DecisioNet: A Binary-Tree Structured Neural Network

Noam Gottlieb[0000−0002−2561−432X] and Michael Werman[0000−0002−0665−967X]

The Hebrew University of Jerusalem, Israel
{noam.gottlieb2,michael.werman}@mail.huji.ac.il

Abstract. Deep neural networks (DNNs) and decision trees (DTs) are both state-of-the-art classifiers. DNNs perform well due to their representational learning capabilities, while DTs are computationally efficient as they perform inference along one route (root-to-leaf) that is dependent on the input data. In this paper, we present DecisioNet (DN), a binary-tree structured neural network. We propose a systematic way to convert an existing DNN into a DN to create a lightweight version of the original model. DecisioNet takes the best of both worlds - it uses neural modules to perform representational learning and utilizes its tree structure to perform only a portion of the computations. We evaluate various DN architectures, along with their corresponding baseline models on the FashionMNIST, CIFAR10, and CIFAR100 datasets. We show that the DN variants achieve similar accuracy while significantly reducing the computational cost of the original network.

Keywords: Neural Network Optimization · Decision Trees.

1 Introduction

Deep neural networks (DNNs) have achieved exceptional performance in various visual recognition tasks in recent years, such as image classification, object detection, and semantic segmentation. That is mostly due to their representational learning capabilities. However, deploying DNN models in an industrial environment is challenging - especially when the computational resources are low (which is the case for many mobile device applications) or where the model's inference time has to be fast enough (*e.g.*, real-time applications). In addition, a DNN is seen in many cases as a "black box" - one cannot easily figure out *why* a final prediction is made.

Another powerful machine learning model is the Decision Tree (DT). A DT model learns a routing function, where each node of the tree routes the data to one of its children until it reaches a leaf with the final output. This conditional computation means that only part of the DT is visited for each input thus achieving high efficiency. Moreover, DTs provide an interpretable structure, allowing the user to understand why a decision was made. On the downside, DTs usually require hand-engineered data features, and they cannot be trained with gradient-based optimization methods, which limits their expressiveness.

In this paper, we propose a novel general model with the benefits of both DNNs and DTs - the DecisioNet (DN). This is a binary-tree structured DNN, derived from any other DNN that we wish to reduce its computational cost. It can be trained end-to-end using backpropagation [18] just like any other DNN. In addition, the DN has routing modules which play the role of the DT within this model, routing the input through the tree. The outcome is a lighter model - in terms of parameters and even more in terms of computational cost - whose performance is at par with the baseline model. We evaluate and compare full baseline models against their DN variants on the FashionMNIST [22], CIFAR10, and CIFAR100 [14] datasets. In this paper, we examine DN only for image classification tasks but this method can be used for other types of tasks as well.

Contributions the contributions of this paper are: i) we propose a systematic way of transforming an existing DNN into a tree-structured version of it (its DecisioNet), yielding a lightweight model with comparable performance. ii) we propose a way of training the DN end-to-end, despite the explicit data routing functions within the DN (which involves non-differentiable operations) using Improved Semantic Hashing. iii) we provide an open-source PyTorch [16] implementation of DecisioNet, available at https://github.com/noamgot/DecisioNet.

2   Related Work

Soft Decision Trees The Soft Decision Tree (SDT) mechanism has been studied in various works, including [20,10,17,5]. SDT is a DT with neural routing nodes sending the data to a sub-tree, multiplied by a probabilistic factor in the range $[0,1]$. The final prediction of the tree is the weighted sum of all leaves, where the weight of each leaf is the probability of arriving at it.

This method differs from traditional DTs, whose decisions are binary and deterministic. Another SDT-like method is the Neural Decision Forest (NDF), proposed in [13]. This method achieved great performance on the Imagenet dataset [4], replacing the last fully-connect layer of a neural network with decision forest nodes. These nodes yield a prediction which is a weighted sum of all the trees' predictions. The major drawback of SDTs is that all paths of the tree must be executed to get a prediction. These methods lack the advantage of traditional decision trees using only a single computational path, and therefore their efficiency is limited. In this work, we focus on hard decision trees only, which perform only one path of the tree.

Neural Trees With Hard Routing Another type of neural DT is proposed in [8,6], where the forward pass utilizes the routing nodes to make a hard decision, and in this way indeed only the relevant nodes of the tree are visited. [8] introduce a routing function which outputs binary values at test time only, such that the tree still performs soft decisions during training, allowing it to train end-to-end.

