



HAL
open science

Deep Neural Network for DrawiNg Networks, (DNN) 2

Loann Giovannangeli, Frédéric Lalanne, David Auber, Romain Giot, Romain Bourqui

► **To cite this version:**

Loann Giovannangeli, Frédéric Lalanne, David Auber, Romain Giot, Romain Bourqui. Deep Neural Network for DrawiNg Networks, (DNN) 2. 12868, Springer International Publishing, pp.375-390, 2022, Lecture Notes in Computer Science, 10.1007/978-3-030-92931-2_27 . hal-03526866

HAL Id: hal-03526866

<https://hal.science/hal-03526866>

Submitted on 17 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Deep Neural Network for DrawiNg Networks, (*DNN*)²

Loann Giovannangeli¹[0000-0002-9395-6495],
Frederic Lalanne¹[0000-0001-9108-0955], David Auber¹[0000-0002-1114-8612],
Romain Giot¹[0000-0002-0638-7504], and Romain Bourqui¹[0000-0002-1847-2589]

¹ LaBRI, UMR CNRS 5800, University Bordeaux, 33405 Talence, France
{firstname}.{lastname}@u-bordeaux.fr

Abstract. By leveraging recent progress of stochastic gradient descent methods, several works have shown that graphs could be efficiently laid out through the optimization of a tailored objective function. In the meantime, Deep Learning (DL) techniques achieved great performances in many applications. We demonstrate that it is possible to use DL techniques to learn a graph-to-layout sequence of operations thanks to a graph-related objective function. In this paper, we present a novel graph drawing framework called (*DNN*)²: Deep Neural Network for DrawiNg Networks. Our method uses Graph Convolution Networks to learn a model. Learning is achieved by optimizing a graph topology related loss function that evaluates (*DNN*)² generated layouts during training. Once trained, the (*DNN*)² model is able to quickly lay any input graph out. We experiment (*DNN*)² and statistically compare it to optimization-based and regular graph layout algorithms. The results show that (*DNN*)² performs well and are encouraging as the Deep Learning approach to Graph Drawing is novel and many leads for future works are identified.

Keywords: Graph Drawing · Deep Learning · Graph Convolutions

1 Introduction

Optimization-based (OPT) and Deep Learning (DL) methods are gaining increasing interest in the information visualization field [30, 34]. From the very design of visualizations to their evaluations, such techniques have shown to perform well and present benefits over standard methods. These advances motivated the exploration of these techniques adaptation to graph drawing. Some studies [1, 17, 35] used OPT approaches to optimize an objective function for a single graph with Stochastic Gradient Descent (SGD) and obtained good results; Zheng *et al.* [35] even outperformed some state-of-the-art layout algorithms. On the other hand, if DL techniques have been applied on graph and graph drawing related problems (*e.g.*, evaluate aesthetic metrics) [9, 11, 20], to the best of our knowledge, only one study made use of this technique to *draw* graphs, *DeepDrawing* [31]. Their framework leverages DL techniques to learn a model to reproduce layouts (*i.e.*, ground truths) by optimizing a Procrustes-based cost

function that compares the produced layout to the ground truth one. A major flaw of optimizing such a cost function by opposition to a graph topology related function is that the model is trained to optimize a similarity to a ground truth graph layout (that can be suboptimal) rather than the emphasis of the topology. This paper presents Deep Neural Networks for DrawiNg Networks, $(DNN)^2$, a graph layout framework relying on *unsupervised* Deep Learning. It proposes to adapt well-proven Convolutional Neural Network architecture to graph context using Graph Convolutions [5, 16]. To the best of our knowledge, it is the first Deep Neural Network (DNN) architecture trained to lay generic graphs out by optimizing a graph-drawing related cost function. We propose an experimentation of $(DNN)^2$ where we use *ResNet* [14] architecture as a basis to optimize the Kruijer *et al.* [17] adaptation of the *Kullback-Leibler* divergence. In DL, as a model performance and its capability to generalize to unseen data are often incompatible, we also study the benefits of pre-training $(DNN)^2$. Finally, we statistically compare $(DNN)^2$ with state-of-the-art methods on aesthetic metrics and find that it competes with them. By efficiently learning a bounded sequence of operations that lays generic graphs out, $(DNN)^2$ experimentation suggests that graph drawing can be modeled as a mathematical function.

The remainder of the paper is organized as follows. Section 2 presents related works on OPT and DL methods in graphs context. Section 3 introduces $(DNN)^2$ and its key concepts while Section 4 presents the results of its experimental evaluation. Section 5 discusses the visual aspect of $(DNN)^2$ layouts and its limitations. Conclusions and leads for future works are presented in Section 6.

2 Related Works

First, we define the conventional notations used in this paper. Let $G(V, E)$ be a graph: V is its set of nodes $\{v_i\}, i \in [1, N], N = |V|$ and $E \subseteq V \times V$ its set of edges. Graphs are considered simple and connected. Let nodes positions be encoded in a vector $X \in \mathbb{R}^{N \times 2}$ where X_i is the 2D position of node v_i , $\|X_i - X_j\|$ relates to the Euclidean distance between points X_i and X_j .

Optimization-based (OPT) and Deep Learning (DL) techniques applications to Graph Drawing are gaining popularity and have been applied to a variety of graph and graph drawing related problems. For instance, Kwon *et al.* [19] used Machine Learning techniques to approximate a graph layout and its aesthetic metrics at the same time. Haleem *et al.* [11] also proposed to predict aesthetic metrics using a DL model. Several studies [10, 22, 27] used OPT to compute a feature vector embedding of a graph nodes. Kwon and Ma [20] proposed a Deep encoder-decoder to learn smooth transitions between different layouts of a graph.

