

# Dynamic Transformer for Efficient Machine Translation on Embedded Devices

Hishan Parry, Lei Xun, Amin Sabet, Jia Bi, Jonathon Hare, Geoff V. Merrett

*School of Electronics and Computer Science*

*University of Southampton*

Southampton, UK

{hmrp1r17, lx2u16, ms4r18, J.Bi, jsh2, g.merrett}@southampton.ac.uk

**Abstract**—The Transformer architecture is widely used for machine translation tasks. However, its resource-intensive nature makes it challenging to implement on constrained embedded devices, particularly where available hardware resources can vary at run-time. We propose a dynamic machine translation model that scales the Transformer architecture based on the available resources at any particular time. The proposed approach, ‘Dynamic-HAT’, uses a HAT SuperTransformer as the backbone to search for SubTransformers with different accuracy-latency trade-offs at design time. The optimal SubTransformers are sampled from the SuperTransformer at run-time, depending on latency constraints. The Dynamic-HAT is tested on the Jetson Nano and the approach uses inherited SubTransformers sampled directly from the SuperTransformer with a switching time of  $< 1s$ . Using inherited SubTransformers results in a BLEU score loss of  $\leq 1.5\%$  because the SubTransformer configuration is not retrained from scratch after sampling. However, to recover this loss in performance, the dimensions of the design space can be reduced to tailor it to a family of target hardware. The new reduced design space results in a BLEU score increase of approximately 1% for sub-optimal models from the original design space, with a wide range for performance scaling between 0.356s - 1.526s for the GPU and 2.9s - 7.31s for the CPU.

**Index Terms**—Dynamic DNNs for NLP, Efficient Transformer, Embedded platform, Runtime Resource Management

## I. INTRODUCTION

Machine translation is a fast-growing application of Natural Language Processing (NLP) tasks on embedded devices. However, the most widely used machine learning architecture for NLP, the Transformer [1], is computationally expensive and hence challenging to deploy on embedded devices with limited computational and energy resources. Various approaches have been proposed to address the computational and energy demands of Transformers on embedded devices, such as the Lite Transformer [2], which combines convolution with the standard attention mechanism of the Transformer followed by pruning, allowing better performance in embedded NLP applications. These approaches, which are static DNNs, provide an optimal Transformer architecture for the target application performance requirements based on the measurement of fixed hardware resources. However, embedded devices often run many applications on several heterogeneous cores, and the resources a Transformer was optimized for may not be available at run-time [3]. Therefore, the performance requirements




Fig. 1. Experimental results to showing how latency constraints set during design time optimization are not met at run-time for a Transformer model on the Jetson Nano, due to a varying availability of computing resources or power.

of the application can be violated. Figure 1 illustrates how the desired latency constraint set for a Transformer model can be violated at run-time on the Jetson Nano. For example, on both the GPU and CPU setups, the latency constraint holds when all resources can be dedicated to the Transformer model at full power. When the Jetson is in low-power mode, the reduced frequency of the GPU or the reduced number of CPU cores are not able to meet the required latency. The latency constraint is also not met when the Jetson is at full power, but is dedicating resources to other applications. The frequency and number of cores used may vary at run-time due to power constraints or other processes running simultaneously and using up resources.

Given that hardware resource availability can change dynamically at run-time, dynamic DNNs [4] [5] have been proposed to address this issue. A dynamic DNN contains various sub-networks, each with its accuracy and resource requirements. Given varying available hardware resources and application requirements, different subnetworks can be selected. Previous approaches contain an additional training pipeline to generate these subnetworks. However, state-of-the-art NAS approaches [6] already have an efficient backbone supernet and hence, the additional training pipeline can be skipped.

