ (but it is not necessarily the case in RecGNN models [36]). Many GNN models can be efficiently described under the powerful umbrella of MPNNs [41], where the node state x_i is iteratively updated at each iteration t , through an aggregation of the exchanged information among neighboring nodes \mathcal{N}_i , undergoing a message passing process. Formally

$$x_{(i,j)}^{(t-1)} = \text{MSG}^{(t)}\left(x_i^{(t-1)}, x_j^{(t-1)}, l_{(i,j)}\right) \quad (10)$$

$$x_i^{(t)} = \text{AGG}^{(t)}\left(x_i^{(t-1)}, \sum_{j \in \mathcal{N}_i} x_{(i,j)}^{(t-1)}, l_i\right) \quad (11)$$

where $x_{(i,j)}^{(t)}$ represent explicitly the message exchanged by two nodes, computed by a learnable map $\text{MSG}^t(\cdot)$.⁶ Afterward, $\text{AGG}^t(\cdot)$ aggregates the incoming messages from the neighborhood, eventually processing also local node information, such as the node hidden state x_i and its features l_i . The messaging and aggregation functions $\text{MSG}^{(t)}(\cdot)$, $\text{AGG}^{(t)}(\cdot)$ are typically implemented via MLPs learned from data. Apart from RecGNN, other GNN models leverage a different set of learnable parameters for each iteration step. Hence, the propagation process of such models can be described as the outcome of a multilayer model, in which, for example, the node hidden representation at layer t , $x_i^{(t)}$, is provided as input to the next layer, $t + 1$. Therefore an ℓ -step message passing scheme can be seen as an ℓ -layered model.

This convenient framework is capable to describe several GNN models [18]. In this work, we focus our analysis on three commonly used GNN model from the literature (i.e., Graph Convolutional Network (GCN) [20], Graph Attention Network (GAT) [21], and GIN [22]) whose implementation is given in Table I, characterized by different kinds of aggregation

⁵The probability of edge creation has been set to $p = 0.01$.

⁶Notice that in the case in which arc features $l_{(i,j)}$ are not available they are removed from the problem formulation.

TABLE I

COMMON IMPLEMENTATIONS OF GNN AGGREGATION MECHANISMS.
SEE THE MAIN TEXT AND THE REFERENCED PAPERS FOR
FURTHER DETAILS ON THE FORMULATIONS

Method: Funct.	Implementation
GCN[20]: Mean	$\sigma\left(c_v W^{(t)} x_v^{(t-1)} + \sum_{u \in \mathcal{N}_v} c_{u,v} W^{(t)} x_u^{(t-1)}\right)$
GAT[21]: Att.	$\sigma\left(\sum_{u \in \mathcal{N}_v} \alpha_{u,v}^{(t-1)} W^{(t)} x_u^{(t-1)}\right)$
GIN[22]: Sum	$\text{MLP}^{(t)}\left((1 + \epsilon) x_v^{(t-1)} + \sum_{u \in \mathcal{N}_v} x_u^{(t-1)}\right)$

mechanisms (degree-norm, attention-based, and injective/sum, respectively). Following the notations in Table I, in GCN $c_{u,v}$ denotes a normalization constant depending on node degrees; in GAT, $\alpha_{u,v}^{(t-1)}$ is a learned attention coefficient which introduces anisotropy in the neighbor aggregation, σ denotes a nonlinearity, and W, W_0, W_1 are learnable weight matrices; and in GIN, ϵ is a learnable parameter (which is usually set to zero).

B. Problem Formulation

Through the GNDs framework, we propose to employ the representational and generalization capability of GNNs to learn to generate graph layouts. We formulate the problem as a *node-focused* regression task, in which for each vertex belonging to the input graph we want to infer its coordinates in a bidimensional plane, conditioned on the graph topology and the target layout/loss function (see Section IV-C). Furthermore, in the GND framework, we propose to employ GNNs to learn to draw by themselves following the guidelines prescribed by Neural Aesthetes (see Section IV-F). In order to be able to properly solve the Graph Drawing task via GNDs, a crucial role is played by the expressive power of the GNN model and the nodal features which are used. In fact, in line with the aforementioned regression task, each node state must be uniquely identified to be afterward mapped to a different 2-D position in the graph layout. This problem is inherently connected with recent studies on the representational capabilities of GNNs (see Section II and [61]). Standard MP-GNNs have been proved to be less powerful than the 1-WL test [44], both due to the lack of expressive power of the used aggregation mechanisms and to the existence of symmetries inside the graph. For instance, local isomorphic neighborhoods create indiscernible unfolding of the GNN computational structure. Hence, the GNN embeds isomorphic nodes to the same point in the high-dimensional space of the states, hindering the Graph Drawing task. Some approaches address this problem by proposing novel and more powerful architectures (WL-GNNs) that, however, tend to penalize the computational efficiency of the GNNs [44]. Moreover, given the fact that we focus on the task of drawing nonattributed graphs, it is even more important to enrich the nodes with powerful features able to identify both the position of nodes inside the graph (often referred to as positional encodings (PEs) [51]) and able to describe the neighboring structure.

Recently, it has been shown that the usage of random nodal features theoretically strengthens the representational capability of GNNs [62], [63]. Indeed, setting random initial node

embedding (i.e., different random values when processing the same input graph) enables GNNs to better distinguish local substructures, learn distributed randomized algorithms, and solve matching problems with nearly optimal approximation ratios. Formally, the node features can be considered as random variables sampled from a probability distribution μ with support $D \subseteq \mathbb{R}^s$

$$l_i \sim \mu \quad \forall i \in \mathcal{V} \quad (12)$$

where μ can be instantiated as the uniform distribution. The main intuition is that the underlying message passing process combines such high-dimensional and discriminative nodal features, fostering the detection of fixed substructures inside the graph [62]. These approaches, which hereinafter we refer to as rGNNs, proved that classification tasks can be tackled in a novel way, with a paradigm shift from the importance of task-relevant information (the values of the features) to the relevance of the relationship among node values. However, the peculiar regression task addressed in this work requires both positional and structural knowledge, which is essential to identify and distinguish neighboring nodes.

