

SONGs: Self-Organizing Neural Graphs

Łukasz Struski¹Tomasz Danel¹Marek Śmieja¹Jacek Tabor¹Bartosz Zieliński^{1,2}¹ Faculty of Mathematics and Computer Science, Jagiellonian University, Kraków, Poland² IDEAS NCBR, Warsaw, Poland

{lukasz.struski;marek.smieja;jacek.tabor;bartosz.zielinski}@uj.edu.pl

{tomasz.danel}@doctoral.uj.edu.pl

Abstract

Recent years have seen a surge in research on combining deep neural networks with other methods, including decision trees and graphs. There are at least three advantages of incorporating decision trees and graphs: they are easy to interpret since they are based on sequential decisions, they can make decisions faster, and they provide a hierarchy of classes. However, one of the well-known drawbacks of decision trees, as compared to decision graphs, is that decision trees cannot reuse the decision nodes. Nevertheless, decision graphs were not commonly used in deep learning due to the lack of efficient gradient-based training techniques. In this paper, we fill this gap and provide a general paradigm based on Markov processes, which allows for efficient training of the special type of decision graphs, which we call Self-Organizing Neural Graphs (SONG). We provide a theoretical study on SONG, complemented by experiments conducted on Letter, Connect4, MNIST, CIFAR, and TinyImageNet datasets, showing that our method performs on par or better than existing decision models.

1. Introduction

Neural networks (NNs) and decision trees (DTs) are two exceptionally powerful machine learning models with a rich and successful history in machine learning. However, they typically come with mutually exclusive benefits and limitations. NNs outperform conventional pipelines by jointly learning to represent and classify data [15]. However, they are widely opaque and suffer from a lack of transparency and explainability [30]. On the other hand, it is easy to explain predictions of DTs because they depend on a relatively short sequence of decisions [38]. However, they usually do not generalize as well as deep neural networks [8]. As a result, a strong focus is recently put on joining the positive aspects of

both models [2, 8, 13, 23, 24, 34, 36, 38]. There are methods that combine NNs and soft decision trees with partial membership in each node [8, 13, 24, 34]. Others use trees to explain NNs [5, 39] or to obtain their optimal hierarchical structure [2, 23, 36]. Finally, some models replace the final softmax layer of a neural network with a hierarchical binary decision tree [20, 21, 38].

While decision trees can increase the performance and interpretability of NNs, they usually suffer from exponential growth with depth [32], repeating nodes [8], and suboptimal structure, often selected manually before training [38]. Hence, more and more attention is put on combining NNs with decision graphs instead of trees [4, 9, 11, 22, 37]. Decision graphs have a few advantages when compared to decision trees. They have a flexible structure that allows multiple paths from the root to each leaf. As a result, nodes are reused, resulting in simpler and smaller models, which solves the replication problem [27]. Moreover, decision graphs require substantially less memory while considerably improving generalization [33]. Nevertheless, decision graphs are not commonly used in deep learning due to a lack of efficient gradient-based training techniques.

In this paper, we introduce Self-Organizing Neural Graphs (SONGs), a special type of decision graphs that generalize methods like Soft Decision Tree (SDT) [8] and Neural-Backed Decision Trees (NBDT) [38], and as a differentiable solution are applicable to any deep learning pipeline. Moreover, in contrast to the fixed structure of the existing methods [8, 38], SONGs can strengthen or weaken an edge between any pair of nodes during training to optimize their structure. We illustrate this process in Figure 1. In the beginning, the edges have random weights. However, in successive steps of training, the structure is corrected with backpropagation, and it gets sparse and converges to the binary directed acyclic graphs [28].

Our contributions can be summarized as follows:

- We introduce Self-Organizing Neural Graphs (SONGs),




Figure 1: Training stages of SONG that uses gradient descent to modify the graph structure and transition probabilities. Based on an input x , the backbone neural network (NN) extracts a vector representation, which is passed to SONG to obtain a prediction for each class (y_1 and y_2). At the beginning of training, a graph has root r , nodes v_1 and v_2 , leaves l_1 and l_2 , and randomly initialized edges (a). In the successive training iterations, the entropy of edge weights grows (b), finally resulting in a sparse binary graph, with two strong edges outgoing from each node (c). Notice that SONG contains two alternative sets of edges between the nodes (dashed blue arrows and solid red arrows, respectively) that are combined based on the input (see Figure 2 for details).

a new paradigm of end-to-end training based on Markov processes that simultaneously learn the optimal graph structure and transition probabilities.

- Our model is fully differentiable and thus suitable for combined training with other deep learning models.
- We prove empirically and theoretically that SONGs during training converge to sparse binary acyclic graphs.
- Our method performs on par or outperforms decision trees trained in a similar setup and does not require the graph/tree structure to be predefined before training.