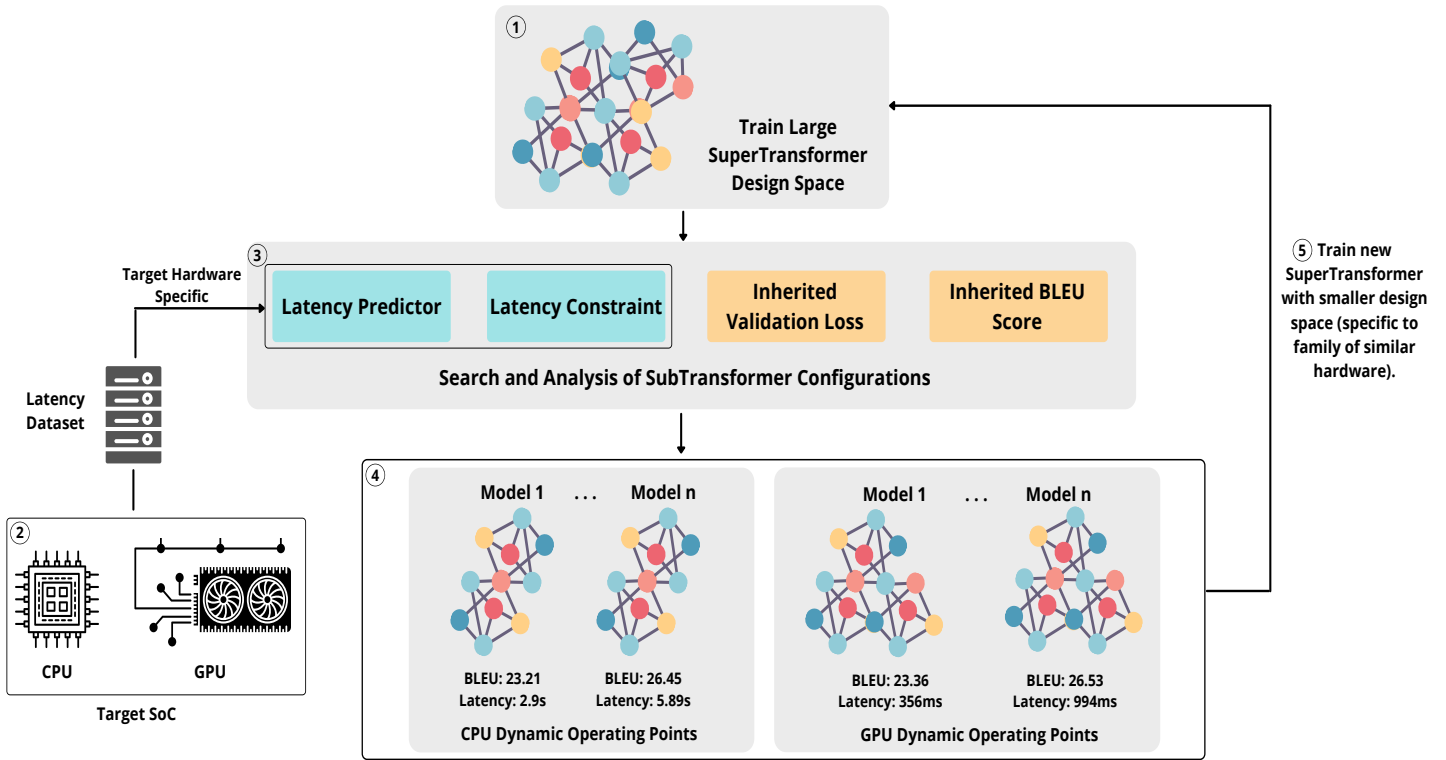


Fig. 2. An overview of the steps in the proposed Dynamic-HAT: (1) A SuperTransformer with a large design space is first trained. (2) Latency datasets are collected from the target hardware. (3) Optimal SubTransformers are searched for the target hardware based on the provided latency constraints. The SubTransformer configurations are analyzed to find the dimensions of the design space preferred by the target hardware (or family of target hardware). (4) These searched SubTransformer configurations are used by the Dynamic-HAT as operating points to meet latency constraints on time-varying hardware resources. (5) The last step is to train a new SuperTransformer with a reduced design space based on the target hardware. This attempts to mitigate the performance loss arising from the use of inherited SubTransformers.

