



FastDeepIoT: Towards Understanding and Optimizing Neural Network Execution Time on Mobile and Embedded Devices

Shuochao Yao
University of Illinois Urbana
Champaign

Yiran Zhao
University of Illinois Urbana
Champaign

Huajie Shao
University of Illinois Urbana
Champaign

ShengZhong Liu
University of Illinois Urbana
Champaign

Dongxin Liu
University of Illinois Urbana
Champaign

Lu Su
State University of New York at
Buffalo

Tarek Abdelzaher
University of Illinois Urbana
Champaign

ABSTRACT

Deep neural networks show great potential as solutions to many sensing application problems, but their excessive resource demand slows down execution time, posing a serious impediment to deployment on low-end devices. To address this challenge, recent literature focused on compressing neural network size to improve performance. We show that changing neural network size does not proportionally affect performance attributes of interest, such as execution time. Rather, extreme run-time nonlinearities exist over the network configuration space. Hence, we propose a novel framework, called FastDeepIoT, that uncovers the non-linear relation between neural network structure and execution time, then exploits that understanding to find network configurations that significantly improve the trade-off between execution time and accuracy on mobile and embedded devices. FastDeepIoT makes two key contributions. First, FastDeepIoT automatically learns an accurate and highly interpretable execution time model for deep neural networks on the target device. This is done without prior knowledge of either the hardware specifications or the detailed implementation of the used deep learning library. Second, FastDeepIoT informs a

compression algorithm how to minimize execution time on the profiled device without impacting accuracy. We evaluate FastDeepIoT using three different sensing-related tasks on two mobile devices: Nexus 5 and Galaxy Nexus. FastDeepIoT further reduces the neural network execution time by 48% to 78% and energy consumption by 37% to 69% compared with the state-of-the-art compression algorithms.

CCS CONCEPTS

• **Human-centered computing** → Ubiquitous and mobile computing; • **Computing methodologies** → Machine learning; • **Computer systems organization** → Embedded and cyber-physical systems;

KEYWORDS

Deep Learning, Execution Time, Model Compression, Mobile Computing, Internet of Things

ACM Reference Format:

Shuochao Yao, Yiran Zhao, Huajie Shao, ShengZhong Liu, Dongxin Liu, Lu Su, and Tarek Abdelzaher. 2018. FastDeepIoT: Towards Understanding and Optimizing Neural Network Execution Time on Mobile and Embedded Devices. In *The 16th ACM Conference on Embedded Networked Sensor Systems (SenSys '18)*, November 4–7, 2018, Shenzhen, China. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3274783.3274840>

1 INTRODUCTION

The proliferation of internetworked mobile and embedded devices with growing sensing and computing capabilities promises to revolutionize the interactions between humans and devices that perform complex sensing and recognition tasks. Much prior work has been dedicated to building

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.

SenSys '18, November 4–7, 2018, Shenzhen, China

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5952-8/18/11...\$15.00

<https://doi.org/10.1145/3274783.3274840>

Table 1: Execution time of convolutional layers with 3×3 kernel size, stride 1, same padding, and 224×224 input image size on the Nexus 5 phone.

	in_channel	out_channel	FLOPs	Time (ms)
CNN1	8	32	452.4 M	114.9
CNN2	32	8	452.4 M	300.2
CNN3	66	32	3732.3 M	908.3
CNN4	43	64	4863.3 M	751.7


smarter and more user-friendly sensing applications in several embedded systems areas, including health and wellness [6, 25, 29, 33], context sensing [7, 21, 23, 31, 39], and object detection and localization [9, 14, 18, 19, 22, 32].

At the same time, recent advances in deep learning have changed the way computing devices process human-centric content, such as images, speech and audio. Neural network models are especially good at fusing multiple sensing modalities and extracting temporal relationships, which have shown remarkable improvements in audio sensing [11, 17], tracking and localization [34], human activity recognition [24, 34, 36], and environment sensing [35]. Applying deep neural networks to mobile and embedded devices could thus bring about a generation of applications capable of performing complex sensing and recognition tasks to support a new realm of interactions between humans and their physical surroundings [37].

The key impediment to wide-spread deployment of deep-learning-based sensing applications remains their high execution time and energy consumption on mobile and embedded devices. Minimizing the execution time of deep neural networks is critical to preserve the real-time properties of such embedded sensing applications as image recognition and object detection in self-driving cars [12, 26]. One promising solution is to compress neural networks into more succinct structures. Traditionally, speeding up neural network execution time is accomplished by reducing the size of model parameters [13, 38]. Most manually designed time-efficient neural network structures for mobile devices use parameter size or FLOPs as the indicator of execution time [15, 16, 40]. Even the official TensorFlow website recommends to use the total number of floating number operations (FLOPs) of neural networks “to make rule-of-thumb estimates of how fast they will run on different devices”.¹

Although significant progress has been made on neural network structure compression to reduce the resource demands, changing neural network structure has a non-linear effect on system performance, opening opportunities for further performance improvements should such nonlinearities be explicitly identified and exploited.

In this paper, we show how a better understanding of the non-linear relation between neural network structure and performance can further improve execution time and energy consumption without impacting accuracy. The rest of this

**Figure 1: The non-linearity of neural network execution time over input/output channel.**

paper is organized as follows. The nonlinear relation between network structure and performance is discussed in Section 2. We present the technical details of FastDeepIoT in Section 3 and system implementation in Section 4. The evaluation is presented in Section 5. Section 6 introduces related work. We conclude in Section 7 introducing avenues for future work.

2 NONLINEARITIES: EVIDENCE AND EXPLOITATION

In practice, counting the number of neural network parameters and the total FLOPs does not lead to good estimates of execution time because the relation between these predictors and execution time is not proportional. On one hand, the fully-connected layer usually has more parameters but takes much less time to run compared to the convolutional layer [27]. On the other hand, one can easily find examples, where increasing the total FLOPs does not translate into added execution time. Caching effects, memory accesses, and compiler optimizations complicate the translation. Table 1 shows that CNN2 takes around $\times 2.6$ the execution time of CNN1, while both have *the same* total FLOPs. Moreover, CNN3 takes *longer* to run compared to CNN4 despite having *fewer* FLOPs. These observations indicate that current rules-of-thumb for estimating neural network execution time are not the best approximations.

FastDeepIoT answers two key questions to better parameterize neural network implementations for efficient execution on mobile and embedded platforms:

- (1) What are the main factors that affect the execution time of neural networks on mobile and embedded devices?
- (2) How to guide existing structure compression algorithms to minimize the neural network execution time properly?

FastDeepIoT consists of two main modules to tackle these two challenging problems, respectively.

Profiling: Due to different code-level optimizations for different network structures within the deep learning library, the execution time of neural network layers can be extremely nonlinear over the structure configuration space. A simple illustration is shown in Figure 1, where we plot the execution time of convolutional layers when changing the size of input and output channels simultaneously. The plot reveals non-monotonic effects, featuring periodic dips in execution time as network size increases.

¹<https://www.tensorflow.org/versions/r1.5/mobile/optimizing>

A simple regression model over the entire space will thus not be a good approximation. Instead, we propose a tree-structured linear regression model. Specifically, we automatically detect key conditions at which linearity is violated and arrange them into a tree structure that splits the overall modeling space into piecewise linear regions. Within each region (tree branch), we use linear regression to convert input structure information into some key explanatory variables, predictive of execution time. The splitting of the overall space and the fitting of subspaces to predictive models are done jointly, which improves both model interpretability and accuracy. The aforementioned modeling is done without specific knowledge of underlying hardware and deep learning library.