Recently, OPT and DL techniques were proposed to lay graphs out and did compete with state-of-the-art layout algorithms. Kruijer *et al.* [17] proposed to optimize the Kullback-Leibler divergence by gradient descent. Kullback-Leibler divergence is a measure of dissimilarity between two probabilities distribution P and Q which was used to visualize data [13, 28] and is defined as: $D_{KL} = \sum_i P(i) \log \frac{P(i)}{Q(i)}$. The proposed optimization framework, *tsNET*, showed

to perform well, although its execution time is extremely high (*i.e.*, several seconds for graphs with $N < 100$). The authors proposed an improved variant of their method for which nodes positions are initialized with PivotMDS rather than randomly. This variant showed to be more efficient in terms of aesthetic metrics and converged faster on larger graphs. S_GD^2 [35] relies on the optimization of *stress* by stochastic gradient descent (SGD). Stress is modeled by a set of constraints between nodes that are relaxed by iteratively moving pairs of nodes. GD^2 [1] also leveraged SGD to optimize a set of aesthetic metrics whose combination can be tuned by associating a weight to each metric.

On the other hand, GraphTSNE [21] learned a shallow Neural Network made of Graph Convolutions to predict a graph layout. The key idea of their work is to train a model for each graph to draw, the train dataset being the graph nodes themselves. Even if their model cannot be described as *deep*, their work confirms that a t-SNE based loss can be optimized by Graph Convolutions networks. *DeepDrawing* [31] is the first method to train a DNN to compute graph layouts. It aims to mimic a *target* algorithm given as ground truth and can be seen as a fast approximation of its *target*. This was also studied by Espadoto *et al.* [7] and both studies raised several limitations to this approach. First, it requires to run the *target* algorithm thousands of times to generate labeled training data. Due to model convergence issue, the labeled data generation should be manually supervised and the model cannot reproduce results of a non-deterministic algorithm either. Second, as the model learns to mimic an algorithm, it cannot produce better results than its *target* baseline and it also learns its defects. Finally, as the function optimized by the model is not related only to its input data, it does not learn features from its input but rather from its combination input–target algorithm. Hence, it is unclear how well it can generalize to unseen data for which no *target* result was ever provided. As opposed to DeepDrawing, $(DNN)^2$ training is *unsupervised* (*i.e.*, no groundtruth layout is provided) and generated graph layouts are evaluated according to a graph topology related cost function based on t-SNE.

3 $(DNN)^2$ Framework Design

3.1 $(DNN)^2$ Architecture

A ResNet-like Basis. The design of $(DNN)^2$ architecture leverages Convolutional Neural Networks (CNNs) by adapting them to a graph context with Graph Convolutions [5, 16]. The architecture reproduces *ResNet* [14], a CNN designed to classify images and reaching a high accuracy on the ImageNet challenge [6]. It is composed of *residual blocks* that contain *shortcuts* which enable the model to work on several levels of abstraction. It is made of 52 spectral *Graph Convolutions* (see Section 3.1) organized in 16 residual blocks as in the ResNet architecture (see Figure 1). In addition, three node-wise fully connected layers with shared weights are added after the last convolution, the final layer being the model output. To handle graphs of varying sizes, the model inputs are fixed to an arbitrary size N_{max} and are padded with *fictive nodes* to fit this size. After each residual

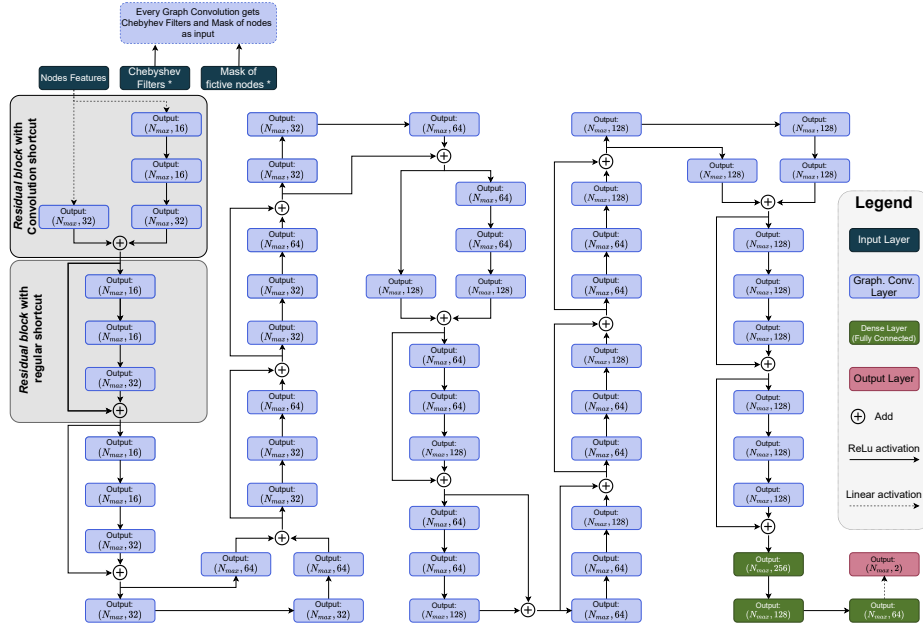


Fig. 1. $(DNN)^2$ architecture based on ResNet50 [14]. Some details have been omitted: (i) Chebyshev filters input is provided up to order 4 to all Graph Convolution layers except the 9 last layers (up to order 2); (ii) features vectors (*i.e.*, convolutions outputs) are normalized after every convolution; (iii) normalized features vectors are applied a mask encoding the *real-fictive* nodes information; and (iv) only the first two residual blocks are emphasized out of the 16 blocks.

block, its resulting features tensor is multiplied with a mask of real-fictive nodes $Mask \in \mathbb{1}^{N_{max}}$ where $Mask_i = 0$ if v_i is a fictive node, 1 otherwise. Padding the model inputs to match the expected shape could create a bias during the training: if fictive nodes are always padded at the same position in the tensors, some trainable weights will mostly see irrelevant features of *fictive nodes* and be underfitted. To avoid this bias, the padded model inputs are randomly permuted.