To address this issue, we keep standard GNN architectures and leverage positional features defined as the Laplacian eigenvectors [64] of the input graph, as introduced recently in GNNs [51]. Laplacian eigenvectors embed the graphs into the Euclidean space through a spectral technique, and are unique and distance preserving (far away nodes on the graph have large PE distance). Indeed, they can be considered hybrid positional-structural encodings, as they both define a local coordinate system and preserve the global graph structure.

Formally, they are defined via the factorization of the graph Laplacian matrix:

$$L = I - D^{-1/2} A D^{-1/2} = U^T \Lambda U \quad (13)$$

where I is the $N \times N$ identity matrix, D is the node degree matrix, A is the adjacency matrix and Λ and U correspond, respectively, to eigenvalues and eigenvectors. As proposed in [51], we use the k smallest nontrivial eigenvectors to generate a k -dimensional feature vector for each node, where k is chosen by grid search. Noticeably, given that the smallest eigenvectors provide smooth encoding coordinates of neighboring nodes, during the message exchange each node receives and implicit feedback on its own positional-structural characteristics from all the nodes with which it is communicating. This process fosters the regression task on the node coordinates, which receives useful information from their respective neighborhood. We believe that this is a crucial component of the model pipeline.

C. Experimental Setup

We test the capabilities of the proposed framework by comparing the performances of three commonly used GNN models (see Table I). In the following, we describe the learning tasks and the datasets employed for testing the different models. In Sections IV-D–IV-F, instead, we will give qualitative and quantitative evaluations for each learning problem, showing the generality of our approach.




Fig. 2. Datasets composition statistics. On the left-hand side, the histogram of the graph order (number of nodes for each graph $|\mathcal{V}|$) for both the analyzed datasets. On the right-hand side, the histogram of the graphs sizes (number of edges $|\mathcal{E}|$). The SPARSE dataset is characterized by a sparse connection pattern.

Given the fact that the outputs of GND are node coordinates, we can impose on such predictions heterogeneous loss functions that can be optimized via BackPropagation. In the proposed experiments, we test the GND performances on the loss functions defined as the following.

- 1) Distance with respect to ground-truth (GT) node coordinates belonging to certain layouts, produced by Graph Drawing packages (Section IV-D).
- 2) Esthetic loss functions (e.g., Stress) (Section IV-E).
- 3) Loss functions provided by Neural Aesthetes (Section IV-F).

We assume to work with solely the graph topology; hence, the node is not characterized by additional features.

We employed two different Graph Drawing datasets with different peculiarities. We chose to address small-size graphs (≤ 100 nodes) to assure the graph layout readability, since prior works highlighted node-link layouts are more suitable for small-size graphs [15], [65]. The former one is the ROME dataset,⁷ a Graph Drawing benchmarking dataset containing 11534 undirected graphs with heterogeneous structures and connection patterns. We preprocessed the dataset by removing three disconnected graphs⁸. Each graph contains a number of nodes between 10 and 100.

We built a second dataset, which we refer to as SPARSE, with the same technique described in Section III-C. We generated 10K Erdős–Rényi graphs following the method presented in [56] for efficient sparse random networks and implemented in NetworkX [57]. We randomly picked the probability of edge creation in the interval (0.01, 0.05) and the number of nodes from 20 to 100. To improve the sparsity and readability, we discarded all the created graphs having both more than 60 nodes and more than 120 edges. Afterward, we selected the connected component of the generated graph having the biggest size (max node number). We report in Fig. 2 a visual description of the datasets composition.

In order to carry out the training process and afterward evaluate the obtained performances, we split each of the datasets into three sets, (i.e., training, validation, and test) with a ratio of (75%, 10%, and 15%).

⁷<http://www.graphdrawing.org/data.html>

⁸Stress-based Graph Drawing techniques cannot take into account disconnected graphs. However, one can easily draw each connected component separately and then plot them side by side.

D. GNNs Learn to Draw From Ground-Truth Examples

The first experimental goal is focused on the task of learning to draw graph layouts given GT node positions produced by Graph Drawing frameworks. Among several packages, we chose NetworkX [57] for its completeness and ease of integration with other development tools. This framework provides several utilities to plot graph appearances. We choose two different classical layouts. The first is the KAMADA–KAWAI node layout [9] computed by optimizing the stress function. In a few words, this force-directed method models the layout dynamic as springs between all pairs of vertices, with an ideal length equal to their graph-theoretic distance. The latter is the SPECTRAL layout, which leverages the unnormalized Laplacian \hat{L} and its eigenvalues to build Cartesian coordinates for the nodes [66], formally

$$\hat{L} = D - A = \hat{U}^T \hat{\Lambda} \hat{U} \quad (14)$$

where $\hat{\Lambda}$ and \hat{U} correspond to the eigenvalues and eigenvectors, respectively, and using the first two nontrivial eigenvectors ($k = 2$) as the actual node coordinates.⁹ We remark that (13) and (14) produce different outputs. This layout tends to highlight clusters of nodes in the graph.¹⁰

Each training graph is enriched by PEs defined as k -dimensional Laplacian eigenvectors (see Section IV-B) and is processed by each of the tested GNN models to predict the node coordinates. Hence, we need a loss function capable to discern if the generated layout is similar to the corresponding GT. Furthermore, trained models should generalize the notion of graph layout beyond a simple one-to-one mapping. For these reasons, we leverage the Procrustes statistic [15] as a loss function since it measures the shape difference among graph layouts independently of affine transformations such as translations, rotations, and scaling. Given a graph composed of N nodes, the predicted node coordinates $P = (p_1, \dots, p_N)$ and the GT positions $\hat{P} = (\hat{p}_1, \dots, \hat{p}_N)$, the Procrustes statistic similarity is defined as the squared sum of the distances between P and \hat{P} after a series of possible affine transformations [15]. Formally

$$R^2 = 1 - \frac{(\text{Tr}(P^T \hat{P} \hat{P}^T P))^{\frac{1}{2}}}{\text{Tr}(P^T P) \text{Tr}(\hat{P}^T \hat{P})} \quad (15)$$

where $\text{Tr}(\cdot)$ denotes the trace operator and the obtained metric R^2 assumes values in the interval $[0, 1]$, the lower the better. We will use the Procrustes statistic-based similarity both as the loss function to guide the model training and to evaluate its generalization capability on the test set.

