Edge Intelligence: On-Demand Deep Learning Model Co-Inference with Device-Edge Synergy

En Li

lien5@mail2.sysu.edu.cn
School of Data and Computer
Science, Sun Yat-sen University

Zhi Zhou

zhouzhi9@mail.sysu.edu.cn
School of Data and Computer
Science, Sun Yat-sen University

Xu Chen*

chenxu35@mail.sysu.edu.cn
School of Data and Computer
Science, Sun Yat-sen University

ABSTRACT

As the backbone technology of machine learning, deep neural networks (DNNs) have quickly ascended to the spotlight. Running DNNs on resource-constrained mobile devices is, however, by no means trivial, since it incurs high performance and energy overhead. While offloading DNNs to the cloud for execution suffers unpredictable performance, due to the uncontrolled long wide-area network latency. To address these challenges, in this paper, we propose **Edgent**, a collaborative and on-demand DNN co-inference framework with device-edge synergy. **Edgent** pursues two design knobs: (1) DNN partitioning that adaptively partitions DNN computation between device and edge, in order to leverage hybrid computation resources in proximity for real-time DNN inference. (2) DNN right-sizing that accelerates DNN inference through early-exit at a proper intermediate DNN layer to further reduce the computation latency. The prototype implementation and extensive evaluations based on Raspberry Pi demonstrate **Edgent**'s effectiveness in enabling on-demand low-latency edge intelligence.

CCS CONCEPTS

• **Computer systems organization** → **Real-time system architecture**; • **Computing methodologies** → *Neural networks*; • **Networks** → Network experimentation;

KEYWORDS

Edge Intelligence, Edge Computing, Deep Learning, Computation Offloading

ACM Reference Format:

En Li, Zhi Zhou, and Xu Chen. 2018. Edge Intelligence: On-Demand Deep Learning Model Co-Inference with Device-Edge Synergy. In *MECOMM'18: ACM SIGCOMM 2018 Workshop on Mobile Edge Communications*, August 20, 2018, Budapest,

*The corresponding author is Xu Chen.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MECOMM'18, August 20, 2018, Budapest, Hungary

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5906-1/18/08...\$15.00

<https://doi.org/10.1145/3229556.3229562>

Hungary. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3229556.3229562>

1 INTRODUCTION & RELATED WORK

As the backbone technology supporting modern intelligent mobile applications, Deep Neural Networks (DNNs) represent the most commonly adopted machine learning technique and have become increasingly popular. Due to DNNs's ability to perform highly accurate and reliable inference tasks, they have witnessed successful applications in a broad spectrum of domains from computer vision [14] to speech recognition [12] and natural language processing [16]. However, as DNN-based applications typically require tremendous amount of computation, they cannot be well supported by today's mobile devices with reasonable latency and energy consumption.

In response to the excessive resource demand of DNNs, the traditional wisdom resorts to the powerful cloud datacenter for training and evaluating DNNs. Input data generated from mobile devices is sent to the cloud for processing, and then results are sent back to the mobile devices after the inference. However, with such a cloud-centric approach, large amounts of data (e.g., images and videos) are uploaded to the remote cloud via a long wide-area network data transmission, resulting in high end-to-end latency and energy consumption of the mobile devices. To alleviate the latency and energy bottlenecks of cloud-centric approach, a better solution is to exploiting the emerging edge computing paradigm. Specifically, by pushing the cloud capabilities from the network core to the network edges (e.g., base stations and WiFi access points) in close proximity to devices, edge computing enables low-latency and energy-efficient DNN inference.

While recognizing the benefits of edge-based DNN inference, our empirical study reveals that the performance of edge-based DNN inference is highly sensitive to the available bandwidth between the edge server and the mobile device. Specifically, as the bandwidth drops from 1Mbps to 50Kbps, the latency of edge-based DNN inference climbs from 0.123s to 2.317s and becomes on par with the latency of local processing on the device. Then, considering the vulnerable and volatile network bandwidth in realistic environments (e.g., due to user mobility and bandwidth contention among various Apps), a natural question is that can we further improve the performance (i.e., latency) of edge-based DNN execution, especially for some mission-critical applications such as VR/AR games and robotics [13].