Spectral Graph Convolutions. Abbreviated Graph Convolutions, they were defined by Kipf and Welling [16] to operate on a graph *signal* encoded as a features vector for every node. The convolution kernel size K is defined to convolve a node with its K -hop neighborhood. The graph topology is provided through the graph spectrum (*i.e.*, eigendecomposition of the normalized Laplacian matrix) [3], approximated with Chebyshev polynomials [12]. Graph Convolutions are formally defined as a function of a signal x :

$$g_\theta \star x = U g_\theta U^T x \quad (1)$$

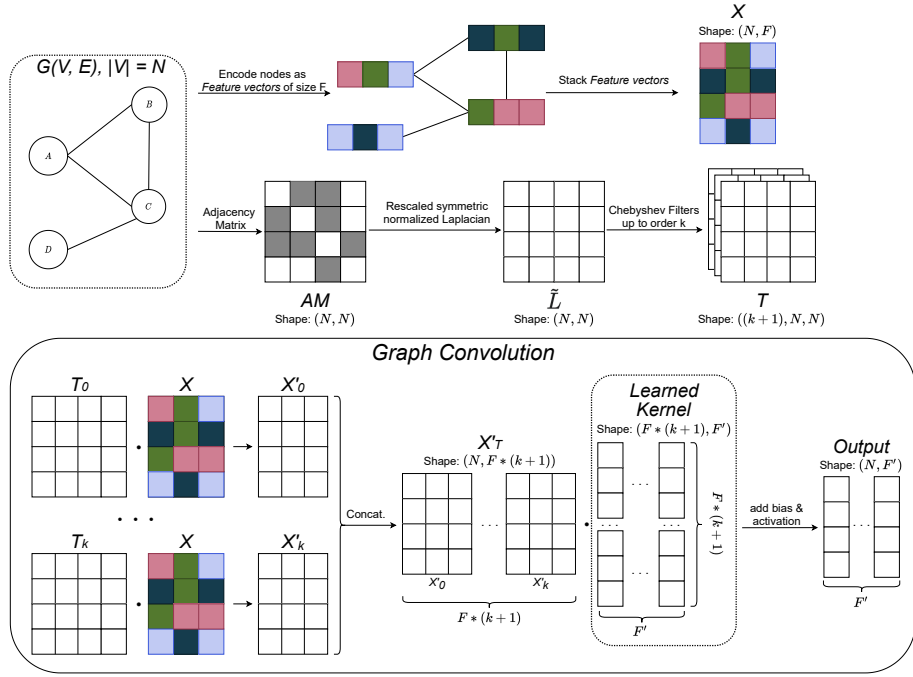


Fig. 2. Graph Convolutional layer diagram. It takes two tensors as input: (i) a feature vector to convolve (X); and (ii) a data structure that encodes the graph topology (T). X can refer to a node features vector at any step of the training.

where U is the matrix of eigenvectors of the symmetric normalized Laplacian matrix L so that $L = U\Lambda U^T$ where Λ are the eigenvalues of L . As the evaluation of Equation 1 and the eigendecomposition of L are expensive, the operation can be approximated [12] using the Chebyshev polynomials $T_k(x)$ up to order K :

$$g_{\theta'} \star x \approx \sum_{k=0}^K \theta'_k T_k(\tilde{L}) x \quad (2)$$

where \tilde{L} is the rescaled symmetric normalized Laplacian $\tilde{L} = \frac{2}{\lambda_{max}} L - I_N$, λ_{max} being the highest eigenvalue in Λ and I_N the identity matrix of size N . $\theta' \in \mathbb{R}^K$ is a vector of Chebyshev coefficients and $T_k(x)$ is the Chebyshev polynomial defined as $T_0(x) = 1$, $T_1(x) = x$ and $T_k(x) = 2xT_{k-1}(x) - T_{k-2}(x)$, $\forall k \geq 2$ and that costs $\mathcal{O}(K|E|)$ to be computed up to order K [12]. The Graph Convolution computation in this paper (illustrated in Figure 2) can be formally defined as:

$$Z = \prod_{k=0}^K T_k(\tilde{L}) \cdot X \cdot \Theta \quad (3)$$

where $X \in \mathbb{R}^{N \times F}$ is the nodes features vectors (*i.e.*, graph signal) where each node has F features, and $\Theta \in \mathbb{R}^{(F \cdot (K+1)) \times F'}$ is the learned graph convolution kernel where F' is the size of the desired output feature vector for every node. The symbol $\|$ is used as a *concatenate* operator on all the $T_k(\tilde{L}) \cdot X$ tensors.

Finally, $(DNN)^2$ is fed with three tensors: the graph signal (nodes feature vectors, defined later in Section 3.3), a mask of real-fictive nodes and the Chebyshev polynomials (also referred to as *Chebyshev filters*). Its output is set to a $N_{max} \times 2$ tensor of nodes positions in the plane. The time complexity of a forward pass in the model is $\mathcal{O}(N_{max})$ as this constant bounds the tensors size.

3.2 Loss function

Unlike *DeepDrawing* [31], $(DNN)^2$ is trained to optimize a loss function that captures the graph layout quality based on its topology. As optimizing a function for a whole dataset is fundamentally different from optimizing it for specific graphs, the loss function should have already been used with standard and OPT methods to lay graphs out so that we can compare their performances. This mainly let us with two possible functions: *stress* and *Kullback-Leibler(KL) minimization* (see Section 2). If we believe both can be optimized by $(DNN)^2$, we selected the *KL minimization* from Kruijer *et al.* [17] as it adapted better to the framework throughout experimentations. The loss is then defined as:

$$C = \lambda_{KL} C_{KL} + \frac{\lambda_c}{2N} \sum_i \|X_i\|^2 - \frac{\lambda_r}{2N^2} \sum_{i,j \in V, i \neq j} \log(\|X_i - X_j\| + \epsilon_r) \quad (4)$$

where C_{KL} is the main topology-related cost term based on the Kullback-Leibler divergence proposed by Kruijer *et al.* [17]. The second and third terms are respectively a *compression* that minimizes the scale of the drawing and a *repulsion* that counter-balances the compression. $(\lambda_{KL}, \lambda_c, \lambda_r)$ are weights used to tune the loss function during the optimization. $\epsilon_r = \frac{1}{20}$ is a regularization constant.