We tested the proposed framework by comparing the test performances obtained by the three different GNN models described in Table I, GCN, GAT, and GIN. All models are characterized by the ReLU nonlinearity. The GAT model is composed of four attention heads. The ϵ variable in the GIN aggregation process is set to 0, as suggested in [22]. We leverage the PyTorch implementation of the models provided by the deep graph library (DGL).¹¹

⁹The NetworkX Spectral layout adds a rescaling of the node coordinates into the range $(-1, 1)$ as a standard step.

¹⁰See the referenced papers for further details on the layout properties.

¹¹<https://www.dgl.ai/>




Fig. 3. KAMADA-KAWAI layout. Qualitative example of the predicted node coordinates for both the ROME dataset (first four column) and the SPARSE dataset (subsequent four columns). Each row depicts the Ground-Truth positions (GT), the graph layout produced by GCN, GAT, and GIN model, left-to-right. We report the predictions on three different test graphs (rows).

We searched for the best hyperparameters selecting the models with the lowest validation error obtained during training, in the following grid of values: size of node hidden states x_i in $\{10, 25, 50\}$; learning rate η in $\{10^{-4}, 10^{-3}, 10^{-2}\}$; the number of GNN layers in $\{2, 3, 5\}$; PE dimension k in $\{5, 8\}$ (20 is added to the grid in the case of the SPARSE dataset, given its greater node number lower bound); and dropout rate in $\{0.0, 0.1\}$. We considered 100 epochs of training with an early stopping strategy given by patience on the validation loss of 20 epochs. For each epoch, we sampled nonoverlapping mini-batches composed of β graphs, until all the training data were considered. We searched for the best mini-batch size β in $\{32, 64, 128\}$. We devised several competitors in order to assess the performance of the proposed approach. Given that Laplacian PEs available at the node level are powerful descriptors of the neighboring graph structure, we leverage an MLP as a baseline. This neural predictor learns a mapping to the node coordinates, solely exploiting the available local information. We compare the performances obtained by GNNs with Laplacian PEs against those achieved by the three corresponding variants of rGNNs, which we denote with rGCN, rGAT, and rGIN. For a fair comparison, we searched in the same hyperparameter space for all the baseline and competitors.

In Figs. 3 and 4, we report a qualitative evaluation obtained by the best performing models for each different GNN architecture on three randomly picked graphs from the test set of each dataset. Fig. 3 shows the aforementioned evaluation in the case of the KAMADA-KAWAI layout supervision, both for the ROME dataset [first four columns, where the first one depicts the GT] and for the SPARSE dataset. Fig. 4 shows

TABLE II
PROCRUSTES STATISTIC SIMILARITY [DEFINED IN (15)] ON THE TEST SPLIT OF THE ROME AND SPARSE DATASET. WE COMPARE THREE GND MODELS WITH TWO GRAPH LAYOUTS GENERATION, KAMADA-KAWAI, AND SPECTRAL. WE REPORT THE AVERAGE VALUES AND STANDARD DEVIATIONS OVER THREE RUNS WITH DIFFERENT WEIGHTS INITIALIZATION

	Model	ROME		SPARSE	
		KAMADA	SPECTRAL	KAMADA	SPECTRAL
COMPET.	MLP	0.291 \pm 0.000	0.144 \pm 0.000	0.282 \pm 0.001	0.131 \pm 0.000
	rGCN	0.612 \pm 0.009	0.527 \pm 0.008	0.592 \pm 0.006	0.532 \pm 0.009
	rGAT	0.475 \pm 0.008	0.437 \pm 0.008	0.465 \pm 0.015	0.425 \pm 0.017
	rGIN	0.685 \pm 0.001	0.590 \pm 0.003	0.675 \pm 0.010	0.603 \pm 0.030
GND	GCN	0.241 \pm 0.001	0.078 \pm 0.002	0.228 \pm 0.000	0.060 \pm 0.000
	GAT	0.186 \pm 0.000	0.057 \pm 0.000	0.177 \pm 0.001	0.045 \pm 0.001
	GIN	0.240 \pm 0.002	0.076 \pm 0.003	0.229 \pm 0.003	0.059 \pm 0.001

the same analysis in the case of the SPECTRAL layouts. The results show the good performances of the GND framework in generating two heterogeneous styles of graph layouts, learning from different GT node coordinates. In order to give a more comprehensive analysis, we report in Table II a quantitative comparison among the global Procrustes statistic similarity values obtained on the test set by the best models, for both datasets. We report the average score and its standard deviation over three runs with different seeds for the weights random number generator.

The strength of the Laplacian PE is validated by the decent performances yielded by the MLP baseline. Conversely, the random features characterizing the rGNNs are not sufficient to solve this node regression task. Some additional structural information is required in order to jointly represent the




Fig. 4. SPECTRAL LAYOUT. Predicted node coordinates for the ROME dataset (first four column) and the SPARSE dataset (subsequent four columns). Each row depicts the Ground-Truth positions (GT), the graph layout produced by GCN, GAT, and GIN model, left-to-right. We report the predictions on three different test graphs (rows).

node position and its surroundings. Indeed, all the models exploiting the proposed solution outperform the competitors. The improved performances with respect to MLP are due to the fact that node states receive implicit feedback on their own position during the message passing steps. The proposed GAT model with Laplacian PE achieves the best performances in all the settings. We believe that the attention mechanism plays a crucial role in the task of distinguishing the right propagation patterns, alongside the fact that the multihead attention mechanism provides a bigger number of learnable parameters with respect to the competitors.

In general, the SPECTRAL layout is easier to be learned by the models. This can be due to the fact that the Laplacian PE represents an optimal feature for this task, given the common spectral approach. Even from a qualitative perspective generated layouts are almost identical to the GT. Vice versa, the KAMADA-KAWAI layout represents a harder task to be learned from GT positions, especially in the case of the SPARSE dataset. As pointed out in [51], Laplacian-based PE have still some limitations given by natural symmetries, such as the arbitrary sign of eigenvectors, that several recent works are trying to solve [48].

E. GNNs Learn to Draw Minimizing Esthetic Loss Functions

In Section IV-D, GNDs explicitly minimize the distances with respect to certain GT node positions, hence learning to draw directly from data according to certain layouts. In this second experimental setting, instead, we want to build GNNs capable to draw at inference time respecting certain esthetic criteria which are implicitly learned during training. We defined our framework in such a way that powerful

PE features are mapped to 2-D coordinates. Given a smooth and differentiable loss function defined on such output, we can leverage the BP algorithm in order to learn to minimize heterogeneous criteria. We investigate the case in which the GNN models minimize the *Stress function* [see (1)] on the predicted node coordinates. Only during the training phase, for each graph, we compute the shortest path d_{ij} among every node couple (i, j) . At inference time, the GND framework process the graph topology (the adjacency matrix) and the node features, directly predicting the node coordinates, without the need for any further information.