This paper proposes Dynamic-HAT, a dynamic Transformer that can change the complexity of its architecture at run-time. The approach uses Hardware-Aware Transformers (HAT) [6] as the backbone. First, a SuperTransformer is trained once. The SuperTransformer has a large design space containing  $10^{15}$  SubTransformers which encompass optimal model architectures for various accuracy-latency trade-offs. The evolutionary search in HAT to find an optimal SubTransformer to fit a given latency constraint is too expensive to perform at run-time. Hence, we pre-select a small set of optimal SubTransformer configurations for fast switching at run-time. This approach is flexible as it allows changing several hidden dimensions of the Transformer architecture through the HAT backbone with no need for retraining before inference, unlike previous dynamic DNN approaches. These Inherited SubTransformers, which are sampled directly from the SuperTransformer, closely approximates a similar configuration fully trained from scratch with a slight decrease in performance. We also explore different techniques to reduce the performance loss when using Inherited SubTransformers. Sampling inherited SubTransformers allows scaling the Transformer model at run-time to fit changing latency constraints, meaning an optimal Transformer model can always be found for different availability of resources on the same device. The contributions of this paper are:

- Dynamic Transformer architecture using HAT as the backbone of a SuperTransformer, allows SubTransformer configurations to be sampled at run-time without additional training.
- A run-time approach to dynamically switch between Transformer architectures with switching time  $< 1s$ . Dynamic model switching allows the device to meet time-varying hardware resource availability, and maintain performance constraints of machine translation applications.
- Dynamic-HAT reduces the SuperTransformer design space for a set of similar target hardware devices, resulting in a BLEU score increase of up to 1%. We demonstrate Dynamic-HAT using the Jetson Nano, with a wide range for performance scaling between 0.356s - 1.526s for the GPU and 2.9s - 7.31s for the CPU.

## II. RELATED WORK

Several methods have been developed to counter the computational complexity of the Transformer model for efficient NLP, which can be divided into static and dynamic approaches.

### A. Static Approaches

Although standard machine translation models such as the original Transformer [1] and BERT [7] are computation and

resource-intensive, there exist off-line model compression approaches that allow deployment on embedded hardware platforms [8] [9]. A popular method of static model compression is pruning, as proposed by Zhu et al. [10], where sparsity is introduced to deep neural networks in an attempt to reduce the number of computations and make a smaller model. Another static approach to make machine learning models more efficient on resource-limited hardware is quantization, following Jacob et al. [11], which reduces the number of floating point operations to lower memory requirements and computational cost. However, most of these works only compress the model to a fixed size (static), such as the Lite Transformer [2], so the proposed model may not be optimal at run-time when resource constraints change.

### B. Dynamic Approaches

An example of a dynamic model for NLP tasks is DynamicBERT [4], a BERT model which can adjust the model size and latency by changing the width and depth of the network. Another approach to dynamic neural machine translation is the scalable transformer [5], which can change the feature dimensions for all sub-layers in the encoder/decoder. Pruning can also be used for dynamic implementations, such as dynamic sparsity neural networks, which have a set number of model sparsity levels that can be changed at run-time [12].

Compared to previous dynamic approaches, the proposed Dynamic-HAT is more flexible due to the use of a HAT backbone [6]. HAT is more similar to the original Transformer, but with a larger number of elastic hidden dimensions. HAT also has the benefit of arbitrary encoder-decoder attention and flexible decoder layers. HAT addresses the environmental concerns of Transformers by training a SuperTransformer (encompassing a large design space) only once. This SuperTransformer is used to sample SubTransformers for different target hardware under varying latency constraints. The design space encompasses about  $10^{15}$  SubTransformer configurations resulting in a wide range of model sizes and latencies, which makes it possible to find a SubTransformer model for a range of hardware platforms spanning from powerful computers with NVIDIA GPU's to tiny IoT devices with ARM CPUs. Instead of just targeting different hardware platforms, our approach moves one step further by using a library of pre-selected optimal SubTransformers to target dynamically available hardware resources at run-time. The small library also enables our approach to have real-time model switching capability, which is not possible using a static HAT model.

