

# Distributed Neural Networks for Internet of Things: The Big-Little Approach

Elias De Coninck<sup>(✉)</sup>, Tim Verbelen, Bert Vankeirsbilck, Steven Bohez,  
Pieter Simoens, Piet Demeester, and Bart Dhoedt

Department of Information Technology, Ghent University – iMinds,  
Gaston Crommenlaan 8/201, 9050 Gent, Belgium  
`elias.deconinck@intec.ugent.be`

**Abstract.** Nowadays deep neural networks are widely used to accurately classify input data. An interesting application area is the Internet of Things (IoT), where a massive amount of sensor data has to be classified. The processing power of the cloud is attractive, however the variable latency imposes a major drawback in situations where near real-time classification is required. In order to exploit the apparent trade-off between utilizing the stable but limited embedded computing power of IoT devices and the seemingly unlimited computing power of Cloud computing at the cost of higher and variable latency, we propose a Big-Little architecture for deep neural networks. A small neural network trained to a subset of prioritized output classes is running on the embedded device, while a more specific classification is calculated when required by a large neural network in the cloud. We show the applicability of this concept in the IoT domain by evaluating our approach for state of the art neural network classification problems on popular embedded devices such as the Raspberry Pi and Intel Edison.

**Keywords:** Deep neural networks · Distributed intelligence · Internet of things

## 1 Introduction

Currently, the Internet of Things (IoT) is a popular paradigm that envisions a world in which all kinds of physical objects or “things” are connected to the Internet, interact with each other and cooperate to reach common goals [1]. This idea goes beyond machine-to-machine communications (M2M), as it covers not only communication protocols, but also the application domains and the services running on top of these connected things. By providing easy access to a myriad of devices such as sensors, surveillance cameras, home appliances, cars, actuators etc., the IoT will enable a new range of applications and use cases in the field of domotics, assisted living, logistics, smart environments, manufacturing, and many more.

In order to create truly smart applications for the IoT, massive amounts of data coming from all connected things will have to be processed and analysed into actionable and contextualized information [8]. Currently, the Cloud is most often the natural choice to perform this data processing, benefiting from the huge computing power and scalability [15]. However, the Cloud is not a silver bullet solution, as a network connection to a Cloud datacenter often suffers from high and variable latency, as well as limited upload bandwidth [2]. Moreover, extensive Cloud processing incurs considerable cost for renting infrastructure. Therefore, our goal is to first use local computing power in the various embedded devices and gateways for data processing, before turning to the Cloud. This addresses the problem associated with network connectivity (latency and bandwidth) while also reducing operation cost.

A very important processing step in IoT applications is classification, i.e., determining the system “state” and its subsequent actions based on (sequences of) sensory data. A promising state-of-the-art technique in this field is the use of a Deep Neural Network (DNN), which offers a biologically inspired trainable architecture that can learn various invariant features [18]. As a neural network is trained by iteratively evaluating input samples and updating the neural network weights, one often relies on efficient GPU implementations that can exploit the parallelism and speedup neural network evaluation. However, the CPU power and internal memory available on an embedded IoT device imposes an upper limit on the size of neural networks that can be evaluated with sufficiently high throughput.

The contributions presented in the current paper is a Big-Little architecture for neural networks, that is tailored to the specific characteristics of IoT environments. The idea is to take a large trained neural network for a certain classification problem, and from that distil a smaller neural network that only classifies a well chosen subset of the output space. The little neural network is suited to be executed locally on the embedded devices, whether the input is also sent to the Cloud, for evaluation with the big neural network, is upto the application specific demands.

The rationale for this idea is the following. Consider a smart home environment that is monitored by a large number of sensors and some actuators to trigger an alarm, start the heating system, detect fire or leaks, etc. Some situations are critical to detect fast while others need a reliable response before taking a direct action. For example stop the robot before bumping into a domestic cat that popped up before it, or close off the water supply when a leak is detected. One can clearly distinguish between critical situations, in which you want a fast or a reliable response, versus non-critical ones, as well as the granularity of the response (a person versus which specific person). By carefully selecting the critical outputs and the granularity of the outputs, we can craft a little local neural network that offers fast response. This local neural network also acts as a filter to limit the number of inputs sent to the big neural network in the cloud. For example, only when a person is detected and one requires a more fine grained classification of the person, the input is evaluated in the Cloud.