2. Related works

2.1. Decision trees

Numerous Decision Tree (DT) algorithms have been developed over the years [18, 19, 29, 31] and after the success of deep learning, much research relates to combining DTs with neural networks. As a result, Soft Decision Tree (SDT) was introduced, allowing for the partial membership of a sample in the nodes that make up the tree structure [34], also trained in distillation setup [8]. This idea was also used in [13] that trains a set of classification trees and a backbone network in an end-to-end manner. Moreover, it was recently used in [24] to faithfully visualize the model using nodes with prototypes [6] instead of classifiers. Trees were also used to explain the previously trained black box models [5, 39]. More advanced methods automatically generate deep networks with a tree structure in a multi-step or an end-to-end manner [1, 2, 23, 38]. Many previous works were generalized by the ANT framework [36], which additionally enriched the tree structure with transformer, a nonlinear func-

tion that maps samples from the previous module to the next one, and allows for training the tree topology. In contrast to these methods, we design a simple and effective model for training decision graphs together with the parameters of the base neural network in the end-to-end manner.

2.2. Decision graphs

A decision graph is a well-studied classifier and has been used to solve many real-world problems [35]. When implemented as Directed Acyclic Graphs (DAG), it leads to accurate predictions while having lower model complexity, subtree replication, and training data fragmentation compared to decision trees [33]. However, most of the existing algorithms for learning DAGs involve training a conventional tree that is later manipulated into a DAG [7, 12, 25, 26] and, as such, are difficult to be directly adopted into neural networks. Hence, alternative approaches were proposed, like [4], which maintains the structure of the standard convolutional neural networks (CNNs) but uses additional routing losses at each layer to maximize the class-wise purity (like in growing decision trees) using data activation according to the class label distribution. Another method [37] introduces identity skip-connections similar to ResNets [9] that are executed or skipped depending on the gate response for an input. A similar gate mechanism was used in [22] to choose branches specialized for different inputs, whose outputs are combined to make the final predictions. Finally, [11] embeds infinitely many filters into low dimensional manifolds parameterized by compact B-splines and maximizes the mutual information between spline positions and class labels to specialize for classification tasks optimally. Such

a mechanism significantly reduces runtime complexity. In contrast to existing methods, SONG is a directed graph that can be adapted to any deep architecture and trained in an efficient gradient-based manner.

3. Self-organizing neural graphs

To adequately describe the Self-Organizing Neural Graph (SONG), we first define a more abstract structure that we call Soft Binary Directed Graph (SBDG). SBDG is considered binary because there are two alternative sets of edges, and soft because those sets are combined into one target set of edges depending on the input. Then, based on SBDG, we define SONG and describe how to use them as a decision model. Finally, we present method limitations and show how to overcome them with additional regularizers. The below definitions correspond to single-label classification for the clarity of description. However, they could be easily extended to other tasks, like multi-label classification or regression.

3.1. Method

3.1.1 Soft binary directed graphs

Soft Binary Directed Graph (SBDG) is a directed graph, which can be viewed as a probabilistic model. It is defined as graph $G = (V, E^0, E^1)$, with V corresponding to a set of nodes and E^0, E^1 corresponding to two alternative sets of edges, where:

- Set V contains two types of nodes:
 - internal nodes v_0, \dots, v_n , with v_0 specified as root r ,
 - leaves l_1, \dots, l_c , each exclusively associated with one class from set $\{1, \dots, c\}$,
- Set E^d , for $d \in \{0, 1\}$, contains all possible edges with weights m_{ji}^d corresponding to the probability of moving from node u_i to $u_j \in V$, as presented in Figure 2a. In the following, the aggregated probabilities of moving from node u_i to other nodes will be called a transition vector and denoted as m_i^d .
- If u_i is a leaf, then $m_{ji}^d = \delta_{ji}$ (Kronecker delta), which means that it is impossible to move out from the leaves.
- Each internal node u_i makes binary decisions $d \in \{0, 1\}$ with probabilities σ_i^d of using edges from set E^d .
- $\sigma_i^0 + \sigma_i^1 = 1$ and G can be transformed to a standard directed graph by combining m_i^0 and m_i^1 using the following formula for each node u_i : $\sigma_i^0 m_i^0 + \sigma_i^1 m_i^1$. This process is presented in Figure 2b.

Notice that if all transition vectors are binary, then after removing the edges with zero probability, SBDG becomes a binary directed graphs [28].