[6] on the other hand introduce the Tree Ensemble Layer (TEL), which is capable of performing hard decisions at test time and even during training, by simply skipping unreachable nodes (*i.e.*, nodes whose ancestors reach probability is 0). Notice that the number of reachable leaves is not necessarily limited to 1, hence a soft routing is still performed. In contrast, our method applies fully hard routing (with a single output leaf) both at training time (where we also use soft decision behavior) and at test time.

Tree-Structured Neural Networks Instead of replacing DTs decision nodes with neural modules, the other direction is also possible: creating a DNN with a structure of a DT, such that the routing behavior of a DT is present. This network consists of multiple sub-networks, that are traversed conditionally based on the routing nodes' decisions. [23] proposed HD-CNN, where a small CNN is first activated to classify inputs into coarse categories. This network has a hierarchical architecture which is essentially an SDT-structured DNN with a single decision node with k branches. HD-CNN is not trained end-to-end but in a modular way: there is a common (shared) network for predicting coarse categories and coarse-category expert networks. Each of these networks is trained alone (the expert modules are trained only with their corresponding classes; The coarse-category labels are achieved by clustering similar classes based on an existing network), and all of them are fined tuned together in the end. [9] extended the tree-structured net of HD-CNN to conditional networks. These are DAG-based CNN architectures with data routing. They distinguish between two types of routing: i) explicit routing - where data is conditionally passed to a node's children (one or more) based on a routing function. This method is similar to DTs. ii) implicit routing - where the data is split into portions that are sent to the node's children unconditionally. In the paper. this method is done using filter groups (*i.e.*, splitting outputs feature maps into groups, where each group goes to a separate route). Another tree-structured NN is Adaptive Neural Trees (ANT) [21]. In ANT, the tree structure is learned together with the model's weights; This approach is different from tree-structured NNs with a static architecture that is commonly used (in this work as well).

Differences and Similarities DecisioNet is, to some extent, a combination of HD-CNN [23] and conditional networks [9]. We create labels for training the routing modules with a method that's based on the one proposed in HD-CNN, but unlike [23], we train deeper, tree-structured DNNs with explicit routing, and we do the training end-to-end. We use explicit routing, similar to the one that was suggested for conditional networks. However, we use a different method for training end-to-end (improved semantic hashing [11]), and more importantly, we show actual results for explicitly-routed DNN trees (in [9] the main results were achieved using implicit routing; The only result that contained explicit routing was a CNN ensemble with a single decision for choosing which models to use for the final prediction). Finally, the approach used by ANT [21] is building the tree-structured net from scratch; We aim to optimize the computational cost of an *existing network* by transforming its structure into a tree.

3   DecisioNet

The proposed model is based on any existing deep neural network (DNN) such as VGG [19], ResNet [7], *etc*. We'll denote such a DNN as the "baseline model".

3.1   Architecture

The DN architecture is based directly on the baseline model's architecture. It is a transformation of the baseline model into a binary decision tree, whose primary goal is to achieve similar performance with fewer computational operations. We begin by choosing *split points* - these are the places the new DN model will decide to route the data into one of two routes. In our method, after the i-th split (Starting with $i = 1$), we split the relevant layers to 2^i equal portions, in terms of the number of filters in convolutional layers, *etc*. A toy example of this idea is displayed in figure 1. This way we get a balanced tree whose number of parameters is fewer than that of the baseline model. Moreover, since at test time our model chooses a single path from the root to a leaf, traversing only through a portion of the nodes (just like a DT), the number of multiply-accumulate (MAC) operations is also smaller. The routing at each split is done using a trainable *routing module* (RM). The RM is a small, lightweight network that makes a binary decision (a detailed explanation is found in section 3.3).


(a) Baseline network   (b) DecisioNet
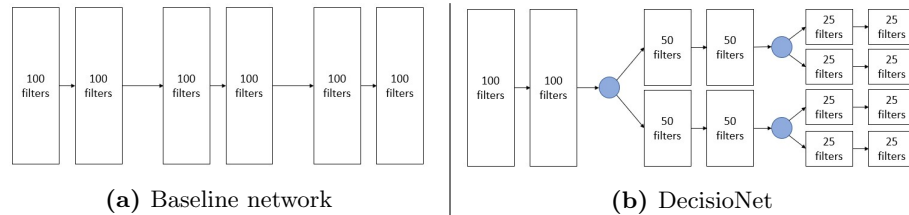
Fig. 1. A toy example of transforming a baseline model (a) into its respective DecisioNet (b). For the sake of clarity, the baseline model consists of 6 convolutional layers, each with 100 filters. The blue circles represent routing modules. In the DN model, the layers after the first split (leftmost blue circle) have half the number of filters compared to their corresponding layers in the baseline model. Similarly, the layers after the second split(s) have a quarter the number of filters. This idea can be easily generalized to an arbitrary number of splits. At test time, each RM picks one of two routes and the data is passed to the next module through this route only.