Kruijer *et al.* [17] defined two *stages* for their *tsNET* algorithm. In the first stage, the three λ factors are set to $(\lambda_{KL} = 1, \lambda_c = 1.2, \lambda_r = 0)$ while in the second stage, they are switched to $(1, 0.01, 0.6)$. They also proposed a variant called *tsNET** with two differences: nodes positions are initialized with PivotMDS [2] and the first stage lambda factors are $(1, 0.1, 0)$. In this paper, $(DNN)^2$ extends both *tsNET* variants and is compared to their implementation¹.

3.3 Graph signal: Initial Nodes Features

The graph signal is defined as a features vector for every node. Some methods already exist to extract a graph signal [10, 22, 27]. As standard layout algorithms achieved to lay graphs out only using their topology [8, 15, 35], we assume it can

¹ <https://github.com/HanKruijer/tsNET>, consulted on February 2021

Table 1. Random and Rome graphs datasets properties.

	Graphs distribution			Dataset size		
	V	E	Degree	Train	Validation	Test
Random Graphs	[2, 128]	[1, 6502]	[1, 118]	127 000	25 400	–
Rome Graphs	[10, 107]	[9, 158]	[1, 13]	8000	1600	1931

be sufficient to feed the model with this information encoded through Chebyshev filters. *Nodes features* are then represented by a tensor $F \in \mathbb{R}^{N \times 2}$ with nodes id to help the model differentiate them and a *random metric* to reduce overfitting.

With this *nodes features* tensor, it can be expected that adding meaningful features should help the model achieving better layouts. We experimented additional features by adding PivotMDS 2D positions such as in *tsNET** variant [17], raising its size to $F \in \mathbb{R}^{N \times 4}$. This tensor is then transformed throughout the model successive Graph Convolution and Dense layers as presented in Figure 1.

As all the nodes features are not necessarily of the same order of magnitude, they are normalized to give them the same importance.

4 Experimentation and Statistical Comparison

4.1 Datasets

Two datasets were considered for this experimentation (see Table 1), both being split for Deep Learning validation purposes (*i.e.*, *hold out* validation). In our terminology, *train* and *validation* sets are used during training to feed the model and evaluate it. *Test* set is used to benchmark models on unseen data.

Random Graphs. Used to pretrain $(DNN)^2$, its *train* set was generated to sample 1000 random graphs for each graph size between 2 and N_{max} . It is noteworthy that by generating 1000 instances of each graph size, the model will see many isomorphic graphs (mainly of small size). It means the model could overfit on small graphs, but this kind of overfitting could be beneficial for it. Since graphs can be decomposed into subgraphs of smaller size, the model capability to layout a small graph g can help it laying out a larger graph G having g as a subgraph. The *validation* set was generated with 200 instances per graph size.

Rome Graphs. Rome is a dataset of undirected graphs provided by the Graph Drawing symposium² made of 11 534 graphs, 3 of them being excluded as they are disconnected. The set was randomly split as presented in Table 1 and the layout methods of this experiment will be evaluated on the Rome *test* set.

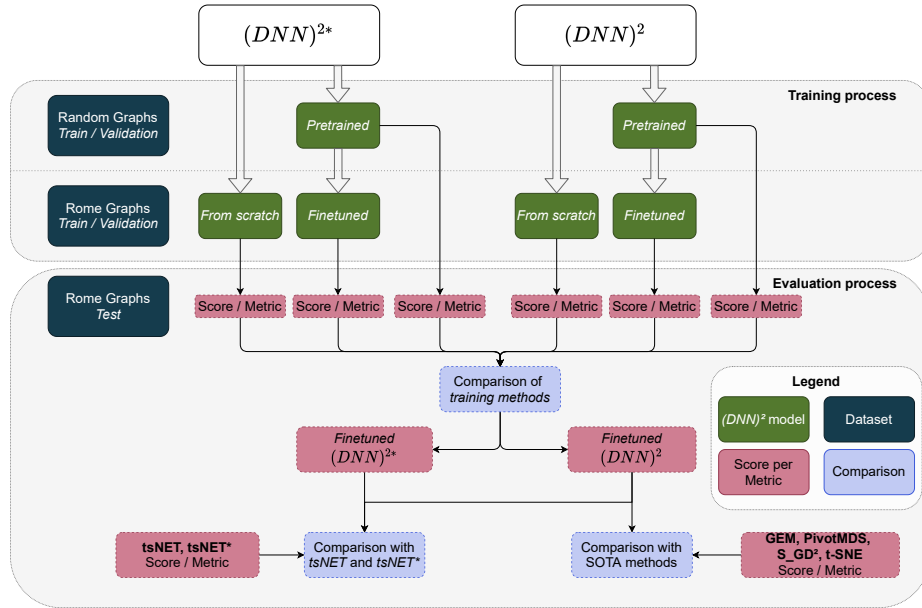


Fig. 3. $(DNN)^2$ training and evaluation pipeline. Six models are initially trained and compared. Then, the best selected models for the two loss variants are compared to $tsNET$, $tsNET^*$ and state-of-the-art layout methods.

4.2 Training

In this experiment, N_{max} was set to 128 to be slightly larger than the biggest graph in the Rome dataset. Transformed features tensors sizes were defined experimentally and are described in Figure 1. Chebyshev filters were computed up to order 4 for all the Graph Convolution layers except the nine last ones which were only provided up to order 2. Therefore, more weight is given to direct neighborhood which minimizes overdraws that are critical for the drawing quality.