To answer the above question in the positive, in this paper we proposed **Edgent**, a deep learning model co-inference framework with device-edge synergy. Towards low-latency edge intelligence¹, **Edgent** pursues two design knobs. The first is DNN partitioning, which adaptively partitions DNN computation between mobile devices and the edge server based on the available bandwidth, and thus to take advantage of the processing power of the edge server while reducing data transfer delay. However, worth noting is that the latency after DNN partition is still restrained by the rest part running on the device side. Therefore, **Edgent** further combines DNN partition with DNN right-sizing which accelerates DNN inference through early-exit at an intermediate DNN layer. Needless to say, early-exit naturally gives rise to the latency-accuracy tradeoff (i.e., early-exit harms the accuracy of the inference). To address this challenge, **Edgent** jointly optimizes the DNN partitioning and right-sizing in an on-demand manner. That is, for mission-critical applications that typically have a pre-defined deadline, **Edgent** maximizes the accuracy without violating the deadline. The prototype implementation and extensive evaluations based on Raspberry Pi demonstrate **Edgent**'s effectiveness in enabling on-demand low-latency edge intelligence.

While the topic of edge intelligence has began to garner much attention recently, our study is different from and complementary to existing pilot efforts. On one hand, for fast and low power DNN inference at the mobile device side, various approaches as exemplified by DNN compression and DNN architecture optimization has been proposed [3–5, 7, 9]. Different from these works, we take a scale-out approach to unleash the benefits of collaborative edge intelligence between the edge and mobile devices, and thus to mitigate the performance and energy bottlenecks of the end devices. On the other hand, though the idea of DNN partition among cloud and end device is not new [6], realistic measurements show that the DNN partition is not enough to satisfy the stringent timeliness requirements of mission-critical applications. Therefore, we further apply the approach of DNN right-sizing to speed up DNN inference.

2 BACKGROUND & MOTIVATION

In this section, we first give a primer on DNN, then analyse the inefficiency of edge- or device-based DNN execution, and finally illustrate the benefits of DNN partitioning and right-sizing with device-edge synergy towards low-latency edge intelligence.

2.1 A Primer on DNN

DNN represents the core machine learning technique for a broad spectrum of intelligent applications spanning computer vision, automatic speech recognition and natural language processing. As illustrated in Fig. 1, computer vision applications use DNNs to extract features from an input image and

Figure 1: A 4-layer DNN for computer vision

classify the image into one of the pre-defined categories. A typical DNN model is organized in a directed graph which includes a series of inner-connected layers, and within each layer comprising some number of nodes. Each node is a neuron that applies certain operations to its input and generates an output. The input layer of nodes is set by raw data while the output layer determines the category of the data. The process of passing forward from the input layer to the out layer is called model inference. For a typical DNN containing tens of layers and hundreds of nodes per layer, the number of parameters can easily reach the scale of millions. Thus, DNN inference is *computational intensive*. Note that in this paper we focus on DNN inference, since DNN training is generally delay tolerant and is typically conducted in an off-line manner using powerful cloud resources.

2.2 Inefficiency of Device- or Edge-based DNN Inference

Figure 2: AlexNet Runtime

Currently, the status quo of mobile DNN inference is either direct execution on the mobile devices or offloading to the cloud/edge server for execution. Unfortunately, both approaches may suffer from poor performance (i.e., end-to-end latency), being hard to well satisfy real-time intelligent mobile applications (e.g., AR/VR mobile gaming and intelligent robots) [2]. As illustration, we take a Raspberry Pi tiny computer and a desktop PC to emulate the mobile device and

¹As an initial investigation, in this paper we focus on the execution latency issue. We will also consider the energy efficiency issue in a future study.