3.1.2 Self-organizing neural graphs

Self-Organizing Neural Graph (SONG) is a fully differentiable adaptation of SBDG that can be combined with various deep architectures. SONG is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E}^0, \mathcal{E}^1)$, where $\mathcal{V}, \mathcal{E}^0, \mathcal{E}^1$ implement V, E^0 , and E^1 of SBDG, and are obtained for input point x in the following way:

- The probability of decision $d = 1$ in node u_i is obtained as $\sigma_i^1(x) = \sigma(xw_i + b_i)$, where σ is the sigmoid logistic function, w_i is a filter function, and b_i is a bias¹.
- The probability of decision $d = 0$ equals $\sigma_i^0(x) = 1 - \sigma_i^1(x)$.
- The probability of moving from internal nodes is defined by two matrices $M^d = [m_{ji}^d] \in \mathbb{R}^{(n+c) \times n}$, for $d = \{0, 1\}$, with positive values and columns summing up to 1. In our implementation, we obtain such matrices by applying softmax to each of their columns.

Notice that $\{w_i\}_{i=1, \dots, n}$, $\{b_i\}_{i=1, \dots, n}$, M^0 , and M^1 are trainable parameters of the model.

Finally, we define a directed graph $\mathcal{G}_x = (\mathcal{V}, \mathcal{E})$ generated for input x where \mathcal{E} corresponds to the combination of matrices M^0 and M^1 :

$$M_x = \mathbb{1} \sigma_x^T \odot M^1 + \mathbb{1} (\mathbb{1} - \sigma_x)^T \odot M^0, \quad (1)$$

where $\sigma_x = [\sigma_0^1(x), \dots, \sigma_n^1(x)]^T$, symbol \odot denotes the Hadamard product, and $\mathbb{1}$ is the all-ones vector of dimension n .

3.1.3 Decision model

Matrix M_x contains the probability of moving from internal nodes to all nodes of the graph. However, to apply the theory of the Markov processes, it needs to be extended by columns corresponding to the leaves (as presented on the left side of Figure 3):

$$P_x = \begin{bmatrix} M_x & \begin{matrix} 0 \\ \vdots \\ 0 \end{matrix} \\ \hline & \mathbf{I} \end{bmatrix} \in \mathbb{R}^{(n+c) \times (n+c)}, \quad (2)$$

where $0 \in \mathbb{R}^{n \times c}$ is zero matrix and $\mathbf{I} \in \mathbb{R}^{c \times c}$ is an identity matrix. As a result, we obtain a square stochastic (transition) matrix used to describe the transitions of a Markov chain. While P_x contains the probability of moving from u_i to u_j in one time step, it can be easily used to obtain a similar probability for N steps by calculating the N -th power of P_x . Finally, the resulting matrix can be multiplied by vector $v = [1, 0, \dots, 0]^T$ to obtain the probability of moving from the root to any node of the graph, including leaves, whose probability is the output of the model. We present a simple example illustrating this process on the right side of Figure 3. More examples are provided in the Supplementary Materials.

¹In practice, this probability could also be obtained with any NN that ends with a sigmoid function.




Figure 2: SONG contains two alternative transition vectors $m_{\cdot i}^0$ and $m_{\cdot i}^1$ that aggregate the probability of moving from a particular node v_i to all other nodes. In (a), they are represented as dashed blue and solid red arrows, respectively. Each node obtains input data x and makes a binary decision with probabilities σ_i^0 and σ_i^1 of using one transition or another. As $\sigma_i^0 + \sigma_i^1 = 1$, SONG can be transformed to a standard directed graph by combining $m_{\cdot i}^0$ and $m_{\cdot i}^1$, as presented in (b). During training, both σ_i^j and $m_{\cdot i}^j$ are trained to obtain the optimal decision graph as presented in Figure 1 of the paper.




Figure 3: Construction of the transition matrix and successive steps of our Markov process. On the left, a graph with its matrices M^0 and M^1 is presented, followed by an exemplary decision vector σ_x and the resulting matrix P_x . On the right, the flow in a graph is depicted for 3 consecutive steps. At first, the probability is entirely placed in the root. However, in the next steps, the distribution splits between nodes according to the transition probabilities, reaching leaves in step 3. The probabilities in the leaves after all steps are class probabilities inferred by the model (the number of steps is considered as a method hyperparameter).

3.1.4 Algorithm complexity

The memory scales quadratically with the number of nodes n due to the necessity of storing transition matrices M^0 and M^1 . For a single image x , the computational complexity is bounded by the Hadamard product used to produce matrix M_x (and P_x consequently), which is $O(n^2)$. Considering also the number of steps N , vector v is N times multiplied by matrix P_x . Thus, the complexity is $O(n^2 + N \cdot (n + c)^2)$, where c is the number of leaves.

3.2. Regularizations

Similarly as in Soft Decision Trees (SDT) [8], we observe that our graphs require additional training regularizers. The

reasons for that are threefold. First, SONG may get stuck on plateaus in which one or more $\sigma_i^d(x)$ is 0 for all input samples x , and the gradient of the sigmoid logistic function for this decision is always very close to zero. Second, if SONG is uncertain of its predictions, it can safely hold the probabilities in internal nodes instead of moving them to leaves, which results in a small accumulated probability in the latter. Third, SONG tends to binarize what is positive in general, but if this binarization appears too early, the model can get stuck in a local minimum. Therefore, to prevent model degeneration, we introduce three types of regularization.