Compression: Using the results of profiling, we then propose a compression steering module that guides existing neural network structure compression methods to better minimize execution time. The execution time model leads compression algorithms to focus more on the layer that takes longer to run instead of treating all layers equally or concentrating on inaccurate total metrics. It is also better able at exploiting non-monotonicity of execution time with respect to network structure size to reduce the former without hurting application-level accuracy metrics.

We evaluate the profiling and compression steering modules in FastDeepIoT on two devices, Nexus 5 and Galaxy Nexus, with the TensorFlow for Mobile library [2]. The profiling module is evaluated on all commonly used network layers, including fully-connected, convolutional, and recurrent layers. The mean absolute percentage error in estimating execution time is around 1% to 7%, which outperforms other complex regression models in most cases. The compression steering module is evaluated with three representative sensing-related tasks, including vision-based interactions and human activity recognition. Compared to the state-of-the-art compression algorithms, FastDeepIoT can speed up the neural network execution time by an additional 48% to 78%, and improve energy consumption by an additional 37% to 69% on all devices without loss of accuracy.

3 SYSTEM DESIGN

As mentioned above, the contribution of FastDeepIoT lies in two modules; the profiling module and the compression steering module. Below, we introduce the technical details of the two modules, respectively.

3.1 Profiling Module

We separate this module into two parts. The first part generates diverse training structures for profiling. The second part builds an accurate and interpretable model predicting the execution time of deep learning components for the corresponding structure information.

Table 2: The scope of our structure configuration for fully-connected (FC), convolutional (CNN), and recurrent (RNN) layers.

Type	Structure configuration scope	
FC	$in_dim \in [1, 4096]$	$out_dim \in [1, 4096]$
CNN	$in_height \in [24, 225]$	$in_width \in [24, 225]$
	$kernel_height \times kernel_width \in \{2 \times 2, 3 \times 3, 4 \times 4, 5 \times 5, 2 \times 3\}$	
	$in_channel \in [1, 256]$	$out_channel \in [1, 256]$
	$padding \in \{valid, same\}$	$stride \in \{1, 2\}$
RNN	$in_dim \in [1, 512]$	$out_dim \in [1, 512]$
	$step \in \{8, 10, 15, 20\}$	

3.1.1 Neural Network Profiling. We introduce the basic system settings and the procedure of generating training structures for profiling here.

FastDeepIoT utilizes TensorFlow benchmark tool [1] to profile the execution time of all deep learning components on the target device. In order to make the profiling results fully reflect the changes on the neural network structures, we fix the frequencies of phone CPUs (processors) to be constants and stop all the power management services that can affect the processor frequency on target devices, such as fixing *mpdecision* on Qualcomm chips.

The next step is to generate diverse neural network structures for time profiling. As a deep learning component, such as a convolutional layer and recurrent layer, the combinations of its structure design choices can form an extremely huge structure configuration space. Therefore, we can only select a small proportion of structure configurations during our time profiling. The scope of our structure configuration is shown in Table 2, from which the network generation code chooses a random combination. Notice that we do not contain the activation function as the profiling choice, because it only occupies around 1% ~ 2% execution time of a deep learning component through empirical observations. By eliminating this insignificant configuration, *i.e.*, `activation_function` $\in \{\text{ReLU}, \text{Tanh}, \text{sigmoid}\}$, we can save the number of profiling components by the factor of 3. Except for some pre-defined cases, such as sigmoid activation function for gate outputs in recurrent layers, we set all activation functions to be ReLU, which is one of the most widely used activation functions. In addition, the order of deep learning components in the network has little impact on their execution time empirically.

In our profiling module, for each target device, we profile around 120 neural networks with about 1300 deep learning components in total. These time profiling results form a time profiling dataset, $\mathcal{D} = \{S_i, y_i\}$, where S_i is the structure configuration and y_i the execution time.

3.1.2 Execution Time Model Building. Due to the code-level optimization for different component configuration choices in the deep learning library, execution-time non-linearity appears over the structure configuration space as shown in Figure 1. The main challenge here is to build a model that can automatically figure out the conditions that cause

the execution-time non-linearity without specific knowledge of underlying library and hardware.

In order to maintain both the accuracy and interpretability, we propose a tree-structure linear regression model. The model can recursively partition the structure configuration space such that the time profiling samples fitting the same linear relationship are grouped together. The intuition behind this model is that the execution time of deep learning component under each particular code-level optimization can be formulated with a linear relationship given a set of well-designed explanatory variables. In addition, different deep learning components, *i.e.*, fully connected, convolutional, and recurrent layer, learn their own execution time models.

Each time profiling data is composed of three elements. The feature vector \mathbf{f} , used for identifying the condition that causes the execution-time non-linearity; the execution time y ; and the explanatory variable vector \mathbf{x} , used for fitting the execution time y .

The basic idea of tree-structure linear regression is to find out the most significant condition causing the execution-time non-linearity within the current dataset recursively. These conditions will form a binary tree structure. In order to figure out key conditions causing the execution-time non-linearity, we take two conditioning functions into account.

- (1) *Range condition* $C_1(\mathbf{f}[j], \tau) := f[j] \leq \tau$: identifies execution-time non-linearity caused by cache and memory hit as well as specific implementation for a certain feature range.
- (2) *Integer multiple condition* $C_2(\mathbf{f}[j], \tau) := f[j] \equiv 0 \pmod{\tau}$: identifies execution-time non-linearity caused by loop unrolling, data alignment, and parallelized operations.

Assume that we are generating node m in the binary tree with dataset \mathcal{D}_m . The model creates a set of conditions $\{\phi\}$. Each of them can partition the dataset into two subsets $\mathcal{D}_m^{(l)}(\phi)$ and $\mathcal{D}_m^{(r)}(\phi)$. Each condition ϕ consists of three elements, $\phi = \{f[j], \tau_m, k\}$, where $k \in \{1, 2\}$ is the conditioning function type.

$$\begin{aligned}\mathcal{D}_m^{(l)}(\phi) &= \mathcal{D}_m | C_k(x_j, \tau_m), \\ \mathcal{D}_m^{(r)}(\phi) &= \mathcal{D}_m \setminus \mathcal{D}_m^{(l)}(\phi).\end{aligned}\quad (1)$$

Node m selects the most significant condition ϕ^* by minimizing the impurity function $G(\mathcal{D}_m, \phi)$,

$$\phi^* = \arg \min_{\phi} G(\mathcal{D}_m, \phi), \quad (2)$$

$$G(\mathcal{D}_m, \phi) = \frac{|\mathcal{D}_m^{(l)}(\phi)|}{|\mathcal{D}_m|} H(\mathcal{D}_m^{(l)}(\phi)) + \frac{|\mathcal{D}_m^{(r)}(\phi)|}{|\mathcal{D}_m|} H(\mathcal{D}_m^{(r)}(\phi)), \quad (3)$$

$$H(\mathcal{D}) = \min_{\mathbf{w}, b} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}} (\mathbf{w}^\top \mathbf{x} + b - y)^2 \quad \text{s.t. } \mathbf{w}, b \geq 0. \quad (4)$$