3.2   Classes Hierarchical Clustering

After deciding the depth of the DecisioNet (*i.e.*, the number of splits), we perform hierarchical clustering of the dataset's classes to obtain intermediate labels for

the routing modules. The goal of this phase is to extract the hierarchical structure of the data, grouping similar classes into nested clusters. The routing labels are derived directly from the clustering. The general idea is demonstrated in figure 2.
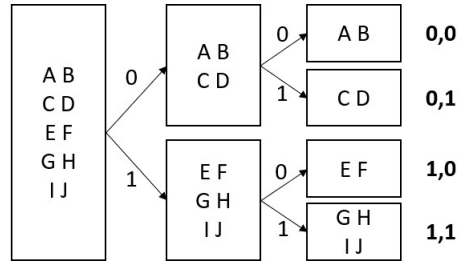
Fig. 2. Hierarchical clustering example - 10 classes (A-J) are clustered into a clustering tree of depth 2. The deepest (rightmost) clusters share the same routing labels, stated to their right. For example, classes E and F labels are 1 for the first routing, and 0 for the second.

We use a method similar to the one used by [23], extending it to handle more than just 2 hierarchical levels (coarse and fine), using a different clustering method. We begin by randomly sampling a balanced subset of images out of the training set. We evaluate the held-out set using the baseline model to obtain a confusion matrix $\mathbf{F}$. We define a distance matrix $\hat{\mathbf{D}}$ as follows:

$$\hat{\mathbf{D}}_{ij} = \begin{cases} 0 & i = j \\ 1 - \mathbf{F}_{ij} & i \neq j \end{cases} \tag{1}$$

To obtain a symmetric distance matrix, we define:

$$\mathbf{D} = \frac{1}{2} \left(\hat{\mathbf{D}} + \hat{\mathbf{D}}^T \right) \tag{2}$$

At this point, $\mathbf{D}_{ij}$ measures the similarity between classes i and j. Having a distance matrix, we use it to perform hierarchical agglomerative clustering. This phase gives us a division of the dataset's classes into 2 clusters, a division of each of these clusters into 2 sub-clusters, and so on. This can be seen as a mapping of a class to its set of clusters $F_c : \{i\}_{i=1}^{C} \to \{0,1\}^k$, where C is the number of classes and k is the depth of the DecisioNet tree.

Unlike [23], we do not allow overlap between same-level clusters, *i.e.*, at each clustering depth, each class can be found in exactly one sub-cluster. The main problem with this choice is that images that are routed to a wrong branch (at any level, even in the deepest routing module) end up in the wrong leaf and are miss-classified. We overcome this problem by allowing each leaf of the DN to predict all classes - even classes that shouldn't have been routed to this leaf. Allowing cluster overlap is possible though, and we leave this for future research.

3.3   Routing

One of the key questions when dealing with tree-structured neural networks is which routing method to use. There are 2 main types of routing found in the literature:

1. *Soft routing* - in this method, the output of each branch is multiplied by a real value (usually between 0 and 1) and the final output of some routing point is the sum of its branches' outputs.
2. *Hard routing* - in this method, each branch either passes its input or not.

The main advantage of the hard-routing method (and the disadvantage of the soft-routing method), is that it allows us to save computational cost by performing only the chosen path of a given input. To decrease the computational cost, DN uses the hard-routing approach.[1]

The Routing Module To choose a route for input, we use a routing module (RM) - this module is a small efficient neural network which outputs a single binary value. Inspired by [3,2], we used a similar routing module. These papers used this module to choose filters of a convolutional layer, hence they needed multiple outputs; However, we use this module to choose a *single* branch, so we use a slightly modified version of it.

Let $x \in \mathbb{R}^{C \times H \times W}$ be the output of the last layer before a split in the tree. Our routing module can be defined as follows:

$$RM(x) = B\left(FC\left(GAP\left(x\right)\right)\right) \tag{3}$$

Where GAP is a Global Average Pooling operation, defined as:

$$GAP(x) = \frac{1}{HW} \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} x_{ij} \tag{4}$$

FC is a linear projection layer (*i.e.*, a fully-connected layer) with C inputs and a single output and B is a binarization function which will be introduced in the next paragraph. A schematic diagram of the proposed routing module can be found in figure 3.

Improved Semantic Hashing We want our module to output a binary value - either 0 or 1 - to choose the next computational path. Performing simple binarization using a threshold is not an option, as this operation is almost always with zero derivative and therefore we will not be able to perform backpropagation properly. Like [3,2], we adopt the Improved Semantic Hashing method, proposed by [11].