Dynamic-OFA [13] is an image classification dynamic DNN based on the Once-for-All (OFA) approach [14], where a large model is first trained from scratch. Several optimal sub-networks on the Pareto front can be pre-sampled from this large model at design time, and adapted to best match the hardware constraints during run-time. Our Dynamic-HAT approach takes inspiration from Dynamic-OFA, whereas our application is NLP.

TABLE I  
PERFORMANCE COMPARISON BETWEEN INHERITED SUBTRANSFORMERS AND SUBTRANSFORMERS TRAINED FROM SCRATCH. THE RESULTS ARE OBTAINED BY RUNNING HAT ARCHITECTURE ON THE JETSON NANO GPU AND COMPARING DIRECTLY SAMPLED SUBTRANSFORMERS TO THEIR RETRAINED COUNTERPARTS.

Latency Constraint (ms)	Inherited BLEU Score	From-Scratch BLEU Score
500	26.17	26.33
1000	26.67	26.84
1500	26.26	27.21

### III. DYNAMIC-HAT

This section introduces our proposed dynamic machine translation model that is built based on HAT [6]. Figure 2 shows an overview of the proposed Dynamic-HAT. A SuperTransformer with a large design space is first trained from scratch. Different SubTransformer configurations can be sampled from the SuperTransformer and used for inference during run-time, depending on resource constraints, which allows the Transformer model to be trained only once but used for a variety of constraints.

#### A. Proposed Dynamic Model for Machine Translation

The final step in the HAT [6] process is to train the searched SubTransformer from scratch before being deployed to the target hardware. Table I shows a performance comparison between sampled Inherited SubTransformers and their counterparts that have been trained from scratch. Retraining before inference increases the performance of the SubTransformer configuration by approximately 1.5% on average. Inherited SubTransformers perform similar to SubTransformers retrained from scratch before inference. This is because all SubTransformers in the SuperTransformer design space are uniformly trained with shared weights. Hence, Inherited SubTransformers have the same relative performance order as SubTransformers trained from scratch.

Using Inherited SubTransformers would be the more favourable approach when designing a dynamic model based on HAT [6] as the model does not need to be retrained before inference. Hence, we get a wide variety of model configuration choices at run-time, with only one training cost. Another option, training the searched SubTransformer configurations again from scratch and switching between these trained individual models, although better in performance, is still memory expensive and results in a bigger carbon footprint for training. As illustrated by Figure 2, our approach to Dynamic-HAT contains the following steps:

- 1) A SuperTransformer model for the specific translation task is trained once, with uniformly trained weight-shared SubTransformers (adapted from HAT [6]).
- 2) A latency dataset is created on the target hardware device.
- 3) An evolutionary search is performed to find the optimal SubTransformer configuration for various latency constraints supported by the target hardware. The part of

the design space used by the target hardware for this latency range is analyzed and recorded.

- 4) The searched set of optimal SubTransformers are used as operating points to switch the SubTransformer model when latency constraints change at run-time.
- 5) A new SuperTransformer design space can be trained, with reduced dimensions specialized for a product family of similar target hardware, leading to performance improvements in BLEU score.

### B. Reducing the Search Space

This section elaborates on Step 5 of the Dynamic-HAT approach detailed in Figure 2. Our dynamic approach using Inherited SubTransformers, as mentioned previously, results in approximately  $\leq 1.5\%$  degradation in BLEU score as the model is not retrained before inference. To mitigate this loss in performance, we attempt to reduce the search space of the SuperTransformer to decrease the number of parameters and model configurations. Traditionally, the HAT approach contains a large design space containing SubTransformers suitable for embedded devices and high-end GPUs. Different devices have varying hardware capabilities, and the SubTransformer model searched for each device at a specific latency constraint will be different. However, devices running on similar hardware configurations perform better using the same type of model structure and, thus, specific parts of the SuperTransformer space. For example, when running an evolutionary search, GPUs find wide but shallow models while CPUs find deep but thin models [6]. Hence, it would be beneficial to reduce the SuperTransformer design space and tailor it to a family of the target hardware, such as ARM CPUs.