edge server respectively, running the classical AlexNet [1] DNN model for image recognition over the Cifar-10 dataset [8]. Fig. 2 plots the breakdown of the end-to-end latency of different approaches under varying bandwidth between the edge and mobile device. It clearly shows that it takes more than 2s to execute the model on the resource-limited Raspberry Pi. Moreover, the performance of edge-based execution approach is dominated by the input data transmission time (the edge server computation time keeps at $\sim 10\text{ms}$) and thus highly *sensitive* to the available bandwidth. Specifically, as the available bandwidth jumps from 1Mbps to 50Kbps, the end-to-end latency climbs from 0.123s to 2.317s. Considering the network bandwidth resource scarcity in practice (e.g., due to network resource contention among users and apps) and computing resource limitation on mobile devices, both of the device- and edge-based approaches are challenging to well support many emerging real-time intelligent mobile applications with stringent latency requirement.

2.3 Enabling Edge Intelligence with DNN Partitioning and Right-Sizing

Figure 3: AlexNet Layer Runtime on Raspberry Pi

DNN Partitioning: For a better understanding of the performance bottleneck of DNN execution, we further break the runtime (on Raspberry Pi) and output data size of each layer in Fig. 3. Interestingly, we can see that the runtime and output data size of different layers exhibit great heterogeneities, and layers with a long runtime do not necessarily have a large output data size. Then, an intuitive idea is DNN partitioning, i.e., partitioning the DNN into two parts and offloading the computational intensive one to the server at a low transmission overhead, and thus to reduce the end-to-end latency. For illustration, we choose the second local response normalization layer (i.e., `lrn_2`) in Fig. 3 as the partition point and offload the layers before the partition point to the edge server while running the rest layers on the device. By DNN partitioning between device and edge, we are able to collaborate hybrid computation resources in proximity for low-latency DNN inference.

DNN Right-Sizing: While DNN partitioning greatly reduces the latency by bending the computing power of the

edge server and mobile device, we should note that the optimal DNN partitioning is still constrained by the run time of layers running on the device. For further reduction of latency, the approach of *DNN Right-Sizing* can be combined with DNN partitioning. DNN right-sizing promises to accelerate model inference through an early-exit mechanism. That is, by training a DNN model with multiple exit points and each has a different size, we can choose a DNN with a small size tailored to the application demand, meanwhile to alleviate the computing burden at the model division, and thus to reduce the total latency. Fig. 4 illustrates a branchy AlexNet with five exit points. Currently, the early-exit mechanism has been supported by the open source framework BranchyNet[15]. Intuitively, DNN right-sizing further reduces the amount of computation required by the DNN inference tasks.

Problem Description: Obviously, DNN right-sizing incurs the problem of latency-accuracy tradeoff — while early-exit reduces the computing time and the device side, it also deteriorates the accuracy of the DNN inference. Considering the fact that some applications (e.g., VR/AR game) have stringent deadline requirement while can tolerate moderate accuracy loss, we hence strike a nice balance between the latency and the accuracy in an on-demand manner. Particularly, given the predefined and stringent latency goal, we maximize the accuracy without violating the deadline requirement. More specifically, the problem to be addressed in this paper can be summarized as: given a predefined latency requirement, how to jointly optimize the decisions of DNN partitioning and right-sizing, in order to maximize DNN inference accuracy.

Figure 4: An illustration of the early-exit mechanism for DNN right-sizing

3 FRAMEWORK

We now outline the initial design of **Edgent**, a framework that automatically and intelligently selects the best partition point and exit point of a DNN model to maximize the accuracy while satisfying the requirement on the execution latency.

3.1 System Overview

Fig. 5 shows the overview of **Edgent**. **Edgent** consists of three stages: offline training stage, online optimization stage and co-inference stage.

At **offline training stage**, Edgent performs two initializations: (1) profiling the mobile device and the edge server

Figure 5: Edgent overview

Table 1: The independent variables of regression models

Layer Type	Independent Variable
Convolution	amount of input feature maps, (filter size/stride) ² *(num of filters)
Relu	input data size
Pooling	input data size, output data size
Local Response Normalization	input data size
Dropout	input data size
Fully-Connected	input data size, output data size
Model Loading	model size

to generate regression-based performance prediction models (Sec. 3.2) for different types DNN layer (e.g., Convolution, Pooling, etc.). (2) Using Branchynet to train DNN models with various exit points, and thus to enable early-exit. Note that the performance profiling is infrastructure-dependent, while the DNN training is application-dependent. Thus, given the sets of infrastructures (i.e., mobile devices and edge servers) and applications, the two initializations only need to be done once in an offline manner.