The remainder of this paper is organized as follows. Section 2 presents related work in scope of distributed neural networks. Section 3 explains the proposed Big-Little architecture for neural networks in the scope of IoT. In Sect. 4 we show some preliminary results that validate our idea using a frequently used neural network evaluation dataset. Section 5 summarizes our conclusions and presents plans for future work.

## 2 Related Work

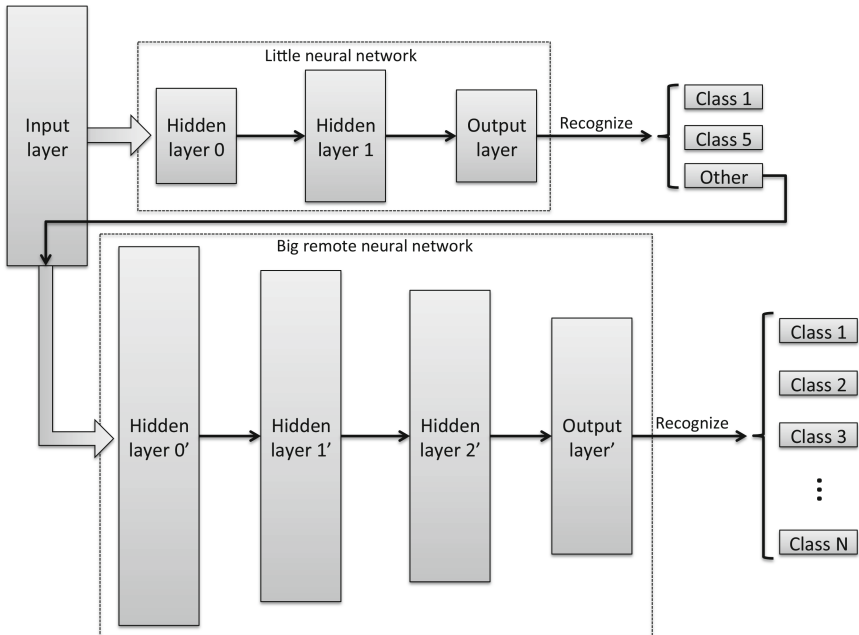
In a highly distributed IoT environment one could try to speed up neural network evaluation by distributing parts of the neural network among all the available devices. However, related work in distributing neural networks shows that the communication overhead quickly becomes the limiting factor [7], hence this approach is mainly used to speed up the training phase on a cluster of nodes connected through a high speed network. Krizhevsky et al. [12] showed how a larger neural network can be trained by spreading the net across two GPUs. This way the communication overhead remains limited, as the GPUs are able to read and write to each other's memory directly. In [10], the authors show that scaling up further to 8 GPUs can lead to a speed up factor of 6.16. Dean et al. [9] presented the DistBelief framework for parallel distributed training of deep neural networks. By adopting new training algorithms they can distribute the training procedure on a large number of CPU nodes, for example achieving a speed up of more than 12x using 81 machines.

As these methods all focus on the training phase and require high end server infrastructure, these are not applicable for speeding up small neural networks in an IoT use case. One approach to optimize distributed neural network execution with parts distributed across embedded devices is simplifying multi-class classification problems to one-vs-all (OVA) or one-vs-one classifiers (OVO). These methods are straightforward and often used to construct a multi-class classifier using binary-class classification [16]. One multi-class classification is split into a set of binary-class classifications for each class and later combine them to the original multi-class classifier.

Another optimization approach in which one part is deployed on an embedded device, and a second part is running in the Cloud, uses a cascade of neural network layers as depicted in [14]. The presented solution exists in augmenting the structure of the neural network to obtain intermediate evaluation output. Then, the evaluation with the remaining neural network layers is pre-empted if the quality of the intermediate output exceeds a given threshold. By deploying only first layers of the net on an embedded device, calls to the Cloud can be limited to the input samples that do not yet result in a good output after these first layers. Our paper proposes a similar approach, but instead of training intermediate layers that classify all outputs with a lower accuracy, we introduce a smaller neural network that is trained to only a subset of the outputs.