Algorithm 1 Execution time model building

```

1: Input: time profiling dataset  $\mathcal{D}_z$ , feature vector  $\mathbf{f}$ , two conditioning functions
    $f[j] \leq \tau$  and  $f[j] \equiv 0 \pmod{\tau}$ , and explanatory variable  $\mathbf{x}_z$ .
2: Fit  $\mathcal{D}_z$  with  $\mathbf{w}_r$  and  $b_r$  according to (4).
3: Save root node  $r = [\emptyset, \mathbf{w}_r, b_r]$ 
4: Initialize:  $\text{que} = [r, \mathcal{D}_z]$ .
5: while  $\text{len}(\text{que}) > 0$  do
6:    $q, \mathcal{D}_q = \text{que.dequeue}()$ 
7:   if  $\mathcal{D}_q$  meets stoping condition then
8:     Continue
9:   end if
10:  Search for the optimal partition  $\phi^* = \{f[j^*], \tau^*, k^*\}$  according to (2) (3) (4).
11:  Generate partitioned dataset  $\mathcal{D}_q^{(l)}(\phi^*)$  and  $\mathcal{D}_q^{(r)}(\phi^*)$  according to (1).
12:  Fit  $\mathcal{D}_q^{(l)}(\phi^*)$  with  $\mathbf{w}_q^{(l)}$  and  $b_q^{(l)}$  according to (4).
13:  Fit  $\mathcal{D}_q^{(r)}(\phi^*)$  with  $\mathbf{w}_q^{(r)}$  and  $b_q^{(r)}$  according to (4).
14:  Save  $q$ 's left child node  $q^{(l)} = [\text{True}, \phi^*, \mathbf{w}_q^{(l)}, b_q^{(l)}]$ 
15:  Save  $q$ 's right child node  $q^{(r)} = [\text{False}, \phi^*, \mathbf{w}_q^{(r)}, b_q^{(r)}]$ 
16:   $\text{que.enqueue}([q^{(l)}, \mathcal{D}_q^{(l)}(\phi^*)])$ 
17:   $\text{que.enqueue}([q^{(r)}, \mathcal{D}_q^{(r)}(\phi^*)])$ 
18: end while

```

The impurity function is designed as the weighted mean square errors of linear regressions over two sub-datasets partitioned by the condition ϕ .

Next, we describe the feature vector \mathbf{f} . Our choice of feature vector \mathbf{f} contains three parts: the structure features, the memory features, and the parameter feature. The structure features refer to *in_dim* and *out_dim* for fully-connected and recurrent layers as well as *in_channel* and *out_channel* for convolutional layers. The memory features include the memory size of input, *mem_in*, the memory size of output, *mem_out*, and the memory size of internal representations, *mem_inter*. The parameter feature refers to the size of parameters, *param_size*. The detailed definitions of memory and parameter features are shown in Table 3. All notations in Table 3 are consistent with the notations of structure configurations in Table 2, except for the height and width of output image, *out_height* and *out_width*, in the convolutional layer. However, we can easily calculate these two values based on other structure information, *i.e.*, *in_height*, *in_width*, *kernel_height*, *kernel_width*, *stride*, and *padding*².

Last, we discuss about our explanatory variable vector \mathbf{x} for linear regression. In this paper, we build an intuitive performance model that the execution time of a program is contributed by three parts, CPU operations, memory operations, and disk I/O operations. For a deep learning component, these parts refer to FLOPs, memory size, and parameter size,

$$\mathbf{x} = [\text{FLOPs}, \text{mem}, \text{param_size}]. \quad (5)$$

where $\text{mem} = \text{mem_in} + \text{mem_out} + \text{mem_inter}$.

With the weight vector \mathbf{w} and the bias term b , the overall execution time of a deep learning component, y , can be modelled as $y = \mathbf{w}^\top \mathbf{x} + b$. Since every term should have a positive contribution to the execution time, we add an additional constraint, $\mathbf{w}, b \geq 0$, as shown in (4).

²https://www.tensorflow.org/api_guides/python/nn#Convolution

Table 3: The definition of parameter and memory information for Fully-Connected layer (FC), Convolutional layer (CNN), Gated Recurrent Unit (GRU), and Long Short Term Memory (LSTM).

Type	$param_size$	mem_in	mem_out	mem_inter
FC	$in_dim \times out_dim + out_dim$	in_dim	out_dim	0
CNN	$kernel_height \times kernel_width \times in_channel \times out_channel + 1$	$in_height \times in_width \times in_channel$	$out_height \times out_width \times out_channel$	$out_height \times out_width \times kernel_height \times kernel_width \times in_channel$
GRU	$3 \times out_dim \times (in_dim + out_dim + 1)$	$step \times in_dim$	$step \times out_dim$	$3 \times step \times out_dim$
LSTM	$4 \times out_dim \times (in_dim + out_dim + 1)$	$2 \times step \times in_dim$	$2 \times step \times out_dim$	$4 \times step \times out_dim$

Table 4: The p-values of explanatory variables.

Type	FLOPs	mem	param_size	step
FC	0.000	0.000	1.000	
CNN	0.037	0.009	1.000	
GRU	0.000	0.000	1.000	0.000
LSTM	0.000	0.000	1.000	0.000

The tree-structure linear regression model builds a binary tree that gradually picks out conditions that cause execution-time non-linearity and breaks the dataset into subsets that contain more “linearity”. Our designed explanatory variable vector \mathbf{x} is able to fit the dataset with linear relationships better level by level, especially for fully-connected and convolutional layer. The recurrent layers, however, still have flaws. We analyze the error and find out that recurrent layers have a constant initialization overhead or set-up time for each step. Therefore, we update explanatory variable vector \mathbf{x} ,

$$\begin{aligned} \mathbf{x}_{fc} &= \mathbf{x}_{cnn} = [\text{FLOPs}, \text{mem}, \text{param_size}], \\ \mathbf{x}_{rnn} &= [\text{FLOPs}, \text{mem}, \text{param_size}, \text{step}]. \end{aligned} \quad (6)$$

We summarize our execution time model building process in Algorithm 1. There is a stopping condition in Line 7 that keeps tree-structure linear regression from growing infinitely. In our case, the stopping condition occurs when a linear regression can fit the current dataset \mathcal{D}_q with a mean absolute percentage error less than 5% or when the size of current dataset is smaller than 15, $|\mathcal{D}_q| < 15$.

3.1.3 Execution Time Model with Statistical Analysis. In this part, we provide an illustration of the FastDeepIoT profiling module on Nexus 5 phone with statistical analysis. The module first profiles and generates the execution time profiling dataset. Then, the module builds an execution time model for each deep learning component based on the tree-structure linear regression in Algorithm 1. Additional evaluations on the execution time model will be shown in Section 5.1.

For fully-connected layers and recurrent layers, including GRU and LSTM, their execution time has a perfect linear relationship with our explanatory variable vector \mathbf{x}_{fc} and \mathbf{x}_{rnn} . However, the execution time model of convolutional layers reflects a strong non-linearity over the structure configuration space. As shown in Figure 2, the execution time of convolutional layer has local minima when $in_channel$ or $out_channel$ is a multiple of 4.

Then we calculate the p-values to evaluate the mathematical relationship between each explanatory variable and the

execution time. The p-value for each explanatory variable tests the null hypothesis that the variable has no correlation with the execution time. Results are shown in Table 4. The p-values of explanatory variables, FLOPs, mem , and $step$, are less than the significance level (0.05) for all deep learning components. So our empirical time profiling data provides enough evidence that the correlation between these explanatory variables and the execution time are statistically significant. However, the p-values for $param_size$ is high for all cases, which shows that the number of parameters has limited correlation with the execution time. This experiment, again, highlights the importance of proposing a compression algorithm targeting on minimizing the execution time instead the number of parameters.

3.2 Compression Steering Module

Profiling and modelling deep learning execution time is not enough for speeding up the model execution. In this section, we introduce the compression steering module that is designed to empower existing deep learning structure compression algorithms to minimize model execution time properly.

We assume that $\mathcal{S} = \{\mathbf{s}_l\}$ and $\mathcal{W} = \{\mathbf{W}_l\}$ for $l = 1, \dots, L$ is structure information and weight matrix of a neural network from layer 1 to layer L respectively. We denote our execution time model as $t_l = \mathcal{T}(\mathbf{s}_l)$, which takes the structure information \mathbf{s}_l as input and predicts the component execution time t_l . For a general neural network structure compression algorithm, we denote the original compression process as,

$$\min_{\mathcal{S}, \mathcal{W}} \mathcal{L}_\theta(\mathcal{S}, \mathcal{W}), \quad (7)$$

where the compression algorithm minimizes a loss function, concerning prediction error or parameter size, with either the gradient descend or searching based optimization method.