We use the same experimental setup, competitors and hyperparameters selection grids of Section IV-D. However, according to a preliminary run of the models which achieved poor performances, we varied the hidden state dimension grid to $\{100, 200, 300\}$. This means that this task needs a bigger representational capability with respect to the previous one, which is coherent with the complex implicit nature of the learning problem. We set the stress normalization factor to $w_{ij} = (1/d_{ij})$ (hence, $\alpha = 1$) and compute the averaged stress function.¹² For this experiment, we use the stress value obtained on the validation split as the metric to select the best performing model. For comparison, we report the stress loss values obtained by three state-of-the-art Graph Drawing methods. Neato¹³ leverage the stress majorization [30] algorithm to effectively minimize the stress. PivotMDS [14] is a deterministic dimension reduction-based approach. Finally,

¹²We used the following *average Stress* definition to avoid potential numerical issues: $\text{STRESS}(P) = (1/D) \sum_{i < j} w_{ij} (\|p_i - p_j\| - d_{ij})^2$, where D is the number of considered node couples.

¹³Implementation available through Graphviz, <https://graphviz.com>




Fig. 5. STRESS MINIMIZATION ON ROME. Qualitative example of the graph layout produced by three GNN models on the test graphs of the ROME dataset. Each row contains one of the same three graphs depicted in the first column of Fig. 3 for comparison with the layout produced by Kamada–Kawai [9].




Fig. 6. STRESS MINIMIZATION ON SPARSE. The same setting of Fig. 5.

ForceAtlas2 [8] generates graph layouts through a force-directed method.

We report in Figs. 5 and 6 some qualitative examples of the graph layouts produced by the best selected GNN models on test samples (the same graphs selected for Fig. 3) of the two datasets, following the aforementioned setting. Noticeably, all three models succeed in producing a layout that adheres to the typical characteristics of graphs obtained via stress minimization. In particular, for reference to the drawing style, the layouts of these same graphs generated via the Kamada–Kawai algorithm (that also minimizes stress) are

TABLE III

AVERAGE STRESS LOSS VALUE OBTAINED ON THE TRAINING SET AND TEST SET BY THE BEST SELECTED MODELS, FOR EACH DATASET. WE REPORT THE MEAN AND STANDARD DEVIATION OBTAINED OVER THREE RUNS INITIALIZED WITH DIFFERENT FIXED SEEDS. WE DO NOT REPORT STANDARD DEVIATIONS FOR NEATO AND PIVOTMDS, BEING DETERMINISTIC ALGORITHMS

Model	ROME		SPARSE		
	TRAIN LOSS	TEST LOSS	TRAIN LOSS	TEST LOSS	
GD	ForceAtl.2	27.44 ± 0.01	26.82 ± 0.02	23.31 ± 0.01	22.69 ± 0.02
	Neato	4.35	4.34	5.61	5.66
	Piv.MDS	16.40	16.65	28.93	29.27
COMPET.	MLP	1.07 ± 0.01	1.06 ± 0.03	1.16 ± 0.02	1.18 ± 0.00
	rGCN	1.25 ± 0.01	1.24 ± 0.04	1.70 ± 0.02	1.72 ± 0.01
	rGAT	0.98 ± 0.02	0.97 ± 0.00	1.34 ± 0.01	1.36 ± 0.01
	rGIN	1.11 ± 0.03	1.49 ± 0.04	1.49 ± 0.07	1.89 ± 0.02
GND	GCN	0.51 ± 0.02	0.53 ± 0.03	0.56 ± 0.05	0.61 ± 0.02
	GAT	0.33 ± 0.02	0.34 ± 0.00	0.30 ± 0.02	0.33 ± 0.01
	GIN	0.49 ± 0.04	0.88 ± 0.05	0.28 ± 0.05	0.81 ± 0.01

depicted in the first and fifth columns of Fig. 3, for ROME and SPARSE datasets, respectively. Comparing the graph layout produced by the various GNN models and the aforementioned ones from Kamada–Kawai, also, in this case, it is easy to see from a qualitative analysis that the GAT model is the best performing one. The peculiar characteristics of the SPARSE dataset (sparse connection patterns, causing many symmetries and isomorphic nodes) entail a hardship in minimizing the stress loss function in some of the reported examples.

A quantitative comparison is reported in Table III, with the stress values obtained by the best models for each competitor and dataset, both at training time and test time, averaged over three runs initialized with different seeds. Once again, GAT performs the best. The metrics obtained by the GIN model highlight an overfit of the training data, given the selected grid parameters. GND models obtain better stress than all the SOTA Graph Drawing packages, with Neato being the best performing one in terms of stress minimization, as expected. A similar conclusion with respect to the previous experiment can be drawn regarding the results obtained by rGNNs and MLP. Indeed, these results show how learning to minimize stress requires both positional and structural knowledge and that the message passing process fosters the discriminative capability of the learned node states, with respect to solely exploiting local information.

Summing up, the experimental campaign showed the generalization capabilities of the proposed framework even in the task of minimizing common esthetic criteria imposed on the GNNs nodewise predictions, such as the *stress function*, on unseen graphs. The GND framework is capable to predict node positions on unseen graphs respecting typical stress minimization layouts, without providing at inference time any explicit graph-theoretic/shortest path information.

F. GNNs Learn to Draw From Neural Esthetes

In Section IV-E, we showed that GNDs are capable to learn to minimize a differentiable smooth function that implicitly guides the node coordinates positioning. In a similar way, the Neural Aesthetes presented in Section III provide a smooth

differentiable function that can be leveraged to find a good gradient descent direction for the learning parameters. In this section, we mix the two proposals in order to build a GND that learns to generate graph layouts thanks to the gradients provided by the edge-crossing Neural Aesthete, and, eventually, to optimize the combination of several esthetic losses.