### 3 Big-Little Neural Network Architecture

A typical feed-forward neural network is composed of an input layer, one or more hidden layers and one output layer. The output layer has one output for each classification class, resulting in a value between 0 and 1, depicting the probability that the input can be classified as such. In order to process as much as possible locally on embedded IoT devices, and limiting communication to the Cloud, we designed a little neural network that can be processed with limited CPU power, which classifies only a subset of the output classes. This results in a Big-Little neural network architecture as depicted in Fig. 1.



**Fig. 1.** Architecture of Big-Little feed-forward neural network: In this example, the little neural network only classifies a subset (2) of the  $N$  output classes, and can be executed locally with limited CPU power. When the little neural network cannot classify the input sample, a big neural network running in the Cloud can be queried.

The hypothesis of this paper is that by limiting the number of output classes, we can also limit the size and amount of hidden layers of the neural network, while maintaining the desired classification accuracy. By crafting the little neural network to classify high priority classes only, a local response is obtained very fast in these critical cases, while retaining the availability of the complete classification set by relating to the cloud, that executes the big neural network. The application developer may decide to directly send the raw input data to the large neural network in the cloud, or only send the input after evaluation by the little

network. Directly sending it to the cloud allows for a better classification but consumes more bandwidth, while waiting for completion of the little network takes some time.

Because the big and little neural network only share the input layer, these two networks can be trained independently from each other using state-of-the-art training techniques in an offline training phase.

## 4 Evaluation

To train and execute our neural networks, we used Theano [3,4], a Python module compatible with GPU/CPU and many computer architectures. In our experiments we use multilayer perceptrons (MLP) with one or more hidden layers to easily increase the number of calculations by increasing the number of neurons in each hidden layer.

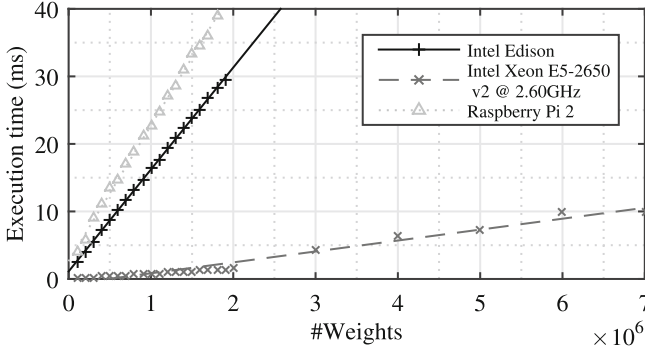
### 4.1 Neural Network Evaluation Time

To assess how the evaluation throughput of a neural network varies with the size of the network, we created fully connected multilayer perceptron (MLP)s with an increasing number of hidden layers. Each hidden layer introduces additional weighted links and thus increase the number of calculations needed to evaluate a single input sample. We then measured the wall clock time needed to evaluate a randomly generated input sample on various types of hardware, as listed in Table 1.

**Table 1.** Hardware specifications.

Name	Architecture	CPU	RAM
Raspberry Pi 2	ARM	Cortex-A7 (quad-core @ 900 MHz)	1 GB
Intel Edison	x86	Intel Atom (dual-core @ 500 MHz)	1 GB
iLab.t server [5]	x86	2x Intel Xeon E5-2650v2 (8-core @ 2.60 GHz)	48 GB

Figure 2 shows how the execution time linearly scales with the number of weights in the neural network. In Fig. 2, we also included experimental results on server-grade hardware from our testbed iLab.t [5] (Table 1). Given a maximum response time of 40 ms, on a Raspberry Pi 2 we can run neural networks of sizes below  $2 \times 10^6$  weights. The size of the large network in the cloud depends on the link delay between the nodes. With a link delay of for example 15 ms, the large neural network evaluated on our server hardware can spend at most  $40ms - 2 * 15ms = 10ms$  on neural network processing, which corresponds with up to  $6.5 \times 10^6$  weights.