In order to enable the compression algorithm to minimize the execution time, our first step is to incorporate the execution time model into the original objective function (7),

$$\min_{\mathcal{S}, \mathcal{W}} \mathcal{L}_\theta(\mathcal{S}, \mathcal{W}) + \lambda \sum_{l=1}^L \mathcal{T}(\mathbf{s}_l), \quad (8)$$

where λ is a hyper-parameter that make the tradeoff between minimizing training loss and minimizing execution time.

Adding execution time to the compression objective function can encourage the compression algorithm to concentrate




Figure 2: The execution time model of convolutional layers on Nexus 5.

Algorithm 2 Layer expansion and local minima searching

```

1: Input: the execution time model  $\mathcal{T}()$  with root node  $r$  and the layer structure  $\{f, x\}$ .
2: Set node =  $r$ , condL = [].
3: while  $\neg(\text{node.left} == \text{None} \ \& \ \text{node.right} == \text{None})$  do
4:   if node.cond is a range condition then
5:     condL.append(node.cond)
6:     if  $f$  obeys node.cond then
7:       node = node.left
8:     else
9:       node = node.right
10:    end if
11:  else
12:     $\hat{f} = f$  and  $\hat{x} = x$ .
13:     $\hat{f}[\text{node.j}] = \text{node.r} \times \lceil \hat{f}[\text{node.j}] / \text{node.r} \rceil$ 
14:    Update  $\hat{x}$  according to  $\hat{f}$ .
15:    if  $\text{node.w}_T^T \hat{x} + \text{node.b}_T > \text{node.w}_F^T \hat{x} + \text{node.b}_F$  &  $\hat{f}$  obeys condL then
16:       $f = \hat{f}$  and  $x = \hat{x}$ .
17:      node = node.left
18:    else
19:      node = node.right
20:    end if
21:  end if
22: end while
23: Return:  $f$ 

```

more on the layers with higher execution time, which helps to speed up the whole neural network.

However, due to the existence of execution-time local minima, compressing neural network structure is not always the optimal choice for minimizing the execution time. As shown in Figure 1, enlarging neural network structure can find a nearby execution-time local minimum that reduces the execution time. Notice that enlarging structure is a lossless operation. We can at least enlarge weight matrices with zeros that keeps performance the same.

In general, utilizing execution-time local minima for speeding up involves two steps:

- (1) Identifying an expanded structure configuration that can trigger a nearby execution-time local minimum.
- (2) Deciding whether the expanded structure can speed up the execution time.

For an execution time model trained with a complex method, such as neural networks, identifying a nearby execution-time local minimum can be almost impossible by blindly searching a large configuration space. However, our tree-structure linear regression can easily identify a nearby local minimum speeding up the neural network execution.

Local extrema, *i.e.*, maxima and minima, are identified by the integer multiple condition, $f[j] \equiv 0 \pmod{\tau}$, in our tree-structure linear regression model. Our compression steering





Figure 3: Simplified execution time model of convolutional layers on Nexus 5.

Figure 4: The square region of safely expanding $in_channel$ for speed up.

module searches for the nearby local maxima by gradually expanding the structure that fits the integer multiple conditions from root node to leaf node in the execution time model.

Assume that node m is under the condition $f[j_m] \equiv 0 \pmod{\tau_m}$ with two sets of linear regression parameters $\{w_T, b_T\}$ and $\{w_F, b_F\}$ used for fitting the dataset that obeying and against the condition respectively. A deep learning layer is denoted with the feature vector f_l and the explanatory variable vector x_l . The compression steering module generates an expanded layer with feature vector \hat{f}_l and explanatory variable vector \hat{x}_l by updating the conditioning feature $\hat{f}[j_m] = \tau_m \lceil f[j_m] / \tau_m \rceil$. Then the module compares the values between $w_T^T \hat{x} + b_T$ and $w_F^T \hat{x} + b_F$ to decide whether it should accept the expansion for speeding-up and go through the corresponding branch.

The layer structure expansion and local minima searching process is summarized in Algorithm 2. The algorithm goes through whole tree structure to find out a nearby local minimum that reduces the execution time.

For a whole neural network, each layer goes through the structure expansion and local minima searching process one by one. It is possible that conflicts exist between expanded structures of two neighbouring layers. The module solves these conflicts sequentially by choosing the one having shorter overall execution time.

In addition, we can further analyze the structure expansion process for a particular component on a particular device for a particular application settings. For example, assume that we are compressing the $in_channel$ and $out_channel$ of a convolutional layer on Nexus 5 with kernel size 3×3 , input image size 24×24 , and the same padding. We are considering the root condition $in_channel \equiv 0 \pmod{4}$ as shown in Figure 2. According to our execution time model, two linear regression models that fit the two datasets in the left and right child of the root node are:

$$\begin{aligned} w_T &= [3.41 \times 10^{-8}, 4.03 \times 10^{-6}, 7.11 \times 10^{-25}] \quad b_T = 8.11, \\ w_F &= [3.11 \times 10^{-8}, 8.03 \times 10^{-6}, 1.52 \times 10^{-34}] \quad b_F = 12.82. \end{aligned} \quad (9)$$

Then we can obtain the execution time as a function of $in_channel$ and $out_channel$ by substituting the explanatory variable vector \mathbf{x} with definitions illustrated in Table 3 as well as the application settings about kernel size, input image size, and padding option.

$$\begin{aligned} y_T(in_c, out_c) &= 3.53 \times 10^{-4} \cdot in_c \cdot out_c + 8.11 + \\ &\quad 2.32 \times 10^{-2} \cdot in_c + 2.32 \times 10^{-3} \cdot out_c, \\ y_F(in_c, out_c) &= 3.23 \times 10^{-4} \cdot in_c \cdot out_c + 12.82 + \\ &\quad 4.63 \times 10^{-2} \cdot in_c + 4.63 \times 10^{-3} \cdot out_c, \end{aligned} \quad (10)$$

where we denote $in_channel$ and $out_channel$ as in_c and out_c for simplicity.

We are interested in the region where expanding the $in_channel$ to a nearby multiple of 4 can speed up the execution. This is equivalent to solving

$$y_T(in_c + 3, out_c) - y_F(in_c, out_c) < 0, \quad (11)$$

where its zero contour line is a hyperbola. Therefore, within the region bounded by $in_channel$ axis, $out_channel$ axis, and zero contour line, we can safely expand $in_channel$ to a multiple of 4 to speed up the convolutional layer execution time.

In order to have a more interpretable result, as shown in Figure 4, we can obtain a square region by finding the intersections between the zero contour line and the function $out_channel = in_channel$. In this case, within the region $in_channel \times out_channel \in [1, 1288] \times [1, 1288]$, we can blindly expand $in_channel$ to a multiple of 4 to speed up. This region is much larger than the region we are interested in. We can keep analyzing the next condition $out_channel \equiv 0 \pmod{4}$ and achieve similar result. Within the region $in_channel \times out_channel \in [1, 808] \times [1, 808]$, we can safely expand $in_channel$ and $out_channel$ to a nearby multiple of 4 to speed up. In the end, we can obtain a simplified execution time model $\hat{\mathcal{T}}$ as shown in Figure 3.

In summary, the compression steering module compresses the neural network structure for reducing overall execution time with three steps.