---

[1] To the best of our knowledge, popular deep learning frameworks, such as PyTorch [16] and TensorFlow [1], do not support tree-structured neural nets when using batches of inputs. It means that in practice, during training and evaluating the data passes through all the nodes of the tree where we zero out the "untraversed" ones

Routing Module
Global Average Pooling
FC-1
Binarization
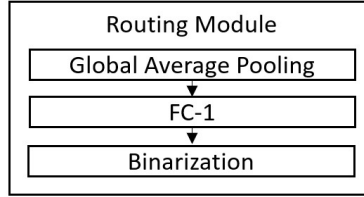
Fig. 3. A schematic diagram of the proposed routing module. FC-1 is a fully-connected layer with a single output neuron. Full details are in sec. 3.3

During training, we draw a random noise ϵ with mean 0 and standard deviation 1. We define the following values:

$$g_\epsilon(z) = z + \epsilon$$
$$g_r(z) = \sigma'\left(g_\epsilon(z)\right) \tag{5}$$
$$g_b(z) = \mathbb{1}\left(g_\epsilon(z) > 0\right)$$

$\mathbb{1}(z)$ is an indicator function (evaluated to 1 when z is true and 0 otherwise); σ' is the saturating sigmoid function from [12]:

$$\sigma'(z) = \max\left(0, \min\left(1, 1.2\sigma(z) - 0.1\right)\right) \tag{6}$$

where σ is the well-known sigmoid function. Notice that g_r is a real-valued number in the range $[0, 1]$ while g_b is a binary value; Moreover, the gradient of g_r is well defined w.r.t g_ϵ, while the gradient of g_b w.r.t g_ϵ is 0.

Finally, we can define our binarization function B as follows:

$$B(z) = g_{b/r}(z) \tag{7}$$

The term $g_{b/r}$ denotes that in the forward pass we use either g_r or g_b, and that is eventually the output of the RM. The choice is done at random with a 50% chance for each. In the backward pass, we *always* use the gradients of g_r to backpropagate meaningful gradients[2]. To allow this dual behavior, we sum the branches' outputs in the following way:

$$f(x) = (1 - (RM(x))) f_L(x) + RM(x) f_R(x) \tag{8}$$

Where f_L, f_R are the functions applied to x by the left and right branches. Notice that if RM outputs a binary value, only one of the branches is used, while in the real-valued case the branches' outputs are mixed (*i.e.*, soft routing is applied).

During evaluation and inference, we perform the same procedure with 2 changes: first, there is no additive noise ($\epsilon = 0$); second, the routing module only outputs binary values, so in this phase, we only use g_b as an output of RM.

To conclude, the Improved Semantic Hashing method allows us to train the DecisioNet model in an end-to-end manner and to apply hard decisions at inference time.

---

[2] In PyTorch, this operation can be performed like this:
```
out = g_b + g_r - g_r.detach()
```

3.4   Loss Function

DecisioNet is trained with 2 sets of labels - classification labels and routing labels (extracted in the clustering phase). During training, we want our model to classify inputs correctly while passing classes to their desired routes, based on the routing labels. Therefore, we design a loss function that balances these 2 demands. This loss is described as follows:

$$\mathcal{L} = \mathcal{L}_{cls} + \beta\mathcal{L}_{\sigma} \tag{9}$$

Where $\mathcal{L}_{cls}$ is a classification loss (*e.g.*, cross-entropy loss) and $\mathcal{L}_{\sigma}$ is a MSE loss for the routing labels. β is a hyper-parameter for balancing the classification accuracy and the routing accuracy. We also tried to use a MSE variant which penalizes routing mistakes with different weights depending on the RM depth. The assumption was that it would make sense to give a higher weight to routing mistakes that happen earlier. However, we did not find evidence that this method gives a significant improvement, hence we use the vanilla MSE instead.

4   Experiments

4.1   FashionMNIST & CIFAR10

Fashion-MNIST [22] consists of a training set of 60K examples and a test set of 10K examples. Each example is a 28×28 grayscale image, associated with a label from 10 classes. The CIFAR10 dataset [14] consists of 60K 32×32 color images in 10 classes, with 6000 images per class. There are 5000 training images and 1000 test images per class. Both datasets were trained using their training images and evaluated on the test images.