Since we want to compare our DL approach to the original $tsNET$ and $tsNET^*$ algorithms, an instance of $(DNN)^2$ is trained for each of them. We refer to these two variants as $(DNN)^2$ and $(DNN)^{2*}$. Similarly to $tsNET$, the models were trained in two stages. First, to optimize the loss C (see Equation 4) with their respective $tsNET$ lambda weights (see Section 3.2). Second, the optimizer is reset and models are trained to optimize C with *second stage* lambda weights.

$(DNN)^2$ variants were trained with three methods to be evaluated on Rome graphs: (i) *pretraining* on Random graphs, (ii) *finetuning* (after the *pretraining*) on Rome graphs, (iii) training *from scratch* on Rome graphs. The goal is to verify if pretraining the model on a large set of random graphs improves its performances, and whether training on a specific dataset leads to better results than on random graphs. There are six $(DNN)^2$ instances in total (see Figure 3).

² <http://www.graphdrawing.org/data.html>, consulted on February 2021

Table 2. Quality metrics used in our benchmark and references to their definition. * represents metrics inverted to allow a *lower is better* reading for all of them.

Metric	Reference
<i>Aspect ratio</i> *	As defined in [1].
<i>Angular resolution</i> *	As defined in [1].
<i>Edge crossings number</i>	Well-known aesthetic metric [1, 24].
<i>Cluster overlap</i>	Autocorrelation metric in [32] with MCL clustering [29].
<i>Neighborhood preservation</i> *	As defined in [17].
<i>Stress</i>	Well-known aesthetic metric [1], normalized by N .

The nodes features are rescaled in $[0; 1]$ based on the *train* set. The random permutation of the model inputs (see Section 3.1) is fixed for the *test* set graphs so that every model is evaluated on the same permuted graphs.

4.3 Metrics and comparison procedure

Designing quality metrics for assessing a graph drawing quality that corroborates how well human subjects understand the drawing is a challenging question [23, 24, 26, 33]. We use a set of common metrics to assess $(DNN)^2$ efficiency and statistically compare it to state-of-the-art methods. Following the recommendations of Purchase [25], some metrics (marked with *) were inverted so that all metrics can be read as *lower is better* (see table Table 2). In addition, we measured Execution times of each algorithm in milliseconds (*ms*).

In the next, the efficiency of different graph drawing techniques are statistically compared on the presented metrics. To assess which method performs significantly better, a Kruskal-Wallis test [18] first verifies whether the differences of performances between all the compared methods on a given metric are significant or not. If so, a post-hoc Conover test [4] is applied to verify which pairs of methods are performing significantly different on that metric. For both tests, the acceptance threshold is set to $\alpha = 0.05$ and all Kruskal-Wallis tests passed.

4.4 Training Methods Evaluation

This section compares the 6 variants of $(DNN)^2$ to determine which training method is the most beneficial for the model. Execution times are not studied here and Figure 4 presents other metrics averages and standard deviations on the Rome test set for each $(DNN)^2$ instance. A red bar indicates that the corresponding model performance is significantly different to *all* others. An arc between two blue bars indicates that the difference of their performance is statistically significant.

Pretrained instances perform significantly worse than others on all metrics but aspect ratio where they lead by a fair margin. *From scratch* instances never perform the best on any metric. Overall, *finetuned* $(DNN)^2$ lead to better scores with most metrics. It could be expected as it is well known in the Image Processing community that initialize weights to *pretrained* values tends to speed up the

training process and to lead to better performances including generalization to unseen data. The idea is that it is easier for the model to learn to solve a specific task if it already knows high-level features. As *finetuned* models results are best, pretraining effectively learned the model such features that helped it to finetune.

In the next, only *finetuned* instances of $(DNN)^2$ are compared to state-of-the-art methods since they perform better on the graph drawing task.




Fig. 4. Comparison of $(DNN)^2$ instances on the Rome test set. A red bar indicates that the corresponding model performance is significantly different to *all* others. An arc between two blue bars indicates pairwise significance.

4.5 Comparison with *tsNET*

In this section, we study how the $(DNN)^2$ adaptation of *tsNET* loss performs compared to the original Optimization-based implementation of *tsNET* as both use the same cost function. The results are presented in Table 3.

$(DNN)^2$ has significantly lower scores than *tsNET* on all metrics but Stress and Execution time. Though it is significant, the difference on Edge Crossings Number is small. The Execution time difference is heavily in favor of $(DNN)^2$: 20.4ms as opposed to 6541ms for *tsNET*. The trends are about the same between $(DNN)^2*$ and *tsNET**, but their scores are closer and the differences in Angular Resolution and Edge Crossings Number are not significantly different anymore. $(DNN)^2*$ is also better on Stress and strongly better on Execution time.

Table 3. $(DNN)^2$ and $(DNN)^{2*}$ pair comparison with their respective $tsNET$ algorithm on the Rome test set. $^{\circ}$ (resp. $^+$) indicates a significant difference with $(DNN)^2$ (resp. $(DNN)^{2*}$). The best *significant* results are bold (*i.e.*, no bold value when the difference is *not* significant).

	Aspect ratio	Angular res.	Cross. number	Cluster overlap	Neighb. preserv.	Stress	Exec. time
$(DNN)^2$ $^{\circ}$	0.3 \pm 0.136	0.963 \pm 0.101	34.9 \pm 40.5	0.55 \pm 0.207	0.451 \pm 0.159	0.117 \pm 0.041	20.4 \pm 10.4
$tsNET$	0.191 \pm 0.091 $^{\circ}$	0.885 \pm 0.155 $^{\circ}$	27.7 \pm 31.8 $^{\circ}$	0.489 \pm 0.229 $^{\circ}$	0.407 \pm 0.1 $^{\circ}$	0.144 \pm 0.155 $^{\circ}$	6541 \pm 5068 $^{\circ}$
$(DNN)^{2*}$ $^+$	0.229 \pm 0.104	0.905 \pm 0.15	30.0 \pm 35.7	0.507 \pm 0.214	0.397 \pm 0.14	0.111 \pm 0.042	24.8 \pm 8.61
$tsNET^*$	0.206 \pm 0.10	0.872 \pm 0.181 $^+$	27.1 \pm 32.0	0.49 \pm 0.218 $^+$	0.386 \pm 0.115 $^+$	0.124 \pm 0.049 $^+$	5836 \pm 5933 $^+$

Table 4. $(DNN)^2$ and $(DNN)^{2*}$ pair comparisons with selected state-of-the-art algorithms. $^{\circ}$ (resp. $^+$) indicates a significant difference with $(DNN)^2$ (resp. $(DNN)^{2*}$). The best *significant* result(s) for each metric is(are) bold (*i.e.*, several bold values when the differences between the *best* algorithms is *not* significant).