- (1) Compressing neural network with a time-aware objective function (8) with execution time model \mathcal{T} .
- (2) Expanding layer structure and searching local minima for further speed up according to Algorithm 2 with execution time model \mathcal{T} or $\hat{\mathcal{T}}$ (if available).
- (3) Depending on the original compression algorithm, freeze the structure and fine-tune the neural network.

4 IMPLEMENTATION

In this section, we briefly describe the hardware, software, and architecture of FastDeepIoT.

4.1 Hardware

In this paper, we test FastDeepIoT on two types of hardware, Nexus 5 phone and Galaxy Nexus phone. Two devices are

profiled for each type of hardware. The Nexus 5 phone is equipped with quad-core 2.3 GHz CPU and 2 GB memory. The Galaxy Nexus phone is equipped with dual-core 1.2 GHz CPU and 1GB memory. We stop the *mpdecision* service and use *userspace* CPU governor for two hardware. We manually set 1.1GHz for the quad-core CPU on Nexus 5, and 700MHz for the dual-core CPU on Galaxy Nexus to prevent overheating caused by the constant time profiling. In addition, all profiling and testing neural network models are run solely on CPU. The execution time model building and the compression steering module are implemented on a workstation connected to two phones.

4.2 Software

FastDeepIoT utilizes TensorFlow benchmark tool [1], a C++ binary, to profile the execution time of deep learning components. For each neural network, the benchmark tool have one warm up run to initialize the model and then profile all components execution time with 20 runs without internal delay. Mean values are taken as the profiled execution time.

We install Android 5.0.1 on Nexus 5 phone and Android 4.3 on Galaxy Nexus phone. All additional background services are closed during the profiling and testing. All energy consumptions are measured by an external power meter.

4.3 Architecture

Given a target device, FastDeepIoT first queries the device and its own database for a pre-generated execution time model with device type and OS version as the key. If the query fails, the profiling module starts its function. FastDeepIoT generates random neural network structures based on the configuration scope in Table 2, pushes the Protocol Buffers (.pb file) to the target device, profiles the execution time of components, fetches back and processes the profiling result. Once the profiling process has finished, FastDeepIoT learns tree-structure linear regression execution time models according to Algorithm 1 based on the time profiling dataset. FastDeepIoT pushes the generated execution time models to the target device and its own database for storage.

Then given an original neural network structure and parameters, the compression steering module can automatically generate a compressed structure to speed up inference time for a target device. FastDeepIoT queries the target device and own database for a pre-generated execution time model, and choose a structure compression algorithm, DeepIoT as a default, to reduce the deep learning execution time according to (8) and Algorithm 2. The resulting compressed neural network is transferred to the target device used locally.

5 EVALUATION

In this section, we evaluate FastDeepIoT through two sets of experiments. The first set evaluates the accuracy of the execution time model generated by our profiling module, while

Table 5: The Mean Absolute Percentage Error (MAPE), Mean Absolute Error (MAE) in millisecond, and Coefficient of determination (R^2) of execution time models.

(a) Nexus 5-Convolutional layer						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	7.6%	233.8%	23.8%	19.7%	10.9%	6.4%
MAE	15.2	227.1	39.2	27.3	20.5	16.4
R^2	0.991	-0.229	0.969	0.985	0.988	0.994
(b) Nexus 5-Gated recurrent unit						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	1.8%	78.7%	9.4%	6.7%	4.8%	2.0%
MAE	0.6	23.6	2.9	1.8	1.5	0.7
R^2	0.999	-0.078	0.986	0.995	0.995	0.999
(c) Nexus 5-Long short term memory						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	2.3%	73.7%	9.0%	4.7%	4.1%	2.8%
MAE	0.6	23.7	3.0	1.4	1.6	0.9
R^2	0.999	-0.223	0.977	0.995	0.993	0.998
(d) Nexus 5-Fully-connected layer						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	1.9%	133.5%	22.5%	12.0%	0.2%	1.9%
MAE	0.19	5.98	1.18	0.38	0.01	0.19
R^2	0.999	0.065	0.977	0.996	0.999	0.999
(e) Galaxy Nexus-Convolutional layer						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	4.1%	164.3%	33.1%	23.0%	15.2%	14.5%
MAE	26.8	878.7	162.9	123.7	114.6	110.1
R^2	0.999	-0.246	0.969	0.980	0.982	0.983
(f) Galaxy Nexus-Gated recurrent unit						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	2.9%	71.5%	10.5%	8.8%	6.0%	4.1%
MAE	1.1	27.8	4.8	4.1	3.2	2.2
R^2	0.997	-0.065	0.968	0.977	0.984	0.989
(g) Galaxy Nexus-Long short term memory						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	2.9%	66.8%	8.4%	7.8%	6.0%	2.9%
MAE	1.4	26.2	3.0	3.3	2.7	1.3
R^2	0.997	-0.196	0.983	0.985	0.987	0.997
(h) Galaxy Nexus-Fully-connected layer						
	FastDeepIoT	SVR	DT	RF	GBRT	DNN
MAPE	4.0%	55.0%	12.3%	11.3%	9.5%	4.1%
MAE	0.3	6.7	1.2	0.9	1.0	0.3
R^2	0.996	-0.629	0.944	0.972	0.949	0.996

the second set evaluates the performance of our compression steering module. In order to evaluate execution time modeling accuracy, we compare our tree-structured linear regression model to other state-of-the-art regression models on two mobile devices. To evaluate the quality of compression, we present a set of experiments that demonstrate the speed-up of the compressed neural network obtained by the compression steering module with three human-centric interaction and sensing applications.

5.1 Execution time Model

We implement the following execution time estimation alternatives:

- (1) **SVR**: support vector regression with radial basis function kernel [8]. This algorithm tries to perform linear separation over a higher dimensional kernel feature space by characterizing the maximal margin.

- (2) **DT**: classification and regression trees [5]. This is an interpretable model. It groups and predicts execution time by the execution time itself.
- (3) **RF**: random forest regression [4]. This algorithm trades the interpretability of regression tree for the predictive performance by ensembling multiple trees with random feature selections.
- (4) **GBRT**: gradient boosted regression trees [10]. This algorithm builds an additive model in a forward stage-wise fashion, which is hard to interpret.
- (5) **DNN**: multilayer perceptron [20]. Deep neural network is a learning model with high capacity. We build a four-layer fully connected neural network with LeRU as the activation function, except for the output layer. We fine-tune the structure and apply dropout as well as L2 regularization to prevent overfitting. DNN is a black-box model.

We train all the baseline models with the dataset generated by the profiling module in FastDeepIoT (75% for training and 25% for testing). For each deep learning component, such as CNN and LSTM, an individual model is trained. We have trained these models with feature vector \mathbf{f} , explanatory variable vector \mathbf{x} , and the concatenate of feature and explanatory variable vectors as inputs, where \mathbf{f} and \mathbf{x} are the same as the definitions in Section 3.1.2. We find that the model trained with explanatory variable vector \mathbf{x} outperforms other choices consistently in all cases, so we only report the results of models trained with \mathbf{x} for simplicity.


We evaluate these models on convolutional layer, gated recurrent unit, long short term memory, and fully-connected layer with mean absolute percentage error, mean absolute error, and coefficient of determination on two hardware. As shown in Table 5, FastDeepIoT is consistently among top 2 predictors for all experiments with all three metrics. FastDeepIoT also outperforms the highly capable deep learning model for more than half of the cases, while FastDeepIoT is much more interpretable. There are two reasons for the remarkable performance of FastDeepIoT. On one hand, FastDeepIoT captures the primary characters of deep learning execution time behaviours, which makes an interpretable and accurate model possible. On the other hand, since the profiled dataset is limited (around one thousand samples for training), complex models such as deep neural networks that require large training dataset may not be the best choice here.