Model We used Network-In-Network (NiN) [15] as our baseline model. The architecture we used is displayed in figure 4a. We treat it as a stack of 3 small networks (which we refer to as "blocks"). The DecisioNet variants of this model are generated using the method described in section 3. We experimented with 4 varieties of DecisioNets:

1. *DN2* - A basic version with 2 splits, located in-between the baseline NiN model blocks. An example of a single DN2 computational path is displayed in figure 4b. Since there are 2 splits in the DN2 tree, there are 4 different computational paths (from root to leaf).
2. *DN1-early* - this version contains a single split point, between the first and second NiN blocks. It is equivalent to the DN2 without the deeper routing modules.
3. *DN1-late* - this version contains a single split point, between the second and third NiN blocks. It is equivalent to the DN2 without the first routing module and replacing the 2 deeper RMs with a single one.

4. *DN2-slim* - A slimmer version of DN2, where we "push" the routing modules to earlier stages compared to DN2. It results in an even lighter version of DN2 (fewer parameters and MAC operations). An example of a single DN2-slim computational path is displayed in figure 4c.

In all cases, the last layer of the DN model outputs 10 values. It means that classes that reached the wrong leaf can still be predicted correctly. This way we overcome the aforementioned choice of not allowing overlapping between clusters.

Preprocessing For each dataset we calculate its training set mean and standard deviation values (per channel) and use them to normalize the images before feeding them into the model. We also created the routing labels using the method described in 3.2. For CIFAR10, we conduct 2 sets of experiments - one with data augmentation (which includes random flips and crops, zero-padded by 4 pixels from each side) and one without it. The data-augmented experiments are marked in the tables with a "+" sign.

Training We apply a similar training method to the one used in [15]: we used stochastic gradient descent with mini-batches of 128 samples, momentum of 0.9, and weight decay of 0.0005. The initial learning rates of FashionMNIST and CIFAR10 were 0.01 and 0.05, respectively. We decrease the learning rate by a factor of 10 when there are 10 epochs without training accuracy improvement. We repeat this process twice during the training period. The training stops if either we reach 300 epochs, or the test loss stops improving for 30 epochs (the latter is to prevent overfitting).

The DecisioNet training process is identical and has the additional hyperparameter β for the loss function (eq. 9). The value of β varies between datasets and models. The specific value that was used for every experiment appears next to the results (table 2).

We initialize the weights of the convolutional layers from a zero-centered Gaussian distribution with a standard deviation of 0.05.

Clustering Interpretability Aside from the performance of the different models, we show that the clustering method partitions the data in an intuitive way. For example, the CIFAR10 first partition is into 2 clusters containing animals and vehicles. The vehicle cluster is then separated into land vehicles (car, truck) and non-land vehicles (plane, ship). In FashionMNIST the first division is into footwear (sandal, snicker, ankle-boot) and non-footwear. The latter is then partitioned into legwear (a singleton cluster with only pants) and non-legwear (T-shirt, dress, bag, etc.). The full hierarchical clustering (for 2 decision levels) is displayed in table 1.
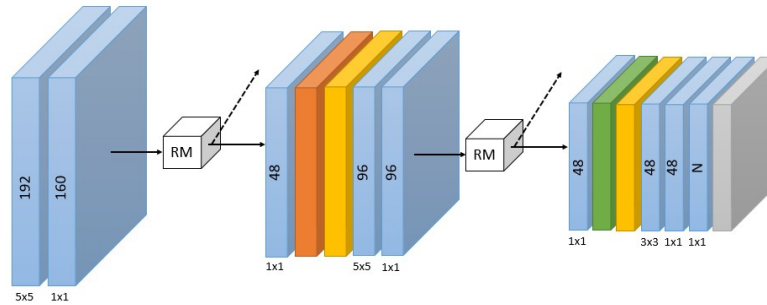
Results and Comparison The results, containing the test accuracy, number of parameters, and number of MACs are presented in table 2. For each experiment
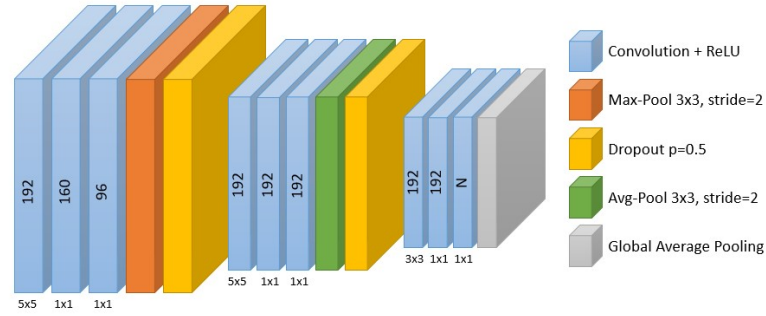
(a)
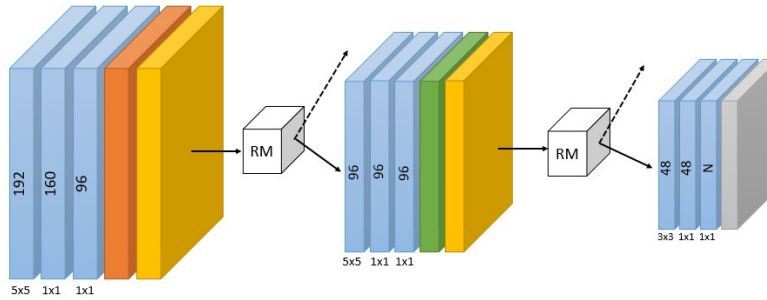
(b)