This process has to be done by analyzing the searched HAT models for the target devices, choosing only the parameters in the best-performing configurations, and eliminating the rest. As this procedure would require a new SuperTransformer to be trained from scratch (with fewer parameter choices), training one SuperTransformer for a large group of devices with similar hardware configurations would be the more environmentally friendly approach. As a result, this step remains optional in our dynamic machine translation model, with the additional benefit of improved translation performance.

## IV. EXPERIMENTAL EVALUATION

### A. Experiment Setups

*a) Datasets:* The experiments are conducted on the WMT'14 En-De translation task consisting of 4.5M pairs of training sentences. We use 32K source-target byte pair encoding for the vocabulary that trains on WMT'16, validate on newstest2013, and test on newstest2014.

*b) Evaluation Metrics:* To evaluate the effectiveness of translated sentences, we use the BLEU score with compound splitting, and case-sensitive tokenization, consistent with [1], [6]. The latency of a model is measured by recording the time taken to translate from a source sentence to a target sentence of the same length (30 words). The latency is measured 300

times, removing the fastest and slowest 10% and taking the average of the other 80%, consistent with [6]. The baseline Transformer used for the SuperTransformer in this experiment is the HAT approach [6].

*c) Hardware Platform Settings:* Our experiments are implemented on an NVIDIA Jetson Nano, which is an embedded computer with an ARM A57 CPU and a 128-core Maxwell GPU for inference. As the Jetson Nano is a compact and powerful computer specifically made to run machine learning models within a low power budget of 5-10 Watts, it allows us to evaluate the performance of SubTransformers for a wide range of latency constraints.

*d) SuperTransformer Setups:* Similar to the original Transformer from [1], the SuperTransformer in HAT consists of an encoder-decoder structure. Each encoder/decoder consists of multiple layers of stacked identical encoders/decoders. It contains Multi-Head Attention, Feed Forward Network (FFN), and Masked Multi-Head Attention sub-layers with elastic dimensions, which allows changing the SubTransformer configuration. The number of encoder/decoder layers and their input embedding dimensions are also elastic.

Unlike the encoder, the decoder cannot be parallelized due to its auto-regressive nature, which results in the decoder performing a larger number of computations than the encoder. As the encoder accounts for less than 5% of the total measured latency of the model, we are suggested by HAT to fix the number of encoder layers at 6 [6]. Following the pre-trained SuperTransformer setup from HAT, the design space has the following choices for each parameter:

- {512, 640} input embedding dimensions for the encoder/decoder
- {1024, 2048, 3072} number of hidden FFN dimensions
- {4, 8} number of attention heads in each Multi-Head Attention sub-layer
- {1, 2, 3, 4, 5, 6} number of decoder layers
- {-1, 1, 2} for arbitrary encoder-decoder attention: -1 means attending to last one encoder layer, 1 means last two encoder layers, 2 means last three encoder layers.




Fig. 3. BLEU Score vs experimentally measured Latency on the Jetson Nano GPU, showing how the number of decoder layers affect translation performance of the sampled SubTransformer model.

TABLE II  
SEARCHED SUBTRANSFORMER OPERATING POINTS FOR THE JETSON NANO WITH INHERITED BLEU SCORES AND VALIDATION LOSS.

Jetson Nano GPU				Jetson Nano CPU			
Measured Latency (ms)	Inherited BLEU Score	Inherited Validation Loss	No. Decoder Layers	Measured Latency (s)	Inherited BLEU Score	Inherited Validation Loss	No. Decoder Layers
356.11	23.66	4.8229	1	2.90	23.21	4.9437	1
608.95	26.23	4.3821	2	3.96	25.73	4.4163	2
854.85	26.17	4.231	3	4.97	25.63	4.2336	3
994.96	26.67	4.155	4	5.90	26.45	4.1478	5
1255.38	26.52	4.1177	5	6.75	26.31	4.114	6
1526.54	26.53	4.1048	6	7.31	26.33	4.1044	6

### B. Linking Decoder Layers to Performance and Latency

Figure 3 shows the relationship between the BLEU score and measured latency for several operating points on the Jetson Nano GPU. It is observed that the performance of a searched SubTransformer configuration does not increase linearly with latency constraints. The actual performance of the model depends significantly on the number of decoder layers being used. The ‘gaps’ in the graph between data points occur when the number of decoder layers change. Adding one decoder layer to the SubTransformer configuration increases its measured latency by around 300 ms on the GPU and around 1000ms on the CPU.