At **online optimization stage**, the DNN optimizer selects the best partition point and early-exit point of DNNs to maximize the accuracy while providing performance guarantee on the end-to-end latency, based on the input: (1) the profiled layer latency prediction models and Branchynet-trained DNN models with various sizes. (2) the observed available bandwidth between the mobile device and edge server. (3) The pre-defined latency requirement. The optimization algorithm is detailed in Sec. 3.3.

At **co-inference stage**, according to the partition and early-exit plan, the edge server will execute the layer before the partition point and the rest will run on the mobile device.

3.2 Layer Latency Prediction

When estimating the runtime of a DNN, Edgent models the per-layer latency rather than modeling at the granularity of a

whole DNN. This greatly reduces the profiling overhead since there are very limited classes of layers. By experiments, we observe that the latency of different layers is determined by various independent variables (e.g., input data size, output data size) which are summarized in Table 1. Note that we also observe that the DNN model loading time also has an obvious impact on the overall runtime. Therefore, we further take the DNN model size as a input parameter to predict the model loading time. Based on the above inputs of each layer, we establish a regression model to predict the latency of each layer based on its profiles. The final regression models of some typical layers are shown in Table 2 (size is in bytes and latency is in ms).

3.3 Joint Optimization on DNN Partition and DNN Right-Sizing

At online optimization stage, the DNN optimizer receives the latency requirement from the mobile device, and then searches for the optimal exit point and partition point of the trained branchynet model. The whole process is given in Algorithm 1. For a branchy model with M exit points, we denote that the i -th exit point has N_i layers. Here a larger index i correspond to a more accurate inference model of larger size. We use the above-mentioned regression models to predict ED_j the runtime of the j -th layer when it runs on device and ES_j the runtime of the j -th layer when it runs on server. D_p is the output of the p -th layer. Under a specific bandwidth B , with the input data $Input$, then we calculate $A_{i,p}$ the whole runtime ($\sum_{j=1}^{p-1} ES_j + \sum_{k=p}^{N_i} ED_k + Input/B + D_{p-1}/B$) when the p -th is the partition point of the selected model of i -th exit point. When $p = 1$, the model will only run on the device then $ES_p = 0$, $D_{p-1}/B = 0$, $Input/B = 0$ and when $p = N_i$, the model will only run on the server then $ED_p = 0$, $D_{p-1}/B = 0$. In this way, we can find out the best partition point having the smallest latency for the model of i -th exit point. Since the model partition does not affect the inference accuracy, we can then sequentially try the DNN inference models with

Table 2: Regression model of each type layer

Layer	Mobile Device model	Edge Server model
Convolution	$y = 6.03e-5 * x_1 + 1.24e-4 * x_2 + 1.89e-1$	$y = 6.13e-3 * x_1 + 2.67e-2 * x_2 - 9.909$
Relu	$y = 5.6e-6 * x + 5.69e-2$	$y = 1.5e-5 * x + 4.88e-1$
Pooling	$y = 1.63e-5 * x_1 + 4.07e-6 * x_2 + 2.11e-1$	$y = 1.33e-4 * x_1 + 3.31e-5 * x_2 + 1.657$
Local Response Normalization	$y = 6.59e-5 * x + 7.80e-2$	$y = 5.19e-4 * x + 5.89e-1$
Dropout	$y = 5.23e-6 * x + 4.64e-3$	$y = 2.34e-6 * x + 0.0525$
Fully-Connected	$y = 1.07e-4 * x_1 - 1.83e-4 * x_2 + 0.164$	$y = 9.18e-4 * x_1 + 3.99e-3 * x_2 + 1.169$
Model Loading	$y = 1.33e-6 * x + 2.182$	$y = 4.49e-6 * x + 842.136$

different exit points(i.e., with different accuracy), and find the one having the largest size and meanwhile satisfying the latency requirement. Note that since regression models for layer latency prediction are trained beforehand, Algorithm 1 mainly involves linear search operations and can be done very fast (no more than 1ms in our experiments).