(c)

Fig. 4. In (a) we display the Network-in-Network (NiN) baseline model architecture. For convolutional layer blocks, the kernel size is stated below each block, and the number of filters is written on the side. We always add zero-padding such that the output and input spatial dimensions are the same (only pooling layers reduce the spatial dimensions). N is the number of output classes (in our case $N = 10$). In (b) and (c) we display a computational path of 2 DecisioNet variants of the NiN model - DN2 and DN2-slim, respectively. The routing module (RM) is the one that is described in sec. 3.3. Notice that the first block after the RM has convolutional layers with half the amount of filters compared to the baseline model; After the second RM, the corresponding block contains a quarter of the filters in the convolutional layers.

Table 1. Hierarchical clustering results (better viewed in color).

Class	1st level	2nd level
plane	0	0
ship	0	0
car	0	1
truck	0	1
bird	1	0
cat	1	0
deer	1	0
dog	1	0
frog	1	0
horse	1	1

Class	1st level	2nd level
T-shirt	0	0
Pullover	0	0
Dress	0	0
Coat	0	0
Shirt	0	0
Bag	0	0
Trouser	0	1
Sneaker	1	0
Ankle-boot	1	0
Sandal	1	1

(a) CIFAR10 clustering   **(b)** FashionMNIST clustering

(row in the table) we measured the average accuracy based on 10 runs with an identical setup (except for the initial weights, of course). For the DN variants, the β values of eq. 9 that were used also appear in the table. The parameters and MAC count were calculated *for a single input image forward pass* using torchinfo [24]. A "+" stands for data augmentation. The DN architectures are described in sec. 4.1.

Analysis: The results (table 2) show that our method works well for these datasets. The performance of the different DN variants is at par with the baseline NiN model (and sometimes even slightly outperforms it). Depending on the exact architecture, we save up to roughly 60% of the computational cost and memory with a negligible decrease in the model's performance.

Comparing DN models whose routing module is placed earlier in the net (DN1-early and DN2-slim) with their similar alternatives (DN1-late and DN2, respectively), we can see a cost-performance tradeoff: when we put the routing module earlier, the model's performance decreases, along with its cost. This phenomenon is another testimony to the well common belief that shallow layers extract features that are important for all classes.

We also would like to emphasize the significance of choosing an appropriate β: when setting $\beta = 0$, we essentially encourage the DN to only optimize the classification loss, ignoring the routing loss. We found that the performance in this case, is similar to a random choice (*i.e.*, DN outputs a constant prediction). This is surprising because in some cases the DN ignored the routing labels at some point, while still improving its classification accuracy.

4.2 CIFAR100

The CIFAR100 dataset [14] is like the CIFAR10, except that it has 100 classes containing 600 images each (500 training images and 100 test images per class).
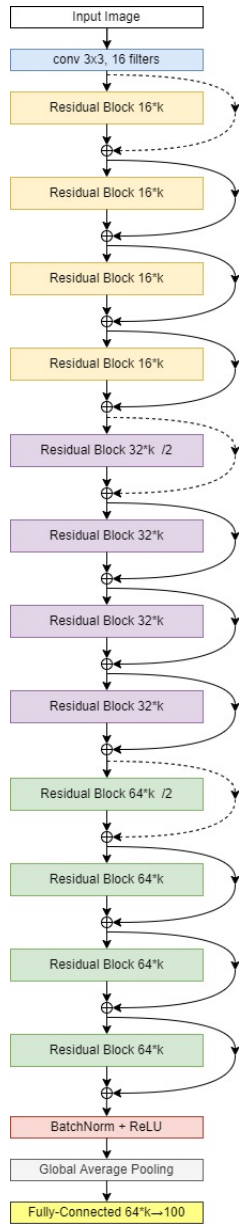
Table 2. Accuracy results, along with parameters and MACs count. For experiments with augmented data (marked with "+") we ommit the parameters and MACs count, as they are identical to those of the non-augmented experiments. More details in sec. 4.1.