**Fig. 2.** Evaluation time of a single input sample for fully connected neural networks of different sizes. Extra cores do not improve performance.

## 4.2 Classification Accuracy vs. Network Size


In our second set of experiments, we want to measure the inevitable loss of accuracy when evaluating a little one-vs-all neural network compared to a big multi-class neural network. By increasing the number of weights in the hidden layer of both networks we can evaluate the difference in reliability. To fairly compare these networks we use the  $F_1$ -score of one priority class:

$$F_1\text{-score} = \frac{2 * \text{true positives}}{2 * \text{true positives} + \text{false positives} + \text{false negatives}}$$

This  $F_1$ -score is the harmonic mean of precision and recall and is independent of the total number of samples. Using the error rate as measure of comparison would yield misleading results because the one-vs-all classifier was trained with an unbalanced set, containing fewer samples for the ‘one’ class than for the ‘all’ class. The  $F_1$ -score of a multi-class neural network can be calculated for each class, but here we only look at the priority class of the little network.

To validate our neural networks, the popular MNIST dataset [13] was used. This dataset consists of a training, validation and test set of handwritten digits with a total of 70000 examples evenly distributed over the classification classes. Starting with the state of the art neural network of Ciresan et al. [6] as the big network that is executed on the cloud, we have distilled a small network by prioritizing some output classes and aggregating the others into one category, thus limiting the number of output classes. For this latter situation, a neural network structure was created that achieves the same accuracy for the prioritized output classes with a much lower neural network size.

Figure 3 shows a comparison between a big and a little fully connected neural network. A more detailed output is given in Tables 2 and 3. Neural networks with more output layers need more weights to reach the same accuracy. Our one-vs-all network outperforms the big neural network for a low number of neural network connections and more or less reaches the same quality of larger networks.



**Fig. 3.**  $F_1$ -score of big and little fully connected neural networks for MNIST classification. The graph shows the  $F_1$ -scores for classifying a ‘5’ character on the big network (ID 3 - Table 2) and the little network (ID 1 - Table 2).

**Table 2.** Execution time of big and little neural networks on the server from Table 1. Each number in the architecture represents the number of neurons in that layer, starting with an input layer of  $28 * 28$  neurons for MNIST. The last layer represents a binary or multi-class classifier layer.

ID	Architecture (number of neurons in each layer)	Weights	Process time for one sample [ms]
1	784, 1 000, 500, 2	1 250 502	0.98
2	784, 1 000, 500, 10	1 254 510	1.00
3	784, 2 500, 2 000, 1 500, 1 000, 500, 10 [6]	11 972 510	16.42

**Table 3.**  $F_1$ -score comparison of a big (ID 3) and a little (ID 1) neural network for two classification classes (1 and 8) of MNIST. Structure of these networks are shown in Table 2. The test error is not applicable for little networks with a unbalanced validation sample set.

ID	Priority class	Test error for best validation [%]	Priority class [%]			
			Recall	Specificity	Precision	$F_1$ -score
1	1	NA	98.94	99.92	99.38	99.16
	8	NA	96.41	99.76	97.71	97.05
3	1	2.11	99.12	99.90	99.21	99.16
	8	2.11	96.82	99.70	97.22	97.02

Deploying the little network (ID 1 from Table 2) with 1 250 502 weights on a Raspberry Pi 2 roughly gives an execution time of 30 ms (from Fig. 2). This little network trained for classifying the priority class ‘8’ has a  $F_1$ -score of 97.05 % which is 0.03 % better than the remote big neural network (ID 3). The overall

performance of this little network is better for this specific class. Other classes yield similar results.

Forwarding the same sample input to the big network gives an execution time of 16.42 ms (from Table 2). From the moment the response time is more than 46.42 ms (sum of execution time of big and little neural network) we can get a faster and equivalent response from the little neural network executed on a Raspberry Pi 2. In most cases this will not make a big difference but in highly critical situations a fast response makes all the difference.