At each learning epoch, GND minimizes the loss function $H(P)$ defined in (7), over the whole edge list \mathcal{E} . The loss function can be computed as follows: the GNN model processes the graph and predicts nodewise coordinates. Given such predicted node positions and the input graph adjacency matrix, the Neural Aesthete (which was trained beforehand as explained in Section III-C) processes couples of arcs and outputs their degree-of-intersection. The overall loss function can then be composed of the contribution given by each of the considered arc-couples, as in (6).

We restrict our analysis to the Rome dataset, exploiting a GAT model with two hidden layers, a hidden size of node state of 25, PE dimension $k = 10$, and learning rate $\eta = 10^{-2}$. We compare the graph layout generated by this model in three randomly picked test graphs, comparing three different loss function definitions: 1) stress loss; 2) neural-esthete edge-crossing-based loss $H(P)$; and 3) a combination of the two losses with a weighing factor $\lambda = 0.5$ acting on the Neural Aesthete loss, in particular

$$\text{LOSS}(P) = \text{STRESS}(P) + \lambda H(P). \quad (16)$$

We report in Fig. 7 some qualitative results on three test graphs (one for each row). We compare the layout obtained by optimizing the stress function (the first column, see Section IV-E), the edge-crossing Neural Aesthete (second column), and the combination of the two losses.

The styles of the generated layout are recognizable with respect to the plain optimization of the Neural Aesthete with gradient descent (see Fig. 1), meaning that the GND framework is able to fit the loss provided by the Neural Aesthete and to generalize it to unseen graphs. Noticeably, the introduction of the combined loss functions (third column in Fig. 7) helps in better differentiating the nodes in the graph with respect to the case of solely optimizing stress. The neural esthete-guided layouts (second and third column) tend to avoid edge intersections, as expected. This opens the road to further studies in this direction, leveraging the generality of the Neural Aesthetes’ approach and the representation capability of GNNs.

G. Computational Complexity

The proposed framework leverages the same computational structure of the underlying GNN model, which we can generally describe, for each parameter update, as linear with respect to the edge number $\mathcal{O}(T(|\mathcal{V}| + |\mathcal{E}|))$, where T is the number of iterations/layers, $|\mathcal{V}|$ the number of nodes and $|\mathcal{E}|$ the number of edges. Through our approach, there is not any increase in the computation related to the graph topology or the edge connection patterns. At inference time, the only additional requirement is the computation of the Laplacian PEs, requiring $\mathcal{O}(\mathcal{E}^{3/2})$, with \mathcal{E} being the number of edges, that, however, can be improved with the Nystrom method [51], [67].

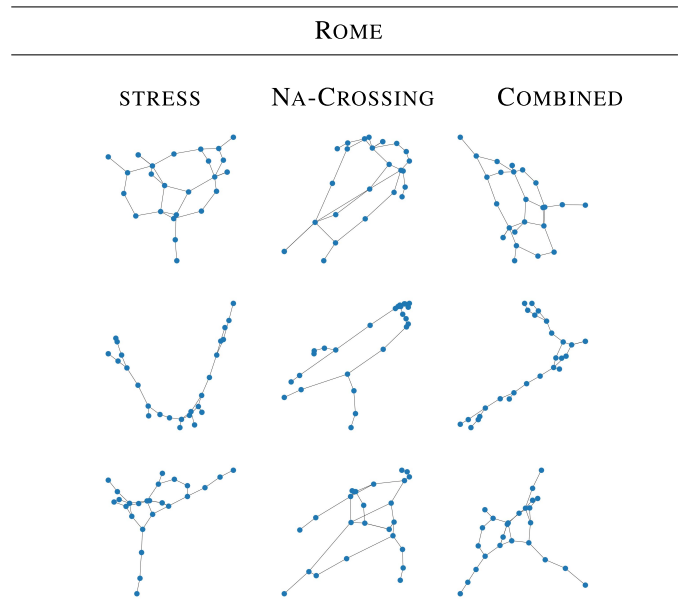


Fig. 7. LEARNING FROM THE NEURAL AESTHETE. We report the layouts obtained on three randomly picked test graphs from the Rome dataset, one for each row. Left-to-right: graph layout generated by optimizing the stress loss function, the edge-crossing Neural Aesthete-based loss (denoted with NA-crossing), the combination of the two losses with a weighing factor $\lambda = 0.5$.

H. Scaling to Bigger Graphs

Common Graph Drawing techniques based on multidimensional scaling [7] or SGD [12] require *ad hoc* iterative optimization processes for each graph to be drawn. Additionally, dealing with large-scale graphs—both in terms of the number of nodes and the number of involved edges—decreases the time efficiency of these approaches. Conversely, once a GND has been learned, the graph layout generation consists solely of the extraction of the Laplacian PE followed by a forward pass on the chosen GNN backbone. In this section, we prove the ability of GND to scale to real-world graphs, providing quantitative results in terms of computational times and qualitative analysis of the obtained graph layouts, with respect to SOTA Graph Drawing techniques. We employed the best-performing GAT model trained to minimize the stress loss on the Rome dataset (Section IV-E). We test the model inference performances on bigger scale graphs from the SuiteSparse Matrix Collection.¹⁴ We report in Fig. 8 the computational times required by the different techniques to generate graph layouts of different scale, from the `dwt_n` graph family. We analyze both the correlation on graph order (left—the varying number of nodes) and size (right—the varying number of edges). We compare the GND execution times against those of the NetworkX-GraphViz implementation of `neato` and `sfdp`, the latter being a multilevel force-directed algorithm that efficiently layouts large graphs. We also tested the Fruchterman–Reingold force-directed algorithm implemented in NetworkX (denoted with FR) and the PivotMDS implementation from the NetworKit C++ framework [68]. The tests were performed in a Linux environment equipped with

¹⁴<https://sparse.tamu.edu>

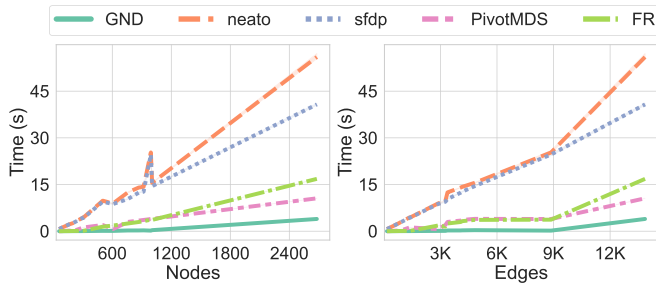


Fig. 8. Computational time comparison on dwt_n graphs. Left: correlation between the number of nodes in the graph and the layout generation timings of the analyzed Graph Drawing methods. Right: correlation between number of edges available in the graphs and the corresponding layout generation timings.