Algorithm 1 Exit Point and Partition Point Search

Input:

M : number of exit points in a branchy model
 $\{N_i | i = 1, \dots, M\}$: number of layers in each exit point
 $\{L_j | j = 1, \dots, N_i\}$: layers in the i -th exit point
 $\{D_j | j = 1, \dots, N_i\}$: each layer output data size of i -th exit point
 $f(L_j)$: regression models of layer runtime prediction
 B : current wireless network uplink bandwidth
Input: the input data of the model
latency: the target latency of user requirement

Output:

Selection of exit point and its partition point

```

1: Procedure
2: for  $i = M, \dots, 1$  do
3:   Select the  $i$ -th exit point
4:   for  $j = 1, \dots, N_i$  do
5:      $ES_j \leftarrow f_{edge}(L_j)$ 
6:      $ED_j \leftarrow f_{device}(L_j)$ 
7:   end for
8:    $A_{i,p} = \underset{p=1, \dots, N_i}{\operatorname{argmin}} (\sum_{j=1}^{p-1} ES_j + \sum_{k=p}^{N_i} ED_k + \text{Input}/B + D_{p-1}/B)$ 
9:   if  $A_{i,p} \leq \text{latency}$  then
10:    return Selection of Exit Point  $i$  and its Partition Point  $p$ 
11:   end if
12: end for
13: return NULL ▷ can not meet latency requirement

```

4 EVALUATION

We now present our preliminary implementation and evaluation results.

4.1 Prototype

We have implemented a simple prototype of **Edgent** to verify the feasibility and efficacy of our idea. To this end, we take a desktop PC to emulate the edge server, which is equipped with a quad-core Intel processor at 3.4 GHz with 8 GB of RAM, and runs the Ubuntu system. We further use Raspberry Pi 3 tiny computer to act as a mobile device. The Raspberry

Pi 3 has a quad-core ARM processor at 1.2 GHz with 1 GB of RAM. The available bandwidth between the edge server and the mobile device is controlled by the WonderShaper [10] tool. As for the deep learning framework, we choose Chainer [11] that can well support branchy DNN structures.

For the branchynet model, based on the standard AlexNet model, we train a branchy AlexNet for image recognition over the large-scale Cifar-10 dataset [8]. The branchy AlexNet has five exit points as showed in Fig. 4 (Sec. 2), each exit point corresponds to a sub-model of the branchy AlexNet. Note that in Fig. 4, we only draw the convolution layers and the fully-connected layers but ignore other layers for ease of illustration. For the five sub-models, the number of layers they each have is 12, 16, 19, 20 and 22, respectively.

For the regression-based latency prediction models for each layer, the independent variables are shown in the Table. 1, and the obtained regression models are shown in Table 2.

4.2 Results

We deploy the branchynet model on the edge server and the mobile device to evaluate the performance of **Edgent**. Specifically, since both the pre-defined latency requirement and the available bandwidth play vital roles in **Edgent**'s optimization logic, we evaluate the performance of **Edgent** under various latency requirements and available bandwidth.

We first investigate the effect of the bandwidth by fixing the latency requirement at 1000ms and varying the bandwidth 50kbps to 1.5Mbps. Fig. 6(a) shows the best partition point and exit point under different bandwidth. While the best partition points may fluctuate, we can see that the best exit point gets higher as the bandwidth increases. Meaning that the higher bandwidth leads to higher accuracy. Fig. 6(b) shows that as the bandwidth increases, the model runtime first drops substantially and then ascends suddenly. However, this is reasonable since the accuracy gets better while the latency is still within the latency requirement when increase the bandwidth from 1.2Mbps to 2Mbps. It also shows that our proposed regression-based latency approach can well estimate the actual DNN model runtime latency. We further fix the bandwidth at 500kbps and vary the latency from 100ms to 1000ms. Fig. 6(c) shows the best partition point and exit point under different latency requirements. As expected, the best exit point gets higher as the latency requirement increases, meaning that a larger latency goal gives more room for accuracy improvement.