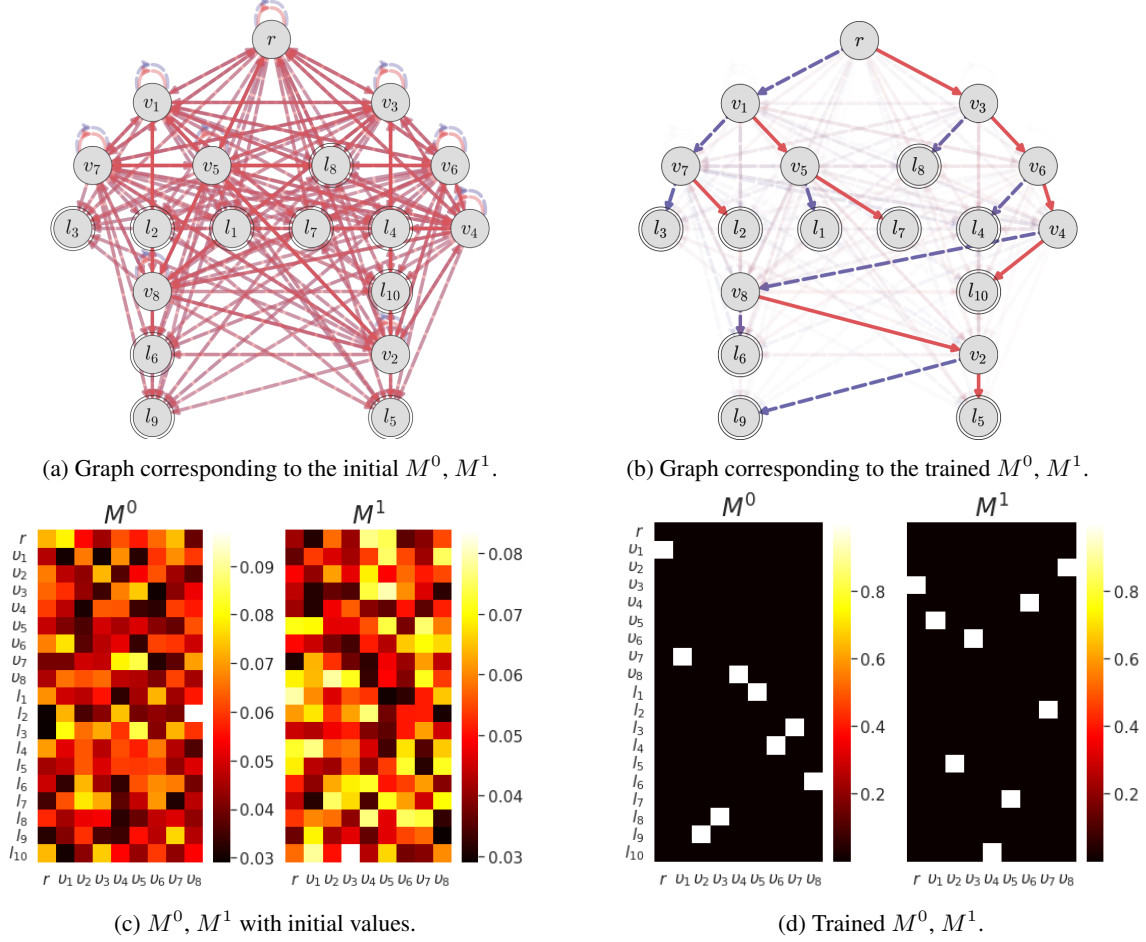


Figure 4: Visualisation of the graphs with edges corresponding to the transition matrices M^0 (blue edges) and M^1 (red edges) of the SONG before and after training on the MNIST dataset with 9 internal nodes. One can observe that SONG models binarize the connections during gradient training.

3.2.1 Node regularization

The node regularization is a direct adaptation of the approach proposed by [8]. It is used to avoid getting stuck at poor solutions by encouraging each internal node to make equal use of both left and right subtrees. In our approach, this regularization encourages each internal node to make equal use of both sets of edges \mathcal{E}^0 and \mathcal{E}^1 . I.e., to send half of the training samples to one direction (using M_0) and half of them to the other direction (with M_1). For this purpose, we calculate the cross entropy between the desired average distribution 0.5, 0.5 for those two sets and the actual average distribution $\alpha_{i,s}, \beta_{i,s}$ in node v_i at step s

$$L_{nodes} = -\frac{\lambda}{2} \sum_{i=1}^n \log(\alpha_{i,s}) + \log(\beta_{i,s}),$$

where

$$\alpha_{i,s} = \frac{\sum_{x \in B} (P_x^s r)_i \cdot (\sigma_i^1(x))^\gamma}{\sum_{x \in B} (P_x^s r)_i},$$

$$\beta_{i,s} = \frac{\sum_{x \in B} (P_x^s r)_i \cdot (\sigma_i^0(x))^\gamma}{\sum_{x \in B} (P_x^s r)_i},$$

