Learning Effective Neural Networks on Resource-Constrained Devices

Franco Maria Nardini
ISTI-CNR, Pisa, Italy
francomaria.nardini@isti.cnr.it

Lorenzo Valerio
IIT-CNR, Pisa, Italy
lorenzo.valerio@iit.cnr.it

Raffaele Perego
ISTI-CNR, Pisa, Italy
raffaele.perego@isti.cnr.it

Andrea Passarella
IIT-CNR, Pisa, Italy
andrea.passarella@iit.cnr.it

Abstract—The digital transformation we are experiencing in recent years is cross-cutting to all sectors of the society. In the industrial scenario, this transformation is leading towards the fourth industrial revolution characterized by i) large amounts of data collected and ii) decentralization of computational resources along the production line. In this context the use of artificial intelligence (AI) is often subordinated to the adoption of distributed solutions characterized by the use of limited capacity hardware. In this paper we describe a new framework for learning neural networks on devices with limited resources. A first experimentation on MNIST datasets confirms the validity of the approach that allows to effectively reduce the size of the network during training without significant losses of its accuracy.

Index Terms—Neural networks, resource-constrained devices

I. INTRODUCTION

Nowadays, global companies are facing an important transition towards the Industry 4.0 paradigm. In this paradigm a central role in the development and the optimization of the production process is given to data management and its analysis in order to extract useful knowledge for improving the industrial process. This transformation will be made possible by i) the increasing availability of data produced at all the levels of the production process, together with ii) the presence of more powerful robots (i.e. with improved sensing and computational computational) distributed over the entire production line. As an example, consider two typical case studies in Industry 4.0: i) process optimization, to improve the efficiency of a production chain adapting the behaviour of the robots according to newly obtained knowledge, ii) predictive maintenance, to optimize the lifetime of an industrial plant using the data of the plant itself. However, given the exponential growth of data produced and collected at the edge of the network and specifically within the industrial premises, it is expected that the capacity of the communication and computing infrastructures (despite the advent of 5G and the evolution of data centers) will not deal with the "data tsunami" expected in the coming years, making the paradigm based solely on cloud-based solutions (i.e. Cloud Robotics [1]) no longer sustainable in the medium-long run [2]. Given these considerations, there is an increasing need for combining AI mechanisms based on centralized cloud systems with decentralized AI solutions that are physically located directly on, or very close to, the devices (e.g. industrial robots) that are directly involved in the production processes. For example, a viable solution would be the adoption of paradigms like edge or fog computing [3].

According to this paradigms in which part of the knowledge extraction is being shifted directly on (or close to) the end devices (i.e. robots), becomes fundamental to develop AI techniques that are suitable for the limited resources at their disposal. In fact, such devices are often characterized by very limited resources in terms of computing capacity, memory, and network, which makes it difficult to apply the typical resource-hungry AI techniques used build models with millions of parameters. Therefore, our goal is to develop AI techniques suitable for the Industry 4.0 scenario that robots can use to train efficiently local accurate models, i.e., using the data they own or sense.

In this paper, we focus machine learning solutions suitable for saving resources during learning phase. To this end, we propose to exploit compression techniques for deep neural networks during the learning process so to allow a significant reduction of the size of the neural network while training it. Our purpose is twofold, on the one hand we want to reduce the computational burden to executing the trained final model (i.e. a smaller neural network is computationally less expensive that a bigger one) and on the other hand we want also to reduce the amount of resources used at training time, i.e., memory. Note that here we focus on the memory because of it is well known that, in terms of power consumption, is accessing the memory is in the range of one order of magnitude more expensive that performing arithmetic operations [].

Compression techniques (pruning) of neural networks follow two main research lines. On the one hand, pruning can be achieved by removing some of the neurons of the NN, while on the other, by removing unnecessary weights. Most of the work so far focus mainly on the second approach. The main contributions in this line are based on: i) the factorization of the weights matrix, ii) limiting the magnitude and/or the representation of the values of the weights and iii) using *distillation* methods. The main criticism moved to the works belonging to this group is that the use of compression is typically employed after training the network.

A recent and alternative research line proposes to integrate compression approaches during the learning phase. Methodologically, these approaches try to learn (for example) the distribution of the neurons' activation during the training. Such information are used to turn on (or off) those neurons that are

deemed as more (or less) useful. In this paper we follow this kind of approach.