Fig. 9. Large-scale graphs from the SuiteSparse Matrix collection. Left to right: layouts produced by a GAT-based GND (trained to minimize stress on the Rome dataset), layout produced by the $sfdp$ algorithm for large-scale graphs and outcome of the PivotMDS method. We report for each row the name of the graph from the dataset collection, its order ($|\mathcal{V}|$) and size ($|\mathcal{E}|$).

an Intel Core i9-10900X CPU @ 3.70 GHz, 128 GB of RAM, and an NVIDIA GeForce RTX 3090 GPU (24 GB).

We report the average execution times over three runs (we omit the variances due to their negligible values). These results confirm the advantages of the proposed approach. While all the competitors require an expensive optimization process that increases their impact with a bigger graph scale,

the fast inference step carried on by GNDs assures small timings even with big graphs. Computing Laplacian PE is scalable and does not hinder the time efficiency of the proposed method. To assess the quality of the generated layouts, we report in Fig. 9 a comparison among the ones yielded by GND the framework, $sfdp$ and PivotMDS on several graphs from the SuiteSparse collection (we report the graph name, its order $|\mathcal{V}|$ and size $|\mathcal{E}|$). While we remark that in this experiment we exploited a GND model trained on a smaller scale dataset (i.e., Rome), the performances show a significant ability of the model to generalize the learned laws (e.g., the stress minimization in this case) to unseen graphs, even when dealing with diverging characteristics. However, we also remark that graphs having very diverse structures from the training distribution may be not correctly plotted. The causes of such performance drop are twofold. First, the intrinsic dependence of neural models on the inductive biases learned during the training process leads to an inability to generalize to unseen graph topologies. On the other hand, the limitations of Laplacian PE to discriminate certain graph symmetries or structures [51] may be further compounded with larger scale datasets, which is an active area of research [48].

V. CONCLUSION

Starting from some very interesting and promising results on the adoption of GNNs for Graph Drawing, which is mostly based on supervised learning, in this article, we proposed a general framework to emphasize the role of unsupervised learning schemes based on loss functions that enforce classic esthetic measures. When working in such a framework, referred to as *GNDs*, we open the doors toward the construction of a novel machine learning-based drawing scheme where the *Neural Aesthete* drives the learning of a GNN toward the optimization of beauty indexes. While we have adopted the *Neural Aesthetes* only from learning to minimize arc intersections, the same idea can be used for nearly any beauty index. We show that our framework is effective also for drawing unlabeled graphs. In particular, we rely on the adoption of Laplacian eigenvector-based positional features [51] for attaching information to the vertexes, which leads to very promising results.

ACKNOWLEDGMENT

The authors would like to thank Giuseppe Di Battista for the insightful discussions and for useful suggestions on the Graph Drawing literature and methods.

REFERENCES

- [1] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, “Geometric deep learning: Grids, groups, graphs, geodesics, and gauges,” 2021, *arXiv:2104.13478*.
- [2] G. D. Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*. Upper Saddle River, NJ, USA: Prentice-Hall, 1998.
- [3] B. Saket, P. Simonetto, S. Kobourov, and K. Börner, “Node, node-link, and node-link-group diagrams: An evaluation,” *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 12, pp. 2231–2240, Dec. 2014.
- [4] R. Tamassia, *Handbook of Graph Drawing and Visualization*. Boca Raton, FL, USA: CRC Press, 2013.