B is a batch of samples used in an iteration, $\gamma \in [1, 2]$, and $(P_x^s r)_i$ corresponds to i th coordinate of vector $(P_x^s r)$. One can observe that our node regularizer is calculated per node and step, and it is different from [8], where additional loss is computed once for each node. Moreover, we penalize model for making uncertain decisions ($\sigma_{i,s}(x) \approx 0.5$) using the parameter γ .

3.2.2 Leaves regularization

The leaves regularization, enforcing the summary probabilities in leaves to be close to 1, is defined as

$$L_{leaves} = -\log \left(\sum_{i=n}^{n+c} (P_x^N r)_i \right), \quad (3)$$

where n is the number of nodes (excluding root indexed with 0), c is the number of leaves (classes), and N is the number

Method	Ex	SO	EE	MNIST	C10	C100	TIN
DDN (NiN)	✗	✓	✗	-	90.32	68.35	-
DCDJ (NiN)	✗	✓	✓	-	-	69.00	-
ANT-A* (n/a)	✓	✓	✗	99.36	93.28	-	-
ResNet18	✗	✗	✗	98.91	94.93	75.82	63.05
DNDF	✗	✗	✗	97.20	94.32	67.18	44.56
DT	✓	✗	✗	-	93.97	64.45	52.09
NBDT	✓	✗	✗	-	94.82	77.09	64.23
NBDT w/o h.	✓	✓	✗	-	94.52	74.97	-
RDT	✓	✓	✓	-	93.12	-	-
SONG (ours)	✓	✓	✓	98.81	95.62	76.26	61.99

Table 1: Comparison of models with deep architecture in terms of model features and accuracy on MNIST, CIFAR10 (C10), CIFAR100 (C100), and TinyImageNet (TIN). ResNet18 was used to extract the vector representation of input images for DNDF [13], DT, NBDT (with and without hierarchy) [38], RDT [2], and SONG. For DDN [23], DCDJ [4], and ANT-A [36], the backbone models are provided in the brackets. “Ex” indicates if the method retains properties such as pure leaves, sequential decisions, and non-ensemble. “SO” indicates if the model is self-organized (does not require a predefined structure). “EE” indicates if the structure and weights of model are trained in an end-to-end continuous manner.

of steps.

3.2.3 Gumbel-softmax

We use Gumbel-softmax [10] instead of softmax to each column of matrices M_0 and M_1 to explore the trajectories of the graph better. In other words, Gumbel-softmax introduces randomness, which results in a wider exploration of the graph structure in the optimization process.

4. Theoretical analysis

In the supplementary materials, we provide theoretical results, showing that the graph structure generated by SONG is binarized during training, as presented in Figure 4.

5. Experiments

In this section, we analyze the accuracy of the SONGs trained on Letter [3], Connect4 [3], MNIST [17], CIFAR10 [14], CIFAR100 [14], and TinyImageNet [16] datasets and compare it with the state of the art methods [2, 4, 13, 23, 36, 38]. We examine how the number of nodes and steps influence the structure of graphs, the number of internal nodes used by the model, the number of back edges, and the distance from the root to leaves. Moreover, we explain how the probability of back and cross edge

Method	Letter	Connect4	MNIST
SDT w/o distillation [8]	78.00 (511)	78.63 (255)	94.45 (255)
SDT [8]	81.00 (511)	80.60 (255)	96.76 (255)
SONG-S-large (ours)	86.25 (511)	82.82 (255)	95.74 (255)
SONG-S-small (ours)	82.95 (64)	80.27 (8)	94.66 (64)

Table 2: Comparison of SDT [8] and shallow SONG (SONG-S) on three datasets, where shallow corresponds to direct flattened inputs (no backbone network used). The accuracy of each model is reported along with the number of internal nodes specified in the parentheses. SONG-S-small contains the minimal number of nodes necessary to match the accuracy of SDT. SONG-S-large uses the same number of internal nodes as SDT. Please notice that SONG models are trained without a distillation mechanism, and they always obtain better results than SDT without distillation.

changes in the successive training steps. Finally, we provide a detailed comparison with SDT [8] and present sample graphs obtained for the MNIST dataset. In all experiments, we use leaves normalization and Gumbel-softmax, and we treat node regularization as a hyperparameter of the model. While this section presents only the most important findings for the sake of clarity, the experimental setup and detailed results can be found in the Supplementary Materials.