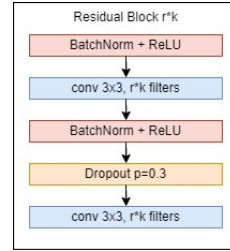| Dataset	Architecture	β	Acc. (%)	Acc. Change (%)	Params (K)	Params change (%)	MACs (M)	MACs change (%)
FashionMNIST	Baseline	-	92.7	-	957.386	-	163.3	
	DN1-early	3.0	92.9	↑ 0.2	736.309	↓23.1	93.64	↓42.7
	DN1-late	3.0	93.4	↑ 0.7	939.157	↓1.9	153.76	↓5.8
	DN2	3.0	92.8	↑ 0.1	727.307	↓24.0	91.24	↓44.1
	DN2-slim	3.0	92.7	0.0	414.027	↓56.8	60.68	↓62.8
CIFAR10	Baseline	-	87.0	-	966.986	-	223.12	-
	DN1-early	3.0	86.6	↓ 0.4	745.909	↓22.9	132.14	↓40.8
	DN1-late	3.0	88.2	↑1.2	948.757	↓1.9	210.66	↓5.6
	DN2	0.5	86.8	↓ 0.2	736.907	↓23.8	129.00	↓42.2
	DN2-slim	1.0	86.2	↓ 0.8	423.627	↓56.2	89.08	↓60.1
CIFAR10+	Baseline	-	88.4	-	-	-	-	-
	DN1-early	3.0	87.8	↓ 0.6	-	-	-	-
	DN1-late	3.0	89.4	↑ 1.0	-	-	-	-
	DN2	0.5	87.6	↓ 0.8	-	-	-	-
	DN2-slim	1.0	86.3	↓ 2.1	-	-	-	-

For this dataset, we experimented with Wide ResNet (WRN) [25] as our baseline network. We used the version with 28 layers and a depth factor of $k = 10$ (denoted as WRN-28-10 in the original paper). We created 3 DN variants out of this network: DN2, whose splits are in between the residual blocks (2 splits), along with DN1-early and DN1-late - each of them with a single split point, positioned at the first or last DN2 split point position. A diagram of the baseline model and the DN2 variant are displayed in figure 5, and the results of our experiments are in table 3.

Training We train the network using the same method as in [25]. That is, we use SGD with Nesterov momentum of 0.9 and a weight decay of 0.0005. The initial learning rate is set to 0.1 and is decreased by a factor of 5 when we reach 60, 120 and 160 epochs. The training is stopped after 200 epochs. The batch size used for training is 128. We augmented the dataset with the same method as we did with CIFAR10, except padding crops with reflections instead of zeros (that is done to follow the training procedure from the original paper).
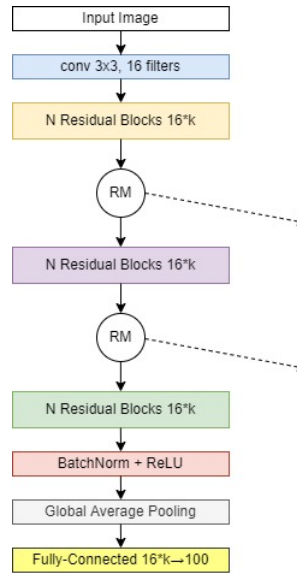
Analysis: we see similar behavior to the former experiments. The top1 accuracy drop is larger than in the former experiments. One possible explanation to this

(b) Residual block architecture. In some cases, the first convolution has a stride of size 2 (In (a) these blocks has the additional "/2" text)

(c) WRN-DN2 single computational branch (we group the residual blocks from (a) into a single block for clarity; Generally, $N = (n - 4)/6$, so in our case $N = 4$). RM is the routing module.

(a) Baseline WRN-n-k architecture with $n = 28$. The dashed arrows are 1x1 convolutions for dimensions matching

Fig. 5. Wide-ResNet-n-k architectures - in this paper we only use $n = 28$ and $k = 10$.

Table 3. CIFAR100 results

Architecture	β	Top-1 Acc. (%)	Top-5 Acc. (%)	Params (M)	MACs (G)
WRN Baseline	-	80.7	95.4	36.537	5.24
WRN-DN1-early	3.0	77.6 ($\downarrow$ 3.1)	93.9($\downarrow$ 1.5)	19.385($\downarrow$ 46.9%)	2.6 ($\downarrow$ 50.4%)
WRN-DN1-late	3.0	78.7 ($\downarrow$ 2.0)	94.5($\downarrow$ 0.9)	23.635 ($\downarrow$ 35.3%)	3.94 ($\downarrow$ 24.8%)
WRN-DN2	5.0	75.7 ($\downarrow$ 5.0)	92.7($\downarrow$ 2.7)	12.935 ($\downarrow$ 64.6%)	2.28 ($\downarrow$ 56.5%)

is the fact that the parameter ratio between the DNs and the baseline model is significantly higher. Different routing module positioning might yield better performance (at some computational cost).