## 5 Conclusion and Future Work

In this paper we investigated the potential of Big-Little neural network architectures, to reduce response time of the overall network, maintaining a comparable accuracy. The results show that prioritizing one class can lower the calculations required to reach the same classification performance for that class. This can benefit use cases where reliability and response time are preferred.

Important points for future work are to deduce little networks and their weights from a pretrained big network and to test the same hypothesis for convolutional neural networks on Cifar [11] and/or ImageNet [17]. In the future we will distribute neural networks on multiple IoT devices in an environment to decrease the cloud usage and increase the independence of this environment.

**Acknowledgements.** Part of the work was supported by the iMinds IoT research program. Steven Bohez is funded by Ph.D. grant of the Agency for Innovation by Science and Technology in Flanders (IWT).

## References

1. Atzori, L., Iera, A., Morabito, G.: The internet of things: a survey. *Comput. Netw.* **54**(15), 2787–2805 (2010)
2. Barker, S.K., Shenoy, P.: Empirical evaluation of latency-sensitive application performance in the cloud. In: *Proceedings of the First Annual ACM SIGMM Conference on Multimedia Systems, MMSys 2010*, pp. 35–46. ACM, New York (2010)
3. Bastien, F., Lamblin, P., Pascanu, R., Bergstra, J., Goodfellow, I.J., Bergeron, A., Bouchard, N., Bengio, Y.: Theano: new features and speed improvements. In: *Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop* (2012)
4. Bergstra, J., Breuleux, O., Bastien, F., Lamblin, P., Pascanu, R., Desjardins, G., Turian, J., Warde-Farley, D., Bengio, Y.: Theano: a CPU and GPU math expression compiler. In: *Proceedings of the Python for Scientific Computing Conference (SciPy)*, vol. 4, p. 3, June 2010. Oral Presentation
5. Bouckaert, S., Becue, P., Vermeulen, B., Jooris, B., Moerman, I., Demeester, P.: Federating wired and wireless test facilities through emulab and OMF: the iLab.t use case. In: *Proceedings of the 8th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities*, pp. 1–16. Department of Information technology, Ghent University (2012)



6. Dan C.C., Meier, U., Gambardella, L.M., Schmidhuber, J.: Deep big simple neural nets excel on handwritten digit recognition. arXiv preprint [arXiv:1003.0358](https://arxiv.org/abs/1003.0358) 22(12) (2010)
7. Coates, A., Huval, B., Wang, T., Wu, D., Catanzaro, B., Andrew, N.: Deep learning with COTS HPC systems. In: Proceedings of the 30th International Conference on Machine Learning, pp. 1337–1345 (2013)
8. Da Li, X., He, W., Li, S.: Internet of things in industries: a survey. *IEEE Trans. Ind. Inform.* **10**(4), 2233–2243 (2014)
9. Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., Le, Q.V., et al.: Large scale distributed deep networks. In: Advances in Neural Information Processing Systems, pp. 1223–1231 (2012)
10. Krizhevsky, A.: One weird trick for parallelizing convolutional neural networks. arXiv preprint [arXiv:1404.5997](https://arxiv.org/abs/1404.5997) (2014)
11. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Technical report, Computer Science Department, University of Toronto (2009)
12. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)
13. LeCun, Y., Cortes, C.: The MNIST database of handwritten digits (1998)
14. Leroux, S., Bohez, S., Verbelen, T., Vankeirsbilck, B., Simoens, P., Dhoedt, B.: Resource-constrained classification using a cascade of neural network layers. In: International Joint Conference on Neural Networks (2015)
15. Parwekar, P.: From internet of things towards cloud of things. In: 2011 2nd International Conference on Computer and Communication Technology (ICCCT), pp. 329–333, September 2011
16. Rifkin, R., Klautau, A.: In defense of one-vs-all classification. *J. Mach. Learn. Res.* **5**, 101–141 (2004)
17. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**(3), 211–252 (2015). doi:[10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y)
18. Schmidhuber, J.: Deep learning in neural networks: an overview. *Neural Netw.* **61**, 85–117 (2015)