5.1. SONG in deep learning setup

In the first experiment, we apply SONG² at the top of the backbone Convolutional Neural Network (CNN) without the final linear layer. CNN takes the input image and generates the representation, which is passed to the SONG. SONG processes the representation and returns the predictions for each class, which are then used with target labels to calculate Binary Cross-Entropy (BCE) loss. As a backbone network, we use ResNet18 for all datasets except MNIST, for which we employ a smaller network (see Supplementary Materials for details).

As presented in Table 1, our method matches or outperforms most of the recent state-of-the-art methods. On CIFAR10, SONG accuracy outperforms all baseline by almost 1 percentage point. On MNIST, it is worse than ANT [36] by around 0.5%, and on CIFAR100 and TinyImageNet, NBDT [38] achieves better results. However, both ANT and NBDT are not trained in an end-to-end continuous manner. Moreover, NBDT requires a hierarchy provided before training, and without such a hierarchy, it obtains accuracy more than 1% lower than SONG on CIFAR100.

5.2. SONG as shallow model

Although SONG can be successfully used in a deep learning setup, it can also be treated as a shallow model. In

²We made the code available at: <https://github.com/gmum/SONGs>

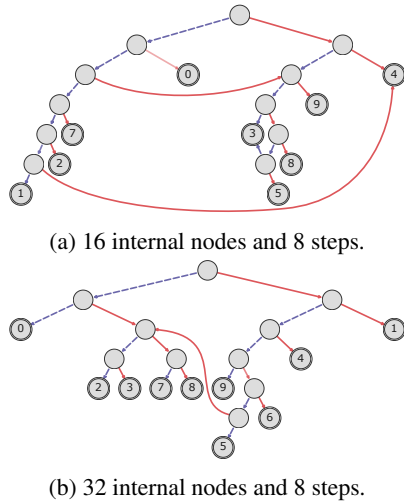


Figure 5: Examples of the graph structures obtained by training SONG on the MNIST dataset. The root is the top-most node in each graph, and double node borders denote the leaves with numbers of the MNIST classes. For each node v_i , we present two edges corresponding to the highest probability from two transition vectors $m_{\cdot,i}^0$ and $m_{\cdot,i}^1$ (represented as dashed blue and solid red arrows, respectively).

this case, SONG directly processes an input sample and returns the predictions passed with target labels to BCE loss. This setup is similar to the one presented in experiments on SDTs [8]. Hence, we compare to SDT on all datasets considered in [8].

Table 2 shows that SONG obtains better results than SDT without distillation on all datasets. Moreover, on Letter and Connect4, SONG outperforms even SDT with distillation. We also observe that SONG requires fewer nodes than SDT and obtains on par results on the Connect4 dataset with 30 times fewer nodes. For Letter and MNIST, similarly good results can be obtained with 30 times fewer nodes. This finding is in line with [33] which shows that decision graphs require dramatically less memory while considerably improving generalization.

5.3. SONG structure

As a fully differentiable model, SONG strengthens or weakens an edge between any pair of nodes during training to constantly optimize the graph’s structure (see Figure 4). Consequently, it can generate any structure that uses all available nodes, or only some of them. In particular, the final structure may be a binary tree or contain back edges. Moreover, the distance from the root to leaves can vary. This variability is visualized in Figure 5, where we present two graphs obtained for MNIST using a different number of internal nodes and steps.

In Figure 6, we provide statistics on multiple SONGs

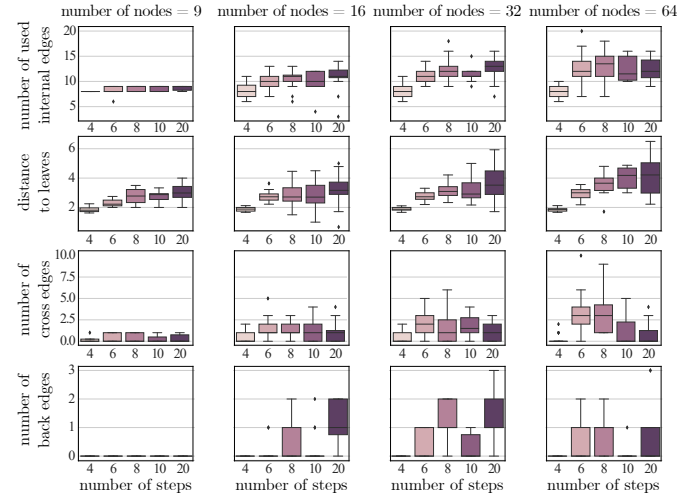


Figure 6: Nodes and edges statistics calculated for SONGs trained on the CIFAR10 dataset. For each combination of the number of internal nodes and steps, 20 graphs are trained and used to plot the distributions of four statistics.

generated for the CIFAR10 dataset. We observe a significant difference in SONG structure depending on the number of internal nodes and steps. First, we note that the number of internal nodes used by the model increases with the increasing number of steps N , and it does not depend on the total number of internal nodes n . As a natural consequence, a similar trend is observed for the distance from the root to the leaves. When it comes to back edges, their number is relatively small, and they appear only for a larger number of steps. At the same time, the cross edges are more often and increase with the increased number of internal nodes.