5.2 Compression Steering Module


In this section, we evaluate the performance of our compression steering module with three sensing applications. We train the neural networks on traditional benchmark datasets as original models. Then, we compress the original models using FastDeepIoT and the three state-of-the-art baseline algorithms. Finally, we test the accuracy, execution time, and energy consumption of compressed models on mobile devices.

Table 6: VGGNet (hidden units) on CIFAR-10 dataset.

Layer	No Execution Time Model				Nexus 5	Galaxy Nexus
	Original	DeepIoT	DeepIoT+localMin	DeepIoT+FLOPs	FastDeepIoT	FastDeepIoT
conv1-1 (3×3)	64	27	28	19	12	16
conv1-2 (3×3)	64	47	48	17	16	24
conv2-1 (3×3)	128	53	56	33	28	36
conv2-2 (3×3)	128	68	68	50	32	44
conv3-1 (3×3)	256	104	104	89	64	72
conv3-2 (3×3)	256	97	100	79	64	56
conv3-3 (3×3)	256	89	92	77	68	72
conv4-1 (3×3)	512	122	124	115	132	96
conv4-2 (3×3)	512	95	96	112	136	80
conv4-3 (3×3)	512	64	64	112	104	120
conv5-1 (2×2)	512	128	128	143	148	116
conv5-2 (2×2)	512	112	112	132	144	108
conv5-3 (2×2)	512	146	148	182	104	92
fc1	4096	27	27	1097	132	132
fc2	4096	161	161	935	152	123
fc3	1000	10	96	72	157	167
Test accuracy	90.6%	90.6%	90.6%	90.6%	90.6%	90.6%
Execution time t (Nexus 5)	328 ms	31 ms	21 ms	28 ms	16 ms	23 ms
Execution time t (Galaxy Nexus)	610 ms	72 ms	63 ms	52 ms	36 ms	34 ms



(a) Tradeoff between testing accuracy and execution time on Nexus 5. (b) Tradeoff between testing accuracy and execution time on Galaxy Nexus. (c) Tradeoff between testing accuracy and compressed parameter size.



(d) Tradeoff between testing accuracy and energy consumption on Nexus 5. (e) Tradeoff between testing accuracy and energy consumption on Galaxy Nexus.

Figure 5: System performance tradeoff for VGGNet on CIFAR-10 dataset

We compare FastDeepIoT with three baseline algorithms:

- (1) **DeepIoT**: This is a state-of-the-art neural structure compression algorithm [38]. The algorithm designs a compressor neural network with adaptive dropout to explore a succinct structure for the original model.
- (2) **DeepIoT+localMin**: We enhance DeepIoT with the ability of expanding layer for finding execution-time local minima. This method takes the compressed model of DeepIoT and expands its layers with zero-value elements that can trigger local minima according to Algorithm 2. We use this almost zero-effort method to show the improvement made on existing compressed models by interpreting deep learning execution time with FastDeepIoT.
- (3) **DeepIoT+FLOPs**: This method enhances DeepIoT by adding a term that minimizes FLOPs to the original objective function (7). Since a large proportion of works use FLOPs as the execution time estimation [15, 16, 40], this method shows to what extent FLOPs can be used to compress neural network for reducing execution time.






5.2.1 Image recognition on CIFAR-10. This is a vision based task, image recognition based on a low-resolution camera. During this experiment, we use CIFAR-10 as our training and testing dataset. The CIFAR-10 dataset consists of 60000 32×32 colour images in 10 classes, with 6000 images per class. There are 50000 training images and 10000 test images.

During the evaluation, we use VGGNet structure as the original network structure [28]. The detailed structure is shown in Table 6, where we also illustrate the best compressed models that keeps the original test accuracy for all algorithms. The compressed model can be even deployed on tiny IoT devices such as Intel Edison.

As shown in Table 6, FastDeepIoT achieves the best performance on two hardware with their corresponding execution time models. Compared with the state-of-the-art DeepIoT algorithm, FastDeepIoT can further reduce the model execution time by 48% to 53%. DeepIoT+localMin outperforms DeepIoT on two hardware, reducing the execution time by

Table 7: VGGNet (hidden units) on ImageNet dataset.

Layer	No Execution Time Model				Nexus 5	Galaxy Nexus
	Original	DeepIoT	DeepIoT+localMin	DeepIoT+FLOPs	FastDeepIoT	FastDeepIoT
conv1-1 (3×3)	64	43	44	23	12	16
conv1-2 (3×3)	64	47	48	32	12	16
conv2-1 (3×3)	128	100	100	65	20	44
conv2-2 (3×3)	128	97	100	67	40	40
conv3-1 (3×3)	256	164	164	116	88	108
conv3-2 (3×3)	256	164	164	135	72	104
conv3-3 (3×3)	256	153	156	70	116	108
conv4-1 (3×3)	512	235	236	72	268	240
conv4-2 (3×3)	512	240	240	181	236	216
conv4-3 (3×3)	512	220	240	258	340	200
conv5-1 (3×3)	512	255	256	261	376	240
conv5-2 (3×3)	512	260	260	303	376	288
conv5-3 (3×3)	512	257	260	47	176	216
fc1	4096	436	436	1594	656	920
fc2	4096	1169	1169	824	1150	1189
fc3	1000	297	297	405	287	402
Test top-5 accuracy	88.9%	88.9%	88.9%	88.9%	88.9%	88.9%
Execution time t (Nexus 5)		1682 ms	1605 ms	968.8 ms	688.8 ms	725.7 ms
Execution time t (Galaxy Nexus)		7773 ms	6991 ms	3930 ms	3211 ms	2930 ms

**(a) Tradeoff between testing accuracy and execution time on Nexus 5.****(b) Tradeoff between testing accuracy and execution time on Galaxy Nexus.****(c) Tradeoff between testing accuracy and compressed parameter size.****(d) Tradeoff between testing accuracy and energy consumption on Nexus 5.****(e) Tradeoff between testing accuracy and energy consumption on Galaxy Nexus.****Figure 6: System performance tradeoff for VGGNet on ImageNet dataset**

12% to 32%. This shows that we can decently reduce the neural network execution time by simply expanding the neural network structure to local execution-time minima. In addition, DeepIoT+FLOPs can speed up the model execution time compared with DeepIoT. However, FastDeepIoT still outperforms DeepIoT+FLOPs by a significant margin. This result highlights that FLOPs is not a proper estimation of time.

Figure 5a and 5b shows the tradeoff between testing accuracy and execution time for different algorithms. FastDeepIoT consistently outperforms other algorithms by a significant margin. Furthermore, the execution time characters on different hardware can affect the final performance. FastDeepIoT (Nexus 5/Galaxy Nexus) performs better on its corresponding hardware. DeepIoT+localMin achieves a better tradeoff compared with DeepIoT. Therefore, utilizing execution-time local minima is a low-cost strategy to speed up neural network execution. In addition, since FLOPs has different degrees of execution time contribution on different hardware,


DeepIoT+FLOPs are not able to achieve a better tradeoff than DeepIoT on all devices.

Figure 5d and 5e shows the tradeoff between testing accuracy and energy consumption for different algorithms. Although FastDeepIoT is not designed to minimize the energy consumption, FastDeepIoT still achieves the best tradeoff. However, we can see that the characters of energy consumption of deep neural network are different from the execution time. FastDeepIoT with the hardware-specific time models are not always the most energy-saving method on the corresponding hardware. Execution-time local minima cannot consistently help DeepIoT+localMin to outperform DeepIoT. Therefore, further studies on understanding and minimizing deep learning energy consumption are needed.