- [5] W. Didimo, G. Liotta, and F. Montecchiani, "A survey on graph drawing beyond planarity," *ACM Comput. Surv.*, vol. 52, no. 1, pp. 1–37, Feb. 2019.
- [6] R. Ahmed, F. De Luca, S. Devkota, S. Kobourov, and M. Li, "Graph drawing via gradient descent, (GD)²," 2020, *arXiv:2008.05584*.
- [7] M. A. Cox and T. F. Cox, "Multidimensional scaling," in *Handbook of Data Visualization*. Berlin, Germany: Springer, 2008, pp. 315–347.
- [8] M. Jacomy, T. Venturini, S. Heymann, and M. Bastian, "ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software," *PLoS ONE*, vol. 9, no. 6, Jun. 2014, Art. no. e98679.
- [9] T. Kamada and S. Kawai, "An algorithm for drawing general undirected graphs," *Inf. Process. Lett.*, vol. 31, no. 1, pp. 7–15, Apr. 1989.
- [10] A. Frick, A. Ludwig, and H. Mehldau, "A fast adaptive layout algorithm for undirected graphs (extended abstract and system demonstration)," in *Graph Drawing*. Berlin, Germany: Springer, 1994, pp. 388–403.
- [11] U. Brandes and C. Pich, "An experimental study on distance-based graph drawing," in *Graph Drawing*. Berlin, Germany: Springer, 2008, pp. 218–229.
- [12] J. X. Zheng, S. Pawar, and D. F. M. Goodman, "Graph drawing by stochastic gradient descent," *IEEE Trans. Vis. Comput. Graph.*, vol. 25, no. 9, pp. 2738–2748, Sep. 2019.
- [13] J. D. Leeuw, "Convergence of the majorization method for multidimensional scaling," *J. Classification*, vol. 5, no. 2, pp. 163–180, Sep. 1988.
- [14] U. Brandes and C. Pich, "Eigensolver methods for progressive multidimensional scaling of large data," in *Graph Drawing*. Berlin, Germany: Springer, 2006, pp. 42–53.
- [15] Y. Wang, Z. Jin, Q. Wang, W. Cui, T. Ma, and H. Qu, "DeepDrawing: A deep learning approach to graph drawing," *IEEE Trans. Vis. Comput. Graph.*, vol. 26, no. 1, pp. 676–686, Jan. 2020.
- [16] O.-H. Kwon and K.-L. Ma, "A deep generative model for graph layout," *IEEE Trans. Vis. Comput. Graph.*, vol. 26, no. 1, pp. 665–675, Jan. 2020.
- [17] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Trans. Neural Netw.*, vol. 20, no. 1, pp. 61–80, Dec. 2009.
- [18] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A comprehensive survey on graph neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 4–24, Mar. 2020.
- [19] X. Wang, K. Yen, Y. Hu, and H.-W. Shen, "DeepGD: A deep learning framework for graph drawing using GNN," *IEEE Comput. Graph. Appl.*, vol. 41, no. 5, pp. 32–44, Sep. 2021.
- [20] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," 2017, *arXiv:1609.02907*.
- [21] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *Proc. ICLR*, 2018, pp. 1–12.
- [22] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" in *Proc. ICLR*, Sep. 2019, pp. 1–17. [Online]. Available: <https://openreview.net/forum?id=ryGs6iA5Km>
- [23] H. Purchase, "Which aesthetic has the greatest effect on human understanding?" in *Graph Drawing*. Berlin, Germany: Springer, 1997, pp. 248–261.
- [24] B. M. Ábrego, S. Fernández-Merchant, and G. Salazar, "The rectilinear crossing number of K_n : Closing in (or are we?)," in *Thirty Essays on Geometric Graph Theory*. New York, NY, USA: Springer, 2013, pp. 5–18.
- [25] A. Shabbeer, C. Ozcaglar, M. Gonzalez, and K. Bennett, "Optimal embedding of heterogeneous graph data with edge crossing constraints," in *Proc. NIPS Workshop Challenges Data Vis.*, 2010, pp. 1–9.
- [26] M. Chrobak, M. T. Goodrich, and R. Tamassia, "Convex drawings of graphs in two and three dimensions (preliminary version)," in *Proc. 12th Annu. Symp. Comput. Geometry (SCG)*, 1996, pp. 319–328.
- [27] M. Kaufmann, A. M. Spallek, and J. Splett, "A heuristic approach towards drawings of graphs with high crossing resolution," in *Graph Drawing*, vol. 11282. Cham, Switzerland: Springer, 2018, p. 271.
- [28] H. Gibson, J. Faith, and P. Vickers, "A survey of two-dimensional graph layout techniques for information visualisation," *Inf. Vis.*, vol. 12, nos. 3–4, pp. 324–357, 2013.
- [29] T. M. J. Fruchterman and E. M. Reingold, "Graph drawing by force-directed placement," *Softw., Pract. Exper.*, vol. 21, no. 11, pp. 1129–1164, 1991.
- [30] E. R. Gansner, Y. Koren, and S. North, "Graph drawing by stress majorization," in *Graph Drawing*. Berlin, Germany: Springer, 2004, pp. 239–250.
- [31] A. Paszke *et al.*, "Automatic differentiation in PyTorch," in *Proc. NIPS*, 2017, pp. 1–4.
- [32] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: Going beyond Euclidean data," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 18–42, Jul. 2017.
- [33] Y. Li, D. Tarlow, M. Brockschmidt, and R. S. Zemel, "Gated graph sequence neural networks," in *Proc. ICLR*, 2016, pp. 1–20.
- [34] H. Dai, Z. Kozareva, B. Dai, A. J. Smola, and L. Song, "Learning steady-states of iterative algorithms over graphs," in *Proc. 35th Int. Conf. Mach. Learn. (ICML)*, vol. 80, Jul. 2018, pp. 1114–1122.
- [35] M. Tiezzi, G. Marra, S. Melacci, M. Maggini, and M. Gori, "A Lagrangian approach to information propagation in graph neural networks," in *Proc. ECAI*, in *Frontiers in Artificial Intelligence and Applications*, vol. 325. Amsterdam, The Netherlands: IOS Press, 2020, pp. 1539–1546.
- [36] M. Tiezzi, G. Marra, S. Melacci, and M. Maggini, "Deep constraint-based propagation in graph neural networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 727–739, Feb. 2022.
- [37] D. Bacciu, F. Errica, A. Micheli, and M. Podda, "A gentle introduction to deep learning for graphs," *Neural Netw.*, vol. 129, pp. 203–221, Sep. 2020.
- [38] T. N. Kipf, "Deep learning with graph-structured representations," Ph.D. thesis, UvA-DARE Digit. Academic Repository, Univ. Amsterdam, Amsterdam, The Netherlands. [Online]. Available: <https://hdl.handle.net/11245.1/1b63b965-24c4-4bcd-aabb-b849056fa76d>
- [39] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein, "Geometric deep learning on graphs and manifolds using mixture model CNNs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5115–5124.
- [40] X. Bresson and T. Laurent, "Residual gated graph ConvNets," 2017, *arXiv:1711.07553*.
- [41] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *Proc. ICML*, 2017, pp. 1263–1272.
- [42] J. Gilmer, S. Schoenholz, P. Riley, O. Vinyals, and D. George, "Message passing neural networks," in *Machine Learning Meets Quantum Physics*. Cham, Switzerland: Springer, 2020, pp. 199–214.
- [43] B. Weisfeiler and A. Lehman, "A reduction of a graph to a canonical form and an algebra arising during this reduction," *Nauchno-Tekhnicheskaya Informatsia*, vol. 2, no. 9, pp. 12–16, 1968.
- [44] C. Morris *et al.*, "Weisfeiler and Leman go neural: Higher-order graph neural networks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 4602–4609.
- [45] H. Maron, H. Ben-Hamu, N. Shamir, and Y. Lipman, "Invariant and equivariant graph networks," in *Proc. ICLR*, 2019, pp. 1–14. [Online]. Available: <https://openreview.net/forum?id=Syx72jC9tm>
- [46] G. Corso, L. Cavalleri, D. Beaini, P. Liò, and P. Veličković, "Principal neighbourhood aggregation for graph nets," in *Proc. NeurIPS*, vol. 33, 2020, pp. 13260–13271.
- [47] C. Bodnar *et al.*, "Weisfeiler and Lehman go cellular: CW networks," 2021, *arXiv:2106.12575*.
- [48] H. Cui, Z. Lu, P. Li, and C. Yang, "On positional and structural node features for graph neural networks on non-attributed graphs," 2021, *arXiv:2107.01495*.
- [49] B. Srinivasan and B. Ribeiro, "On the equivalence between node embeddings and structural graph representations," in *Proc. ICLR*, 2020, pp. 1–24.
- [50] R. Murphy, B. Srinivasan, V. Rao, and B. Ribeiro, "Relational pooling for graph representations," in *Proc. ICML*, 2019, pp. 4663–4673.
- [51] V. P. Dwivedi, C. K. Joshi, A. T. Luu, T. Laurent, Y. Bengio, and X. Bresson, "Benchmarking graph neural networks," 2020, *arXiv:2003.00982*.
- [52] G. D. Battista, P. Eades, R. Tamassia, and I. G. Tollis, "Algorithms for drawing graphs: An annotated bibliography," *Comput. Geometry*, vol. 4, no. 5, pp. 235–282, Oct. 1994.
- [53] V. Yoghoudjian *et al.*, "Exploring the limits of complexity: A survey of empirical studies on graph visualisation," *Vis. Informat.*, vol. 2, no. 4, pp. 264–282, Dec. 2018.
- [54] L. Vismara, G. Di Battista, A. Garg, G. Liotta, R. Tamassia, and F. Vargiu, "Experimental studies on graph drawing algorithms," *Softw., Pract. Exper.*, vol. 30, no. 11, pp. 1235–1284, 2000.
- [55] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015, pp. 1–15.
- [56] V. Batagelj and U. Brandes, "Efficient generation of large random networks," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 71, no. 3, Mar. 2005, Art. no. 036113.