5   Conclusion

We introduced DecisioNet (DN), a binary-tree structured network with conditional routing. We proposed a systematic way of building it based on an existing baseline network. DN utilizes Improved Semantic Hashing for training end-to-end, applying conditional routing both during training and evaluation. DN takes advantage of its tree structure, passing inputs only through a portion of the net's neural modules saving a lot of the computational cost. We evaluated multiple DN variants along with their baseline models on multiple image classification datasets and showed that DN is capable of achieving similar performance (w.r.t. its baseline model) while significantly decreasing the computational cost.

Future Research One clear drawback of our method is the use of additional labels. In addition, these labels are generated using a pre-trained model. Learning hard routing without such auxiliary labels is challenging (as we saw when we set $\beta = 0$), and it might be an interesting direction for future research. Another direction is creating DN with unbalanced trees: in this paper, we only examined balanced DNs, while our clustering trees were not necessarily balanced. A possible idea is to divide the computational power of branches proportionally to the size of the data they should process (*e.g.*, a dataset with 100 classes which is split into clusters of size 70 and 30, would have one branch with 70% of the computational power and the other would have 30%).

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), https://www.tensorflow.org/, software available from tensorflow.org 6
2. Chen, J., Zhu, Z., Li, C., Zhao, Y.: Self-adaptive network pruning (2019) 6
3. Chen, Z., Li, Y., Bengio, S., Si, S.: You look twice: Gaternet for dynamic filter selection in cnns (2019) 6
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009) 2
5. Frosst, N., Hinton, G.: Distilling a neural network into a soft decision tree. arXiv preprint arXiv:1711.09784 (2017) 2
6. Hazimeh, H., Ponomareva, N., Mol, P., Tan, Z., Mazumder, R.: The tree ensemble layer: Differentiability meets conditional computation. In: International Conference on Machine Learning. pp. 4138–4148. PMLR (2020) 2, 3
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition (2015) 4
8. Hehn, T.M., Hamprecht, F.A.: End-to-end learning of deterministic decision trees. In: German conference on pattern recognition. pp. 612–627. Springer (2018) 2
9. Ioannou, Y., Robertson, D., Zikic, D., Kontschieder, P., Shotton, J., Brown, M., Criminisi, A.: Decision forests, convolutional networks and the models in-between. arXiv preprint arXiv:1603.01250 (2016) 3
10. Jordan, M.I., Jacobs, R.A.: Hierarchical mixtures of experts and the em algorithm. Neural computation **6**(2), 181–214 (1994) 2
11. Łukasz Kaiser, Bengio, S.: Discrete autoencoders for sequence models (2018) 3, 6
12. Łukasz Kaiser, Sutskever, I.: Neural gpus learn algorithms (2016) 7
13. Kontschieder, P., Fiterau, M., Criminisi, A., Bulo, S.R.: Deep neural decision forests. In: Proceedings of the IEEE international conference on computer vision. pp. 1467–1475 (2015) 2
14. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) 2, 8, 11
15. Lin, M., Chen, Q., Yan, S.: Network in network. arXiv preprint arXiv:1312.4400 (2013) 8, 9
16. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019), http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf 2, 6
17. Rota Bulo, S., Kontschieder, P.: Neural decision forests for semantic image labelling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 81–88 (2014) 2

18. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. nature **323**(6088), 533–536 (1986) 2
19. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2015) 4
20. Suárez, A., Lutsko, J.F.: Globally optimal fuzzy decision trees for classification and regression. IEEE Transactions on Pattern Analysis and Machine Intelligence **21**(12), 1297–1311 (1999) 2
21. Tanno, R., Arulkumaran, K., Alexander, D., Criminisi, A., Nori, A.: Adaptive neural trees. In: International Conference on Machine Learning. pp. 6166–6175. PMLR (2019) 3
22. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for bench-marking machine learning algorithms. arXiv preprint arXiv:1708.07747 (2017) 2, 8
23. Yan, Z., Zhang, H., Piramuthu, R., Jagadeesh, V., DeCoste, D., Di, W., Yu, Y.: Hd-cnn: Hierarchical deep convolutional neural network for large scale visual recognition (2015) 3, 5
24. Yep, T.: torchinfo (3 2020), https://github.com/TylerYep/torchinfo 11
25. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016) 12