Figure 5c shows the tradeoff between testing accuracy and left proportion of model parameters. Since there is no algorithm targeting at minimizing model parameters, all methods show comparable performances. However, from another perspective, the execution time model learnt by FastDeepIoT

Table 8: DeepSense (hidden units) on HHAR dataset.


Layer		No Execution Time Model										Nexus 5		Galaxy Nexus	
		Original		DeepIoT		DeepIoT+localMin		DeepIoT+FLOPs		t_{min}		FastDeepIoT		FastDeepIoT	
conv1a	conv1b (2 × 9)	64	64	20	19	20	20	26	25	4	4	16	8	16	16
conv2a	conv2b (1 × 3)	64	64	20	14	20	16	19	17	4	4	8	12	20	16
conv3a	conv3b (1 × 3)	64	64	23	23	24	24	22	22	4	4	16	16	16	16
conv4 (2 × 8)		64		10		12		9		4		12		16	
conv5 (1 × 6)		64		12		12		13		4		16		16	
conv6 (1 × 4)		64		17		18		18		4		12		16	
gru1		120		27		27		11		1		15		10	
gru2		120		31		31		15		1		17		10	
Test accuracy		94.6%		94.7%		94.7%		94.7%		16.7%		94.7%		94.7%	
Execution time t (Nexus 5)		26.2 ms		19.5 ms		17.9 ms		18.3 ms		14.1 ms		15.3 ms		15.8 ms	
$t - t_{min}$ (Nexus 5)		12.1 ms		5.4 ms		3.8 ms		4.2 ms		18.4 ms		1.2 ms		1.7 ms	
Execution time t (Galaxy Nexus)		70.9 ms		30.1 ms		27.4 ms		28.2 ms				22.6 ms		22.0 ms	
$t - t_{min}$ (Galaxy Nexus)		52.5 ms		11.7 ms		9.0 ms		9.8 ms				4.2 ms		3.6 ms	




(a) Tradeoff between testing accuracy and execution time on Nexus 5.




(b) Tradeoff between testing accuracy and execution time on Galaxy Nexus.



(c) Tradeoff between testing accuracy and compressed parameter size.



(d) Tradeoff between testing accuracy and energy consumption on Nexus 5.



(e) Tradeoff between testing accuracy and energy consumption on Galaxy Nexus.

Figure 7: System performance tradeoff for DeepSense on HHAR dataset

empowers existing compression algorithms to reduce more execution time with almost the same amount of parameters.

5.2.2 Large-scale image recognition on ImageNet. This is a large-scale vision based task, image recognition based on a high-resolution camera. During this experiment, we use ImageNet as our training and testing dataset. The ImageNet dataset consists of 1.2 million 224×224 color images in 1000 classes with 100,000 images for testing.

During the evaluation, we still use VGGNet structure as the original network structure. The detailed structures of best compressed models without accuracy degradation of all algorithms are shown in Table 7. Note that the original VGGNet for 224×224 colour image input is too large for running on two testing hardware. FastDeepIoT achieves the best performance on the execution time among all methods. Compared with the state-of-the-art DeepIoT method, FastDeepIoT can further reduce the execution time by 59% to 62%. DeepIoT+localMin still outperforms DeepIoT by reducing around 5% to 10% of execution time. In addition, FastDeepIoT can further reduce 25% to 29% of execution time compared with DeepIoT+FLOPs.

Figure 6a and 6b shows the tradeoff between testing top-5 accuracy and execution time for all algorithms. FastDeepIoT consistently outperforms all other algorithms by a significant margin. With the help of execution-time local minima,

DeepIoT+localMin can still outperform DeepIoT in all cases. DeepIoT+FLOPs performs better than DeepIoT in this case.

Figure 6d and 6e illustrates the tradeoff between testing top-5 accuracy and energy consumptions. FastDeepIoT outperforms all algorithms with a large margin. However, FastDeepIoT with the Galaxy Nexus execution time model is not the most energy-saving compression method on the Galaxy Nexus device. Also, DeepIoT+localMin cannot consistently outperform DeepIoT on energy saving. These two observations witness the discrepancies between the execution time and energy modeling on mobile devices. Figure 6c shows the tradeoff between testing accuracy and left proportion of model parameters. Again, all methods show the similar tradeoff, which indicates that FastDeepIoT is a parameter-efficient method on execution time reduction.

5.2.3 Heterogeneous human activity recognition. This is a human-centric context sensing application, recognizing human activities with accelerometer and gyroscope. Especially, we are considering the heterogeneous human activity recognition (HHAR). This task focuses on the generalization ability with human who has not appeared in the training dataset. During this experiment, we use the dataset collected by Allan et al. [30]. During this evaluation, we use DeepSense structure as the original network structure [34]. Table 8 illustrates

the detailed structure of the original network and final compressed networks generated by four algorithms with no degradation on testing accuracy. As shown in Table 8, FastDeepIoT achieves the best performance on two devices with the corresponding execution time models. Compared with DeepIoT, FastDeepIoT can further reduce the model execution time by 22% to 42%. During the compressing process, we observe that all compressed models tend to approach a model execution time lower bound, which has not been seen in the previous two experiments. In order to obtain the lower bound, we build a DeepSense structure with all hidden units that equal to 1, and then applies Algorithm 2 to find the structure that triggers local minimum. The resulted structure is illustrated in Table 8 denoted by t_{\min} . If we calculate the deductible model execution time by subtracting t_{\min} from the model execution time, compared with DeepIoT, FastDeepIoT can reduce the deductible execution time by 69% to 78%.

Furthermore, we can attempt to deduce the fundamental cause of the lower bound with our execution time model. As shown in (6), the execution time of recurrent layer is partially controlled by the number of *step*, which can be interpreted as an initialization overhead for each *step* in the recurrent layer. We can use an example to illustrate the relationship between the *step* overhead and this lower bound. In our experiment, there are 20 steps in the GRU. The coefficient of *step* on Nexus 5 is 0.666 ms. Therefore, the lower bound is $14.1 \approx 20 \times 0.666$ ms. Thus, only algorithms dealing with reducing recurrent-layer steps can help further reducing the model execution time. Unfortunately, to the best of our knowledge, there is no existing work that solves this problem. However, our empirical observation and execution time model reveal an interesting problem that requires future research.

The tradeoffs between testing accuracy and execution time for different algorithms are illustrated in Figure 7a and 7b. FastDeepIoT still achieves the best tradeoff for all cases. The tradeoffs between testing accuracy and energy consumption are illustrated in Figure 7d and 7e. FastDeepIoT performs better than all other baselines in almost all cases. The tradeoffs between testing accuracy and remaining proportion of model parameters are illustrated in Figure 7c. All algorithms show comparable results.

6 RELATED WORK

A key direction in embedded sensing literature is to speed up progressively more complex and interesting applications on resource-constraint embedded and mobile devices. Recent studies start focusing on speeding up deep neural networks through model compression. Han et al. propose a magnitude-based compression algorithm, illustrating promising results on resource-efficient deep neural networks with model compression [13]. Bhattacharya et al. design a sparse-coding and matrix factorization based solution to factorize neural networks into low-complexity structure for reducing resource

consumption [3]. Yao et al. propose a reinforcement learning based adaptive dropout solution to explore the less-redundant network structure for mobile and embedded devices [38]. All these previous compression algorithms focus on reducing the model parameters, while taking execution time speed-up as a by-product. Therefore, these compression methods inevitably show inferior performance on execution time reduction. To the best of our knowledge, FastDeepIoT is the first framework to understand the impact of changing neural network structure on model execution time, and to empower existing compression algorithms to reduce the execution time on mobile and embedded devices properly.