	Aspect ratio	Angular res.	Cross. number	Cluster overlap	Neighb. preserv.	Stress	Exec. time
$(DNN)^2$ $^{\circ}$	0.294 \pm 0.134 $^+$	0.969 \pm 0.092 $^+$	36.3 \pm 39.9 $^+$	0.58 \pm 0.197 $^+$	0.468 \pm 0.154 $^+$	0.128 \pm 0.06 $^+$	21.0 \pm 10.3 $^+$
$(DNN)^{2*}$ $^+$	0.229 \pm 0.105 $^{\circ}$	0.917 \pm 0.138 $^{\circ}$	30.6 \pm 34.8 $^{\circ}$	0.541 \pm 0.206 $^{\circ}$	0.409 \pm 0.136 $^{\circ}$	0.115 \pm 0.046 $^{\circ}$	25.1 \pm 8.36 $^{\circ}$
t-SNE	0.276 \pm 0.158 $^{\circ+}$	0.97 \pm 0.038 $^{\circ+}$	69.1 \pm 48.9 $^{\circ+}$	0.598 \pm 0.252 $^{\circ+}$	0.584 \pm 0.097 $^{\circ+}$	0.56 \pm 0.771 $^{\circ+}$	166 \pm 71.7 $^{\circ+}$
PivotMDS	0.298 \pm 0.125 $^+$	0.978 \pm 0.088 $^{\circ+}$	38.7 \pm 43.6 $^+$	0.623 \pm 0.202 $^{\circ+}$	0.49 \pm 0.17 $^{\circ+}$	0.104 \pm 0.035 $^{\circ+}$	0.546 \pm 0.478 $^{\circ+}$
GEM	0.573 \pm 0.197 $^{\circ+}$	0.972 \pm 0.034 $^{\circ+}$	54.4 \pm 61.2 $^{\circ+}$	0.722 \pm 0.162 $^{\circ+}$	0.617 \pm 0.123 $^{\circ+}$	0.24 \pm 0.062 $^{\circ+}$	5.22 \pm 3.83 $^{\circ+}$
$S-GD^2$	0.263 \pm 0.123 $^{\circ+}$	0.812 \pm 0.208 $^{\circ+}$	32.2 \pm 36.6 $^{\circ}$	0.583 \pm 0.204 $^+$	0.439 \pm 0.181 $^{\circ+}$	0.066 \pm 0.027 $^{\circ+}$	1.13 \pm 0.91 $^{\circ+}$

It is noteworthy that $tsNET$ and $tsNET^*$ suffers from a significantly high Execution time standard deviation, meaning that the methods hardly converge on some graphs. In addition, 450 out of 1931 (*i.e.*, 23%) graphs were excluded from the test set as the $tsNET^*$ implementation would not complete on these.

We can conclude that the $(DNN)^2$ implementation that adapts $tsNET$ to a Deep Learning approach is faster but does not lead to better drawings according to most of the metrics. Although quality metrics differences are significant, they remain small and should undeniably be alleviated by future works.

$tsNET$ is designed to optimize a specific input graph at a time whereas, with the Deep Learning approach, we aim at optimizing the model and not the drawing of a single graph. If the DL training process is computationally expensive, the resulting model should be capable of computing the layout without any further need for optimization. In fact, if a DL model learns well to lay graphs out by optimizing a generic cost function, it suggests that there exists a bounded sequence of operations that efficiently projects a graph in a 2D space.

4.6 Comparison with State-of-the-art Layout Algorithms

This section studies how $(DNN)^2$ performs compared to selected layout algorithms from the literature: t-SNE [28], since we leverage the Kullback-Leibler divergence, PivotMDS [2], a deterministic Multidimensional Scaling used by $(DNN)^{2*}$ and $tsNET^*$, GEM [8], a well-established force-directed technique and $S-GD^2$ [35], a *stress* Optimization-based approach with SGD. The methods are compared on

the Rome test set and the results are reported in Table 4. $(DNN)^2$ scores are slightly different from Table 3 since all test graphs are taken into account here.

$(DNN)^{2*}$ performs better than $(DNN)^2$ as all aesthetic metrics are significantly in its favor. $(DNN)^{2*}$ is slower due to the extra processing of PivotMDS it requires. This outcome was expected in view of *tsNET* variants comparisons in [17].

$(DNN)^2$ is better than GEM on all quality metrics; and is significantly better than t-SNE on all metrics but Aspect ratio. It performs better than PivotMDS on Angular resolution, Cluster overlap and Neighborhood preservation, but is outperformed on Aspect ratio and Stress, while the difference is not significant on Edge Crossings Number. Finally, S_GD^2 performs significantly better than $(DNN)^2$ on all metrics but Cluster overlap.

Overall, $(DNN)^{2*}$ is significantly better on Aspect ratio, Cluster overlap and Neighborhood preservation than all the other considered methods. It is also the best in Edge Crossings Number with S_GD^2 . While it was observed to be better than *tsNET** on Stress, it is here outperformed by PivotMDS and S_GD^2 .

As for Execution time, we can see that both $(DNN)^2$ variants are slower than other methods except t-SNE. However, they are less sensible to graph size variations: $(DNN)^2$ variants execution time standard deviations are 33% and 47% of their average, while they range between 43% and 87% for other methods. It is important to note that a forward pass time in $(DNN)^2$ is almost constant and only takes 1.4ms (*i.e.*, 6% of its total execution time), the remaining time being used to pre-process data for the model inputs.