We also analyze the relationship between the number of nodes and steps and prediction accuracy SONG. As can be seen in Table 3, the performance on MNIST constantly increases as the number of nodes and steps grow. This is not the case of CIFAR-10, where the accuracy remains similar for all combinations of the parameters. It can be caused by the smaller dimension of the representation vector in MNIST (50) than in CIFAR10 (512).

5.4. SONG structure during training

We analyze the relationship between BCE loss and the probability of back and cross edges in the successive epochs of the training.

We present the mean over multiple models and all test samples (as each test sample x has its graph represented by matrix P_x). We observe that the probability of back edges decreases together with decreasing BCE loss, both for simple MNIST and more complicated CIFAR100 dataset (see Figure 7).

Moreover, in Figure 8, we present the mean distances between transition matrices (P_x) obtained for samples of


nodes	steps				
	4	6	8	10	20
9	95.66	97.29	97.25	97.95	97.56
16	97.31	97.83	98.23	98.43	98.56
32	96.82	97.74	98.35	98.65	98.62
64	96.29	98.12	98.12	98.47	98.68

(a) MNIST.


nodes	steps				
	4	6	8	10	20
9	94.48	94.86	94.92	94.94	94.93
16	94.88	94.95	94.86	94.87	94.89
32	94.99	94.95	94.95	94.90	94.98
64	94.90	94.87	94.88	94.94	94.93

(b) CIFAR10.

Table 3: Results of SONG in a deep learning setup. One can observe that for the MNIST dataset (a), the performance increases with the increasing number of nodes and steps. In contrast to CIFAR10 (b), where the performance is relatively similar for all combinations of the parameters.




(a) SONG with 256 internal nodes and 10 steps trained for CIFAR100.



(b) SONG with 64 internal nodes and 10 steps trained for MNIST.

Figure 7: BCE loss as well as the number of back and cross edges in the successive training epochs of SONG. One can observe that number of back edges decrease together with decreasing BCE loss.



(a) 16 internal nodes.

(b) 64 internal nodes.

Figure 8: Mean distances between transition matrices P_x for pairs of MNIST input samples represented by a distance matrix (the larger distance, the brighter color). The rows and columns correspond to 0-9 digits.

the same and different classes. One can observe, among others, that the diagonal is visibly darker than the rest of the matrix, which means that inputs from the same class have more similar transition matrices. This confirms that we obtain similar transition matrices for similar inputs.

6. Conclusions

In this work, we introduce Self-Organizing Neural Graphs (SONGs), a new type of decision graphs applicable in any deep learning pipeline. They optimize their structure during training by strengthening or weakening graph edges using gradient descent. Thanks to the graph structure, SONG can reuse the decision nodes and obtain state-of-the-art results with a significantly smaller number of nodes than existing methods. Moreover, the introduced general paradigm based on Markov processes allows for efficient training, and SONG converges to the binary acyclic directed graphs. Hence, we believe that our work opens a plethora of research pathways towards more effective applications of decision graphs in a deep learning setup.

Acknowledgments

The work of J. Tabor and B. Zieliński is supported by the National Centre of Science (Poland) Grant No. 2021/41/B/ST6/01370, the work of Ł. Struski is supported by the National Centre of Science (Poland) Grant No. 2020/39/D/ST6/01332, and the work of T. Danel is supported by the National Centre of Science (Poland) Grant No. 2020/37/N/ST6/02728. The research of M. Śmieja was carried out within the research project “Bio-inspired artificial neural network” (grant no. POIR.04.04.00-00-14DE/18-00) within the Team-Net program of the Foundation for Polish Science co-financed by the European Union under the European Regional Development Fund.