A higher number of decoder layers does not necessarily lead to better performance in the model, since it greatly depends on the hardware that is running the model. As mentioned before, GPUs perform better with wide but shallow models while CPUs prefer deep but thin models [6]. As GPUs are less sensitive to changes in hidden dimensions and embedded dimensions, while CPUs are very sensitive, a GPU can reduce the number of decoder layers while increasing the dimensions of its sub-layers to keep the same performance and reduce latency. Figure 3 shows the best performing Dynamic-HAT model on the Jetson GPU with a BLEU score of 26.67 has 4 decoder layers. In contrast, the best performing Dynamic-HAT model on the Jetson CPU with a BLEU score of 26.45 has 5 decoder layers. These observations can verify that GPUs search for wide but shallow models – the best Dynamic-HAT model for the GPU has only 4 decoder layers but the largest possible input embedding and hidden dimensions.

Table II displays the relationship between the BLEU score and validation loss of inherited SubTransformer configurations and Figure 4 shows validation loss of the SubTransformer against measured latency. As the latency constraint increases, validation loss of the searched SubTransformer configuration decreases, as seen in Figure 4, but this change is also not linear. A reduction in validation loss generally improves model performance, leading to higher BLEU scores – this can be observed by comparing Figure 3 and Figure 4. In most cases, at latency constraints above 1100ms (for the GPU), 5 or 6 decoder layers are chosen by the evolutionary search – even though the best performing model in terms of BLEU score has 4 decoder layers. As the evolutionary search looks at validation loss, a further decrease in loss would mean that a better performing model has been searched. Although validation loss




Fig. 4. Validation Loss vs experimentally measured Latency on the Jetson Nano GPU, showing how number of decoder layers affect inherited validation loss of the sampled SubTransformer model.

TABLE III  
COMPARING THE SWITCHING TIMES OF STATIC AND DYNAMIC TRANSFORMERS ON THE JETSON NANO.

Model	Switching Time	Procedure
Static HAT [6]	minutes to hours	load weights / search / retrain
Dynamic-HAT	< 1s	sample

continues to decrease with the addition of decoder layers and increasing latency constraints, this change is not always reflected in the BLEU score, as seen in Table II.

### C. Performance Improvements Using New Design Space

The default SuperTransformer design space for HAT is large, as it needs to encompass enough model sizes to span different hardware for a range of latencies. However, we postulate that it may be possible to further optimize SubTransformer sampling and BLEU scores by reducing the dimensions of the design space for a particular platform or product family. As fewer parameters need to be moved around during sampling and inference, the complexity of computation and the number of possible SubTransformer configurations are reduced. We can examine the configurations for the top-5 best performing SubTransformers on the Jetson Nano GPU, and observe which portion of the design space it uses for its optimal SubTransformers. The best performing Jetson Nano models for the GPU, across its entire dynamic range, did not use 640 input embedding dimensions or 1024 hidden dimensions for




Fig. 5. Comparing BLEU scores of operating points for the reduced SuperTransformer design space compared to the original design space.

the FFN layers. The arbitrary encoder-decoder attention sub-layer also only attends to the last one or two encoder layers for optimal SubTransformers on the Jetson GPU. After training for approximately 800 GPU hours on the Nvidia GTX 1080Ti, keeping the same SuperTransformer training setup from HAT [6], the following reduced design space is constructed by eliminating the unused parameter choices:

- {512, 640} input embedding dimensions for the encoder/decoder
- {1024, 2048, 3072} number of hidden FFN dimensions
- {4, 8} number of attention heads in each Multi-HeadAttention sub-layer
- {1, 2, 3, 4, 5, 6} number of decoder layers
- {-1, 1, 2} for arbitrary encoder-decoder attention: -1 means attending to last one encoder layer, 1 means last two encoder layers, 2 means last three encoder layers.