Although $(DNN)^2$ is not the best performing variant, its results indicate that a Deep Learning framework, without any knowledge of what is a graph layout, can learn a sequence of operations that lays graphs out. $(DNN)^{2*}$ leveraged its PivotMDS input and drawn better layouts according to the quality metrics. Its performances make it a good trade-off between *tsNET** and S_GD^2 . The latter performed surprisingly well, while GEM underperformed in this evaluation.

5 Discussion

5.1 Visual evaluation

Graph layout examples of $(DNN)^2$ are presented in Figure 5 alongside *tsNET** and S_GD^2 ones. S_GD^2 drawings being all pleasing and only a few defects away from being perfect, we can use them to see how the layouts should look like. For $(DNN)^2$, the dodecahedron and the grid graph structures can be observed but are severely distorted. It seems that *topologically equivalent* nodes (*i.e.*, nodes that can be mapped to each other by an automorphism) are grouped together. Both drawings are therefore folded, which also emphasizes their symmetry. From what we experienced, this behavior might be caused by the *compression* of the first stage training (see Section 3.2) and too similar *nodes features* and Chebyshev filters. Another explanation might be that the model had seen such small *patterns* more often during the training stage and somehow overfitted on them. On the two Rome graphs, the model has successfully laid the graphs structures out, but

	Dodecahedron	Grid	Rome graph 1 <i>id: 138</i>	Rome graph 2 <i>id: 10082</i>
$(DNN)^2$				
$(DNN)^{2*}$				
$tsNET^*$				
$S-GD^2$				

Fig. 5. Layout examples for $(DNN)^2$, $(DNN)^{2*}$, $tsNET^*$ and $S-GD^2$.

its tendency to group *topologically equivalent* nodes leads to unbalanced edge lengths, edge crossings and overplots. On the other hand, $(DNN)^{2*}$ layouts are visually more pleasing. The dodecahedron structure can clearly be identified. Despite a lack of regularity, the grid layout is also acceptable. The two Rome graph layouts demonstrate $(DNN)^{2*}$ good performances. The model was able to separate topologically equivalent nodes, though they could have been repulsed a little more. $tsNET^*$ also produces nodes overlaps where neighborhoods are similar, as it can be observed in the top right of its *Rome graph 1* layout and on the right side of *Rome graph 2*.

Overall $(DNN)^2$ and in particular $(DNN)^{2*}$ performed well even compared to OPT methods. The latter optimizing their cost function for a specific graph at a time, we could expect them to provide better results than a DL approach.

Nevertheless, $(DNN)^2$ results acts as a proof-of-concept showing that we can learn unsupervised DL models to lay graphs out. It is therefore encouraging as we believe there is still a large room for improvement.

5.2 Limitations

The main limitation of $(DNN)^2$ is the technical need to set a maximum graph size N_{max} so that the architecture tensors size is static. Setting N_{max} to an arbitrarily high number might not be an option either since it would significantly increase the data pre-processing cost which is the most expensive step in $(DNN)^2$. It is also not certain such model would learn if it is only fed with *small* graphs (*i.e.*, with $N \ll N_{max}$), as each graph convolution kernel weight would be underfitted.

Another limitation is the resources required to obtain a well-trained model. First, if the use of the trained model is straightforward, the model *training* relies on many design choices that can only be efficiently made through a trials and errors process by an informed expert. Second, the computational resources required to train the model can be prohibitive. If Deep Learning-designed computers can easily handle small to mid-scale training (*e.g.*, *finetuning*), heavier training (*e.g.*, *pretraining*) can require to generate hundreds of thousands of graphs, which required the use of a Big Data platform in this experiment.

Finally, $(DNN)^2$ has not been tested on disconnected, weighted or directed graphs. Though the handling of these graph properties is straightforward with this framework, it is not part of the scope of this study.

6 Conclusion

We introduced $(DNN)^2$, a Deep Learning based framework for graph drawing. $(DNN)^2$ proposes to adapt well-established Deep Neural Network architectures in image classification to compute the layout of an input graph by using Graph Convolutions. To the best of our knowledge, it is the first DL model trained to lay graphs out by directly optimizing a graph topology related cost function.

We provided an experimentation of the framework and compared its performances to graph drawing algorithms from the literature. The experiment showed that $(DNN)^2$ performs well compared to these algorithms despite some Deep Learning related limitations. The results highly suggest that Deep Learning is a promising approach for the future of graph drawing. It also implies that there exists a mathematical function that efficiently projects any graph structure into a drawing, and that can be learned by Deep Learning models.

Future work leads include trying out other $(DNN)^2$ implementations, meaning other Deep Neural Network architectures, loss functions (*e.g.*, *stress*) and input node features (*e.g.*, *node2vec*, *DeepWalk*). Another interesting direction is to train and evaluate $(DNN)^2$ on other graph datasets or on specific graph families. Expanding the scale of the graphs size $(DNN)^2$ can handle, or apply it to other specific graph drawing applications are also promising leads for future work.