Figure 6: The results under different bandwidths and different latency requirements

Figure 7: The accuracy comparison under different latency requirement

In Fig. 7, under different latency requirements, it shows the model accuracy of different inference methods. The accuracy is negative if the inference can not satisfy the latency requirement. The network bandwidth is set to 400kbps. Seen in the Fig. 7, at a very low latency requirement (100ms), all four methods can't satisfy the requirement. As the latency requirement increases, inference by **Edgent** works earlier than the other methods that at the 200ms to 300ms requirements, by using a small model with a moderate inference accuracy to meet the requirements. The accuracy of the model selected by **Edgent** gets higher as the latency requirement relaxes.

5 CONCLUSION

In this work, we present **Edgent**, a collaborative and on-demand DNN co-inference framework with device-edge synergy. Towards low-latency edge intelligence, **Edgent** introduces two design knobs to tune the latency of a DNN model: DNN partitioning which enables the collaboration between edge and mobile device, and DNN right-sizing which shapes the computation requirement of a DNN. Preliminary implementation and evaluations on Raspberry Pi demonstrate the effectiveness of **Edgent** in enabling low-latency edge intelligence. Going forward, we hope more discussion and effort can be stimulated in the community to fully accomplish the vision of intelligent edge service.

ACKNOWLEDGE

This work was supported in part by the National Key Research and Development Program of China under grant No.2017YFB1001703; the National Science Foundation of China (No.U1711265); the Fundamental Research Funds for the Central Universities (No.17lgjc40), and the Program for Guangdong Introducing Innovative and Entrepreneurial Teams (No.2017ZT07X355).

REFERENCES

- [1] K. Alex, I. Sutskever, and G. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*.
- [2] China Telecom et al. 2017. 5G Mobile/Multi-Access Edge Computing. (2017).
- [3] Forrest N. Iandola et al. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size. *arXiv:1602.07360* (2016).
- [4] J. Wu et al. 2016. Quantized Convolutional Neural Networks for Mobile Devices. (2016).
- [5] S. Han, H. Mao, and W. Dally. 2015. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. *Fiber* (2015).
- [6] Y. Kang, J. Hauswald, C. Gao, A. Rovinski, T. Mudge, J. Mars, and L. Tang. 2017. Neurosurgeon: Collaborative Intelligence Between the Cloud and Mobile Edge. In *2017 ASPLOS*.
- [7] Y. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin. 2016. Compression of Deep Convolutional Neural Networks for Fast and Low Power Mobile Applications. In *ICLR*.
- [8] Alex Krizhevsky. 2009. Learning Multiple Layers of Features from Tiny Images. (2009).
- [9] N. Lane, S. Bhattacharya, A. Mathur, C. Forlivesi, and F. Kawsar. 2016. DXTK: Enabling Resource-efficient Deep Learning on Mobile and Embedded Devices with the DeepX Toolkit.
- [10] V. Mulhollon. 2004. WonderShaper. <http://manpages.ubuntu.com/manpages/trusty/man8/wondershaper.8.html>. (2004).
- [11] Preferred Networks. 2017. Chainer. <https://github.com/chainer/chainer/tree/v1>. (2017).
- [12] Aaron Van Den et al. Oord. 2016. WaveNet: A Generative Model for Raw Audio. *arXiv:1609.03499* (2016).
- [13] Qualcomm. 2017. Augmented and Virtual Reality: the First Wave of 5G Killer Apps. (2017).
- [14] Christian et al. Szegedy. 2014. Going deeper with convolutions. (2014).
- [15] S. Teerapittayanon, B. McDanel, and H. T. Kung. 2016. BranchyNet: Fast inference via early exiting from deep neural networks. In *2016 23rd ICPR*.
- [16] Di Wang and Eric Nyberg. 2015. A Long Short-Term Memory Model for Answer Sentence Selection in Question Answering. In *ACL and IJCNLP*.