The above design space contains  $10^{11}$  SubTransformer configurations compared to  $10^{15}$  in the original design space. We can observe from Figure 5 that the trend of operating points, including sub-optimal points, for the two design spaces follow each other closely. The time to sample a different SubTransformer model from the SuperTransformer, as seen in Table III, is approximately  $<1s$  on the Jetson Nano, depending on available resources and model size. The new design space increases the BLEU score of SubTransformers for all latency constraints except 500ms, which appears sub-optimal. All other sub-optimal models from the original design space receive a BLEU score increase of approximately 1% in the reduced design space.

## V. CONCLUSION

Static approaches to optimize a Transformer model for efficient neural machine translation are based on fixed hardware resources. However, it is challenging to consistently meet the performance targets on resource-constrained embedded devices with time-varying availability of hardware resources. Our Dynamic-HAT provides a Transformer model that can dynamically scale to meet latency constraints in different runtime scenarios. A SubTransformer configuration that meets a

specified latency constraint can be sampled from the SuperTransformer without retraining in  $<1s$  on the Jetson Nano. The proposed Dynamic-HAT uses Inherited SubTransformers. Hence there is a loss in BLEU score  $\leq 1.5\%$ . To mitigate the BLEU score loss, we reduce the design space of SuperTransformer. Parameters from the design space that are not used by the Jetson Nano hardware are eliminated, and the SuperTransformer is retrained. The new design space results in a BLEU score increase of approximately 1% for sub-optimal operating points from the original design space.

## VI. ACKNOWLEDGEMENT

This work was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/S030069/1. Code will be available open-source on <https://github.com/UoS-EEC>.

## REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention Is All You Need," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [2] Z. Wu, Z. Liu, J. Lin, Y. Lin, and S. Han, "Lite Transformer with Long-Short Range Attention," in *International Conference on Learning Representations (ICLR)*, 2020.
- [3] L. Xun, L. Tran-Thanh, B. M. Al-Hashimi, and G. V. Merrett, "Optimising Resource Management for Embedded Machine Learning," in *Design, Automation & Test in Europe Conference (DATE)*, 2020.
- [4] L. Hou, Z. Huang, L. Shang, X. Jiang, X. Chen, and Q. Liu, "DynaBERT: Dynamic BERT with Adaptive Width and Depth," in *Conference on Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, 2020.
- [5] P. Gao, S. Geng, X. Wang, J. Dai, and H. Li, "Scalable Transformers for Neural Machine Translation," *arXiv preprint arXiv:2106.02242*, 2021.
- [6] H. Wang, Z. Wu, Z. Liu, H. Cai, L. Zhu, C. Gan, and S. Han, "HAT: Hardware-Aware Transformers for Efficient Natural Language Processing," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [8] S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," in *International Conference on Learning Representations (ICLR)*, 2015.
- [9] T.-J. Yang, A. Howard, B. Chen, X. Zhang, A. Go, M. Sandler, V. Sze, and H. Adam, "NetAdapt: Platform-Aware Neural Network Adaptation for Mobile Applications," in *European Conference on Computer Vision (ECCV)*, 2018.
- [10] M. Zhu and S. Gupta, "To Prune, or Not to Prune: Exploring the Efficacy of Pruning for Model Compression," in *International Conference on Learning Representations (ICLR)*, 2018.
- [11] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 2704–2713.
- [12] Z. Wu, D. Zhao, Q. Liang, J. Yu, A. Gulati, and R. Pang, "Dynamic Sparsity Neural Networks for Automatic Speech Recognition," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6014–6018.
- [13] W. Lou, L. Xun, A. Sabet, J. Bi, J. Hare, and G. V. Merrett, "Dynamic-OFA: Runtime DNN Architecture Switching for Performance Scaling on Heterogeneous Embedded Platforms," in *Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2021, pp. 3110–3118.
- [14] H. Cai, C. Gan, T. Wang, Z. Zhang, and S. Han, "Once-for-All: Train One Network and Specialize it for Efficient Deployment," in *International Conference on Learning Representations (ICLR)*, 2020.