References

1. Ahmed, R., De Luca, F., Devkota, S., Kobourov, S., Li, M.: Graph Drawing via Gradient Descent, (GD)². arXiv preprint arXiv:2008.05584 (2020)
2. Brandes, Ulrik and Pich, Christian: Eigensolver methods for progressive multidimensional scaling of large data. In: International Symposium on Graph Drawing. pp. 42–53. Springer (2006)
3. Cohen-Steiner, David and Kong, Weihao and Sohler, Christian and Valiant, Gregory: Approximating the spectrum of a graph. In: Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining. pp. 1263–1271 (2018)
4. Conover, William Jay and Iman, Ronald L: On multiple-comparisons procedures. Tech. rep., Technical report, Los Alamos Scientific Laboratory (1979)
5. Defferrard, Michaël and Bresson, Xavier and Vandergheynst, Pierre: Convolutional neural networks on graphs with fast localized spectral filtering. arXiv preprint arXiv:1606.09375 (2016)
6. Deng, Jia and Dong, Wei and Socher, Richard and Li, Li-Jia and Li, Kai and Fei-Fei, Li: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
7. Espadoto, Mateus and Hirata, Nina Sumiko Tomita and Telea, Alexandru C: Deep learning multidimensional projections. *Information Visualization* **19**(3), 247–269 (2020)
8. Frick, Arne and Ludwig, Andreas and Mehldau, Heiko: A fast adaptive layout algorithm for undirected graphs (extended abstract and system demonstration). In: International Symposium on Graph Drawing. pp. 388–403. Springer (1994)
9. Giovannangeli, Loann and Bourqui, Romain and Giot, Romain and Auber, David: Toward automatic comparison of visualization techniques: Application to graph visualization. *Visual Informatics* **4**(2), 86–98 (2020)
10. Grover, Aditya and Leskovec, Jure: node2vec: Scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 855–864 (2016)
11. Haleem, Hammad and Wang, Yong and Puri, Abishek and Wadhwa, Sahil and Qu, Huamin: Evaluating the readability of force directed graph layouts: A deep learning approach. *IEEE computer graphics and applications* **39**(4), 40–53 (2019)
12. Hammond, David K and Vandergheynst, Pierre and Gribonval, Rémi: Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis* **30**(2), 129–150 (2011)
13. Hinton, Geoffrey and Roweis, Sam T: Stochastic neighbor embedding. In: NIPS. vol. 15, pp. 833–840. Citeseer (2002)
14. K. He and X. Zhang and S. Ren and J. Sun: Deep Residual Learning for Image Recognition. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (June 2016)
15. Kamada, Tomihisa and Kawai, Satoru and others: An algorithm for drawing general undirected graphs. *Information processing letters* **31**(1), 7–15 (1989)
16. Kipf, Thomas N and Welling, Max: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
17. Krüger, Johannes F and Rauber, Paulo E and Martins, Rafael M and Kerren, Andreas and Kobourov, Stephen and Telea, Alexandru C: Graph Layouts by t-SNE. In: Computer Graphics Forum. vol. 36, pp. 283–294. Wiley Online Library (2017)

18. Kruskal, William H and Wallis, W Allen: Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* **47**(260), 583–621 (1952)
19. Kwon, Oh-Hyun and Crnovrsanin, Tarik and Ma, Kwan-Liu: What would a graph look like in this layout? a machine learning approach to large graph visualization. *IEEE transactions on visualization and computer graphics* **24**(1), 478–488 (2017)
20. Kwon, Oh-Hyun and Ma, Kwan-Liu: A deep generative model for graph layout. *IEEE Transactions on visualization and Computer Graphics* **26**(1), 665–675 (2019)
21. Leow, Yao Yang and Laurent, Thomas and Bresson, Xavier: GraphTSNE: a visualization technique for graph-structured data. *arXiv preprint arXiv:1904.06915* (2019)
22. Perozzi, Bryan and Al-Rfou, Rami and Skiena, Steven: Deepwalk: Online learning of social representations. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 701–710 (2014)
23. Purchase, Helen: Which aesthetic has the greatest effect on human understanding? In: *International Symposium on Graph Drawing*. pp. 248–261. Springer (1997)
24. Purchase, Helen C: Metrics for graph drawing aesthetics. *Journal of Visual Languages & Computing* **13**(5), 501–516 (2002)
25. Purchase, Helen C: *Experimental human-computer interaction: a practical guide with visual examples*. Cambridge University Press (2012)
26. Purchase, Helen C and Cohen, Robert F and James, Murray: Validating graph drawing aesthetics. In: *International Symposium on Graph Drawing*. pp. 435–446. Springer (1995)
27. Tang, Jian and Qu, Meng and Wang, Mingzhe and Zhang, Ming and Yan, Jun and Mei, Qiaozhu: Line: Large-scale information network embedding. In: *Proceedings of the 24th international conference on world wide web*. pp. 1067–1077 (2015)
28. Van der Maaten, Laurens and Hinton, Geoffrey: Visualizing data using t-SNE. *Journal of machine learning research* **9**(11) (2008)
29. Van Dongen, Stijn Marinus: *Graph clustering by flow simulation*. Ph.D. thesis (2000)
30. Wang, Qianwen and Chen, Zhutian and Wang, Yong and Qu, Huamin: Applying Machine Learning Advances to Data Visualization: A Survey on ML4VIS. *arXiv preprint arXiv:2012.00467* (2020)
31. Wang, Yong and Jin, Zhihua and Wang, Qianwen and Cui, Weiwei and Ma, Tengfei and Qu, Huamin: DeepDrawing: A deep learning approach to graph drawing. *IEEE Transactions on Visualization and Computer Graphics* **26**(1), 676–686 (2019)
32. Wang, Yong and Shen, Qiaomu and Archambault, Daniel and Zhou, Zhiguang and Zhu, Min and Yang, Sixiao and Qu, Huamin: Ambiguityvis: Visualization of ambiguity in graph layouts. *IEEE Transactions on Visualization and Computer Graphics* **22**(1), 359–368 (2015)
33. Ware, Colin and Purchase, Helen and Colpoys, Linda and McGill, Matthew: Cognitive measurements of graph aesthetics. *Information visualization* **1**(2), 103–110 (2002)
34. Wu, Aoyu and Wang, Yun and Shu, Xinhuan and Moritz, Dominik and Cui, Weiwei and Zhang, Haidong and Zhang, Dongmei and Qu, Huamin: Survey on Artificial Intelligence Approaches for Visualization Data. *arXiv preprint arXiv:2102.01330* (2021)
35. Zheng, Jonathan X and Pawar, Samraat and Goodman, Dan FM: Graph drawing by stochastic gradient descent. *IEEE transactions on visualization and computer graphics* **25**(9), 2738–2748 (2018)