II. Hard Pruning of Neural Networks

In this work we take inspiration from a state-of the art approach where the network pruning takes place by learning the activation/deactivation distribution of the network neurons [4]. The proposed approach trains a Deep Neural Network minimizing a regularised loss function. The loss function is composed by two terms that are a function of the parameters of the network $\boldsymbol{\theta}$: i) the accuracy of the network and ii) a L_0-norm regularizer $\|\boldsymbol{\theta}\|_0$ through which we seek the minimum set of parameters, i.e., neurons, that maximizes the accuracy of the network. In details,

$$R(\boldsymbol{\theta}) = \frac{1}{N}\left(\sum_{i=1}^{N} L(h(x_i; \boldsymbol{\theta}), y_i)\right) + \lambda\|\boldsymbol{\theta}\|_0 \qquad (1)$$

where, h represents a neural network whose parameters are identified by $\boldsymbol{\theta}$. Moreover, N is the size of the dataset composed by independent and identically distributed pairs $\{(x_1, y_1), (x_2, y_2), \ldots, (x_N, y_N)\}$. Equation 1 is minimized in the number of parameters $\boldsymbol{\theta}$.

Since the learning of the network is done via *backpropagation*, it is not possible to directly use the loss function in Equation 1 due to the non-derivability of the L_0 regularizer. This problem has been overcome through a probabilistic approach, i.e., by approximating the distribution of active/deactive neurons through a continuous and derivable distribution, namely the *Hard Concrete Distribution* [4]. As a result, it is possible to jointly minimize the prediction error and the number of active neurons, epoch by epoch. However, the reduction introduced by this pruning solution is "virtual" and commonly named *soft pruning*. In fact, although the number of active neurons decreases, due to the random extraction process, the active neurons are always different epoch by epoch. From the point of view of the use of resources, this does not guarantee any memory savings since the size of the network is constant, i.e., neurons are simply turned on and off epoch by epoch.

The problem above can be solved by introducing *hard pruning*. The idea is that, during learning, if a neuron is turned off it cannot be turned on again. Training with hard pruning has the effect of clearing, at training time, some of the weights and consequently parts of the network. Differently from soft pruning, this mechanism allows an effective reduction of the memory occupation of the network. In fact, to effectively hard prune the network, we remove the rows of the weights matrices whose neurons are zeroed to obtain matrices with a lower number of elements. However, since with hard pruning neurons a turn off definitively during training, we must be pay attention neurons to not interfere too much the training process in order to not degrade the prediction capabilities of the network. To this end adopt the following probabilistic approach: during each epoch we train the network using soft-pruning and in the meanwhile we compute the empirical distribution of the neurons' activations. In this way we obtain, for each neuron, its

probability of "being active". From the probabilities obtained, we create a binary mask that is applied to the network during the following epochs. This mask, updated from epoch to epoch, helps to shut down the nodes of the network that were previously switched on while maintaining the overall state of hard pruning during training.

III. Preliminary Results and Conclusions

We evaluated the effectiveness of our hard pruning mechanism described in Section II on a multi-label classification task: the recognition of a handwritten number from an image. The task asks to assign to each image of the dataset a class that identifies the handwritten number in the image itself. To this regard, we used the well-known MNIST dataset [5]. We trained a Multilayer Perceptron (MLP) for 500 epochs with two hidden layers of 300 and 100 nodes, respectively. The network consists of 266.610 parameters for a total of about 1 MiB of memory used. The results of the network trained with hard pruning are then compared with those of the MLP trained using soft pruning.

The preliminary results obtained show that the application of hard pruning reduces the size of the network during training with a performance degradation of accuracy of less than 0.2% compared to the equivalent network trained using soft-pruning.

As future work we intend to refine the proposed method and to conduct a more comprehensive evaluation on its performance.

References

[1] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg, "A survey of research on cloud robotics and automation," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 2, pp. 398–409, 2015.

[2] I. Cisco Systems, "Cisco global cloud index: Forecast and methodology, 2016-2021," *White Pap.*, p. 46, 2018.

[3] M. Conti, A. Passarella, and S. K. Das, "The internet of people (iop): A new wave in pervasive mobile computing," *Pervasive Mob. Comput.*, vol. 41, pp. 1–27, 2017.

[4] C. Louizos, M. Welling, and D. P. Kingma, "Learning sparse neural networks through l_0 regularization," *arXiv preprint arXiv:1712.01312*, 2017.

[5] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner *et al.*, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.