References

- [1] Karim Ahmed, Mohammad Haris Baig, and Lorenzo Torresani. Network of experts for large-scale image categorization. In *European Conference on Computer Vision*, pages 516–532. Springer, 2016.
- [2] Stephan Alaniz, Diego Marcos, Bernt Schiele, and Zeynep Akata. Learning decision trees recurrently through communication. In *34th IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2021.
- [3] Arthur Asuncion and David Newman. Uci machine learning repository, 2007.
- [4] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Deep convolutional decision jungle for image classification. *arXiv preprint arXiv:1706.02003*, 2017.
- [5] Osbert Bastani, Carolyn Kim, and Hamsa Bastani. Interpreting blackbox models via model extraction. *arXiv preprint arXiv:1705.08504*, 2017.
- [6] Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin. This looks like that: deep learning for interpretable image recognition. *arXiv preprint arXiv:1806.10574*, 2018.
- [7] Philip A. Chou. Optimal partitioning for classification and regression trees. *IEEE Computer Architecture Letters*, 13(04):340–354, 1991.
- [8] Nicholas Frosst and Geoffrey Hinton. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784*, 2017.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [11] Cem Keskin and Shahram Izadi. Splinenets: Continuous neural decision graphs. *arXiv preprint arXiv:1810.13118*, 2018.
- [12] Ron Kohavi and Chia-Hsin Li. Oblivious decision trees, graphs, and top-down pruning. In *IJCAI*, pages 1071–1079. Citeseer, 1995.
- [13] Peter Kotschieder, Madalina Fiterau, Antonio Criminisi, and Samuel Rota Buló. Deep neural decision forests. In *Proceedings of the IEEE international conference on computer vision*, pages 1467–1475, 2015.
- [14] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [16] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7:7, 2015.
- [17] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database, 2010.
- [18] Wei-Yin Loh. Classification and regression trees. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 1(1):14–23, 2011.
- [19] Manish Mehta, Rakesh Agrawal, and Jorma Rissanen. Sliq: A fast scalable classifier for data mining. In *International conference on extending database technology*, pages 18–32. Springer, 1996.
- [20] Tomáš Mikolov, Stefan Kombrink, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Extensions of recurrent neural network language model. In *2011 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5528–5531. IEEE, 2011.
- [21] Frederic Morin and Yoshua Bengio. Hierarchical probabilistic neural network language model. In *Aistats*, volume 5, pages 246–252. Citeseer, 2005.
- [22] Ravi Teja Mullapudi, William R Mark, Noam Shazeer, and Kayvon Fatahalian. Hydranets: Specialized dynamic architectures for efficient inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8080–8089, 2018.
- [23] Venkatesh N Murthy, Vivek Singh, Terrence Chen, R Manmatha, and Dorin Comaniciu. Deep decision network for multi-class image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2240–2248, 2016.
- [24] Meike Nauta, Ron van Bree, and Christin Seifert. Neural prototype trees for interpretable fine-grained image recognition. *arXiv preprint arXiv:2012.02046*, 2020.
- [25] Arlindo L Oliveira and Alberto Sangiovanni-Vincentelli. Using the minimum description length principle to infer reduced ordered decision graphs. *Machine Learning*, 25(1):23–50, 1996.
- [26] Jonathan Oliver. *Decision graphs: an extension of decision trees*. Citeseer, 1992.
- [27] JJ Oliver. Decision graphs - an extension of decision trees. In *Proc. 4th International Conference on Artificial Intelligence and Statistics, Miami, FL, 1993*, 1993.
- [28] John C Platt, Nello Cristianini, John Shawe-Taylor, et al. Large margin dags for multiclass classification. In *nips*, volume 12, pages 547–553, 1999.
- [29] J Ross Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.
- [30] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- [31] John Shafer, Rakesh Agrawal, and Manish Mehta. Sprint: A scalable parallel classifier for data mining. In *Vldb*, volume 96, pages 544–555. Citeseer, 1996.
- [32] Jamie Shotton, Ross Girshick, Andrew Fitzgibbon, Toby Sharp, Mat Cook, Mark Finocchio, Richard Moore, Pushmeet Kohli, Antonio Criminisi, Alex Kipman, et al. Efficient human pose estimation from single depth images. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2821–2840, 2012.
- [33] Jamie Shotton, Toby Sharp, Pushmeet Kohli, Sebastian Nowozin, John Winn, and Antonio Criminisi. Decision jungles: Compact and rich models for classification. 2016.
- [34] Alberto Suárez and James F Lutsko. Globally optimal fuzzy decision trees for classification and regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(12):1297–1311, 1999.
- [35] Hiroki Sudo, Koji Nuida, and Kana Shimizu. An efficient private evaluation of a decision graph. In *International Conference on Information Security and Cryptology*, pages 143–160.

- Springer, 2018.
- [36] Ryutaro Tanno, Kai Arulkumaran, Daniel Alexander, Antonio Criminisi, and Aditya Nori. Adaptive neural trees. In *International Conference on Machine Learning*, pages 6166–6175. PMLR, 2019.
 - [37] Andreas Veit and Serge Belongie. Convolutional networks with adaptive inference graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
 - [38] Alvin Wan, Lisa Dunlap, Daniel Ho, Jihan Yin, Scott Lee, Henry Jin, Suzanne Petryk, Sarah Adel Bargal, and Joseph E Gonzalez. Nbd: neural-backed decision trees. *ICLR 2021*, 2020.
 - [39] Quanshi Zhang, Yu Yang, Haotian Ma, and Ying Nian Wu. Interpreting cnns via decision trees. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6261–6270, 2019.