- [57] A. Hagberg, P. Swart, and D. S. Chult, "Exploring network structure, dynamics, and function using NetworkX," Los Alamos Nat. Lab., Los Alamos, NM, USA, Tech. Rep. LA-UR-08-05495, 2008.
- [58] J. Jumper *et al.*, "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, no. 7873, pp. 583–589, Aug. 2021.
- [59] L. Zhao *et al.*, "T-GCN: A temporal graph convolutional network for traffic prediction," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 9, pp. 3848–3858, Sep. 2019.
- [60] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. ICLR*, 2017, pp. 1–14. [Online]. Available: <https://openreview.net/forum?id=SJU4ayYgl>
- [61] J. You, R. Ying, and J. Leskovec, "Position-aware graph neural networks," in *Proc. ICML*, 2019, pp. 1–10.
- [62] R. Sato, M. Yamada, and H. Kashima, "Random features strengthen graph neural networks," in *Proc. SIAM SDM*. Philadelphia, PA, USA: SIAM, 2021, pp. 333–341.
- [63] R. Abboud, I. I. Ceylan, M. Grohe, and T. Lukasiewicz, "The surprising power of graph neural networks with random node initialization," in *Proc. 30th Int. Joint Conf. Artif. Intell.*, Aug. 2021, pp. 1–14.
- [64] M. Belkin and P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural Comput.*, vol. 15, no. 6, pp. 1373–1396, 2003.
- [65] M. Ghoniem, J.-D. Fekete, and P. Castagliola, "A comparison of the readability of graphs using node-link and matrix-based representations," in *Proc. IEEE Symp. Inf. Visualizat.*, Oct. 2004, pp. 17–24.
- [66] B. Beckman, "Theory of spectral graph layout," Microsoft Res., Redmond, WA, USA, Tech. Rep., MSR-TR-94-04, 1994.
- [67] C. Fowlkes, S. Belongie, F. Chung, and J. Malik, "Spectral grouping using the Nystrom method," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 2, pp. 214–225, Feb. 2004.
- [68] C. L. Staudt, A. Sazonovs, and H. Meyerhenke, "NetworKit: A tool suite for large-scale complex network analysis," *Netw. Sci.*, vol. 4, no. 4, pp. 508–530, Dec. 2016.



Matteo Tiezzi received the M.S. degree (*cum laude*) in computer and automation engineering from the University of Siena, Siena, Italy, in 2017, and the Ph.D. degree (Hons.) in computer science (information engineering) from the University of Siena in 2021.

He is currently a Post-Doctoral Researcher with the Siena Artificial Intelligence Laboratory, Siena. His research interests include machine-learning applications to computer vision, foundational studies of learning algorithms, and non-Euclidean deep learning, with an emphasis on graph neural networks.

Dr. Tiezzi also serves as a Reviewer for several journals, such as IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE (TPAMI), the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING (TKDE), *Neurocomputing*, *Artificial Intelligence*, and *Knowledge-Based Systems*.



Gabriele Ciravegna received the master's degree (Hons.) in computer engineering from the Polytechnic University of Turin, Turin, Italy, in 2018, and the Ph.D. degree (Hons.) from the University of Siena, Siena, Italy, in 2022, under the supervision of Prof. Marco Gori.

He is currently a Post-Doctoral Researcher with INRIA, Université Côte d'Azur, Nice, France. His research interests include overcoming the intrinsic limits of machine learning and neural networks, above all understandability. Besides machine learning, he also likes tennis, volleyball, and playing the piano.

Dr. Ciravegna presented works in important international venues, such as Association for the Advancement of Artificial Intelligence (AAAI), International Joint Conferences on Artificial Intelligence (IJCAI), International Joint Conference on Neural Networks (IJCNN), and journals, including *Neural Networks*, *Neurocomputing*, and IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE (TPAMI). He is also a Reviewer of the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS journal and the AAAI conference.



Marco Gori (Fellow, IEEE) received the Ph.D. degree from the Università di Bologna, Bologna, Italy, in 1990.

He is working partly with the School of Computer Science, McGill University, Montreal, QC, Canada. He is currently a Professor with the Università di Siena, Siena, Italy, where he is leading the Siena Artificial Intelligence Laboratory. He has recently published the monograph *Machine Learning: A Constraint-Based Approach* (MK, 560 pp., 2018), which contains a unified view of his approach. He was one of the first people involved in the European project on Artificial Intelligence CLAIRE. His pioneering role in neural networks has been emerging especially from the recent interest in graph neural networks, which he contributed to introduce in the seminal paper "Graph Neural Networks," IEEE-TNN, 2009.

Professor Gori is a fellow of EurAI, International Association of Pattern Recognition (IAPR), and the Machine Learning Association European Laboratory for Learning and Intelligent Systems (ELLIS). He is on the Scientific Committee of Istituto di Calcolo e Reti ad Alte Prestazioni-Consiglio Nazionale delle Ricerche (ICAR-CNR) and the President of the Scientific Committee of Fondazione Bruno Kessler-Center for Information and Communication Technology (FBK-ICT). He currently holds the International 3IA Chair at the Université Côte d'Azur.