7 CONCLUSION AND FUTURE WORK

In this paper, we introduced FastDeepIoT, a framework for understanding and minimizing neural network execution time on mobile and embedded devices. We proposed a tree-structured linear regression model to figure out the causes of execution-time nonlinearity and to interpret execution time through explanatory variables. Furthermore, we utilized the execution time model to rebalance the focus of existing structure compression algorithms to reduce the overall execution time properly. We evaluated FastDeepIoT with three representative sensing tasks on two devices, where FastDeepIoT outperformed the state-of-the-art algorithms on reducing execution time and energy consumption with a large margin.

This work is just a first step into the exploration of neural network compression for performance optimization. More profiling results are needed with the different choices of hardware, OS versions, load factors, power scaling, and deep learning libraries. Currently, FastDeepIoT can only support deep learning *structure* compression algorithms. More work is needed to support other deep learning compression methods, such as parameter quantization and pruning [13]. The execution time model shows that the setup overhead of recurrent layers imposes a lower bound on efficacy of compression. It is a function of recurrent neural network steps, offering another dimension to compress for speeding up recurrent layers. These insights offer avenues for future research on system performance oriented neural network compression for sensing applications.

ACKNOWLEDGMENTS

Research reported in this paper was sponsored in part by NSF under grants CNS 16-18627 and CNS 13-20209 and in part by the Army Research Laboratory under Cooperative Agreements W911NF-09-2-0053 and W911NF-17-2-0196. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, NSF, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] Tensorflow benchmark tool. <https://github.com/tensorflow/tensorflow/tree/r1.4/tensorflow/tools/benchmark>.
- [2] Tensorflow mobile. https://www.tensorflow.org/mobile/mobile_intro.
- [3] S. Bhattacharya and N. D. Lane. Sparsification and separation of deep learning layers for constrained resource inference on wearables. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, pages 176–189. ACM, 2016.
- [4] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [5] L. Breiman. *Classification and regression trees*. Routledge, 2017.
- [6] N. Bui, A. Nguyen, P. Nguyen, H. Truong, A. Ashok, R. Deterding, and T. Vu. Pho2: Smartphone based blood oxygen level measurement systems using near-ir and red wave-guided light. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. ACM, 2017.
- [7] B. Chen, V. Yenamandra, and K. Srinivasan. Tracking keystrokes using wireless signals. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, pages 31–44. ACM, 2015.
- [8] H. Drucker, C. J. Burges, L. Kaufman, A. J. Smola, and V. Vapnik. Support vector regression machines. In *Advances in neural information processing systems*, pages 155–161, 1997.
- [9] M. Eichelberger, K. Luchsinger, S. Tanner, and R. Wattenhofer. Indoor localization with aircraft signals. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, 2017.
- [10] J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- [11] P. Georgiev, S. Bhattacharya, N. D. Lane, and C. Mascolo. Low-resource multi-task audio sensing for mobile and embedded devices via shared deep neural network representations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):50, 2017.
- [12] R. B. Girshick. Fast r-cnn. *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015.
- [13] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [14] Y. He, X. Shen, Y. Liu, L. Mo, and G. Dai. Listen: Non-interactive localization in wireless camera sensor networks. In *Real-Time Systems Symposium (RTSS), 2010 IEEE 31st*, pages 205–214. IEEE, 2010.
- [15] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017.
- [16] F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size. *CoRR*, abs/1602.07360, 2016.
- [17] N. D. Lane, P. Georgiev, and L. Qendro. Deeppear: robust smartphone audio sensing in unconstrained acoustic environments using deep learning. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 283–294. ACM, 2015.
- [18] K. Langendoen and N. Reijers. Distributed localization in wireless sensor networks: a quantitative comparison. *Computer Networks*, 43(4):499–518, 2003.
- [19] P. Lazik, N. Rajagopal, O. Shih, B. Sinopoli, and A. Rowe. Alps: A bluetooth and ultrasound platform for mapping and localization. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 73–84. ACM, 2015.
- [20] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- [21] C.-Y. Li, Y.-C. Chen, W.-J. Chen, P. Huang, and H.-h. Chu. Sensor-embedded teeth for oral activity recognition. In *Proceedings of the 2013 international symposium on wearable computers*, pages 41–44. ACM, 2013.
- [22] M. Mirshekari, S. Pan, P. Zhang, and H. Y. Noh. Characterizing wave propagation to improve indoor step-level person localization using floor vibration. In *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2016*, volume 9803, page 980305. International Society for Optics and Photonics, 2016.
- [23] S. Nirjon, J. Gummeson, D. Gelb, and K.-H. Kim. Typingring: A wearable ring platform for text input. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, pages 227–239. ACM, 2015.
- [24] V. Radu, C. Tong, S. Bhattacharya, N. D. Lane, C. Mascolo, M. K. Marina, and F. Kawsar. Multimodal deep learning for activity and context recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(4):157, 2018.
- [25] T. Rahman, A. T. Adams, R. V. Ravichandran, M. Zhang, S. N. Patel, J. A. Kientz, and T. Choudhury. Dopplesleep: A contactless unobtrusive sleep sensing system using short-range doppler radar. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 39–50. ACM, 2015.
- [26] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [27] R. Rigamonti, A. Sironi, V. Lepetit, and P. Fua. Learning separable filters. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2754–2761. IEEE, 2013.
- [28] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [29] J. M. Sorber, M. Shin, R. Peterson, and D. Kotz. Plug-n-trust: practical trusted sensing for mhealth. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 309–322. ACM, 2012.
- [30] A. Stisen, H. Blunck, S. Bhattacharya, T. S. Prentow, M. B. Kjærgaard, A. Dey, T. Sonne, and M. M. Jensen. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 127–140. ACM, 2015.
- [31] B. Wei, W. Hu, M. Yang, and C. T. Chou. Radio-based device-free activity recognition with radio frequency interference. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, pages 154–165. ACM, 2015.
- [32] H. Wen, Z. Xiao, N. Trigoni, and P. Blunsom. On assessing the accuracy of positioning systems in indoor environments. In *European Conference on Wireless Sensor Networks*, pages 1–17. Springer, 2013.
- [33] Y. Xiang, R. Piedrahita, R. P. Dick, M. Hannigan, Q. Lv, and L. Shang. A hybrid sensor system for indoor air quality monitoring. In *Distributed Computing in Sensor Systems (DCOSS), 2013 IEEE International Conference on*, pages 96–104. IEEE, 2013.
- [34] S. Yao, S. Hu, Y. Zhao, A. Zhang, and T. Abdelzaher. DeepSense: a unified deep learning framework for time-series mobile sensing data processing. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 2017.
- [35] S. Yao, Y. Zhao, H. Shao, A. Zhang, C. Zhang, S. Li, and T. Abdelzaher. Rdeepsense: Reliable deep mobile computing models with uncertainty estimations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(4):173, 2018.
- [36] S. Yao, Y. Zhao, H. Shao, C. Zhang, A. Zhang, S. Hu, D. Liu, S. Liu, L. Su, and T. Abdelzaher. Sensegan: Enabling deep learning for internet of things with a semi-supervised framework. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):144, 2018.
- [37] S. Yao, Y. Zhao, A. Zhang, S. Hu, H. Shao, C. Zhang, L. Su, and T. Abdelzaher. Deep learning for the internet of things. *Computer*, 51(5):32–41, 2018.

- [38] S. Yao, Y. Zhao, A. Zhang, L. Su, and T. Abdelzaher. Deepiot: Compressing deep neural network structures for sensing systems with a compressor-critic framework. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. ACM, 2017.
- [39] H. Zhang, W. Du, P. Zhou, M. Li, and P. Mohapatra. Dopenc: acoustic-based encounter profiling using smartphones. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 294–307. ACM, 2016.
- [40] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *CoRR*, abs/1707.01083, 2017.