



# RepCNN: Micro-sized, Mighty Models for Wakeword Detection

Arnav Kundu\*, Prateeth Nayak\*, Priyanka Padmanabhan, Devang Naik

Apple Inc.

{a.kundu,prateethvnayak,priyanka-padmanabhan,naik.d}@apple.com

## Abstract

Always-on machine learning models require a very low memory and compute footprint. Their restricted parameter count limits the model's capacity to learn, and the effectiveness of the usual training algorithms to find the best parameters. Here we show that a small convolutional model can be better trained by first refactoring its computation into a larger redundant multi-branched architecture. Then, for inference, we algebraically re-parameterize the trained model into the single-branched form with fewer parameters for a lower memory footprint and compute cost. Using this technique, we show that our always-on wake-word detector model, RepCNN, provides a good trade-off between latency and accuracy during inference. RepCNN re-parameterized models are 43% more accurate than a uni-branch convolutional model while having the same runtime. RepCNN also meets the accuracy of complex architectures like BC-ResNet, while having 2x lesser peak memory usage and 10x faster runtime.

**Index Terms:** speech recognition, efficientML

## 1. Introduction

The design and development of efficient machine learning models for always-on mobile applications have progressed significantly, witnessing an overall trend of reduced FLOPs (floating-point operations) and memory usage while maintaining or improving accuracy [1, 2, 3]. This poses a challenge of achieving higher accuracy within limited memory and compute. Specifically related to speech and vision applications on mobile devices, novel architecture optimizations [4, 5, 6] are proposed to achieve high accuracy, but their components, such as residual connections make the inference graph complex (although FLOP is low) and lead to increased runtime memory usage and end-to-end inference latency [7, 8]. This makes it challenging to deploy these models on resource-constrained devices, ultimately leading to higher latency or unsupported configurations. Due to these strict memory and compute demands, existing models developed for always-on keyword spotting have often been simplistic inference graphs [9, 10, 11].

Although these simpler models are efficient in terms of memory usage, they often fall short in achieving comparable accuracy to more complex multi-branched architectures. One reason for this is that multi-branch topology, like ResNet, create an implicit ensemble of shallower models. This ensemble effect inherently mitigates the gradient vanishing problem, which can be a challenge for single-branch models, allowing more accurate predictions [12]. To this end, we propose a re-parameterizable fully convolutional encoder architecture (RepCNN) that is:

- Multi-branched during training to accommodate easier gradient flow.
- Simple single branched inference graph to ensure less compute and memory usage during inference execution flow.

Considering that the benefits of multi-branch architectures are predominantly seen during training and contribute to increased inference costs, we have proposed a novel approach to decouple the training and inference model graphs for keyword-spotting. By doing so, we can leverage the advantages of multi-branch training while reducing its impact on inference. Further, we hypothesize that using multiple identical kernels in parallel could introduce new learning paths, potentially enhancing the model's accuracy. During inference, these redundant kernels can be combined into a single kernel, without any increase in latency.

## 2. Related Works

Traditional wake-word detection approaches are largely based on Hidden Markov Models which includes acoustic DNN model that computes phoneme probabilities followed usually by *viterbi*-based HMM decoder with output observation scores[9, 13, 14, 15, 16], but these systems often suffer from loss-metric mismatches and are difficult to train. With advances in recurrent architectures, there have been complex networks[17, 18, 19] which are easier to train and more accurate, however are not suitable for mobile inference due to high latency concerns. With accuracy gains and decent computational complexities, several End-to-End CNN-based models[2, 12, 4, 1, 20] developed for mobile-vision applications provide a great motivation to apply such architectures to wake-word detection. Additionally, convolutional models have the advantage of having streaming audio with input audio length being arbitrary. In this direction, there have been specific single-path CNN-based hand-crafted architectures[10, 11, 21] that have shown good performance at low latency and runtime memory usage. However, only with complex architectures[3, 22, 23, 24] with high parameter count, bottlenecks or skip-connections we are able to see huge gain in accuracy. Novel residual connections[5] in these work provide further gains, but presence of these residual structures during inference on mobile-applications not only adds computational overhead, but also increase runtime memory size required on-device to store these branched intermediate outputs. Models with such inference graphs become unsuitable when operating in strict memory and computation budget.

Recently, structural re-parameterization has shown significant promise in the field of computer vision with [8, 7, 25, 26]. These models use a multi-branched architecture during training and are converted to a single branched architecture during inference. Motivated by these methods, we extend the discrim-

\*Equal Contribution

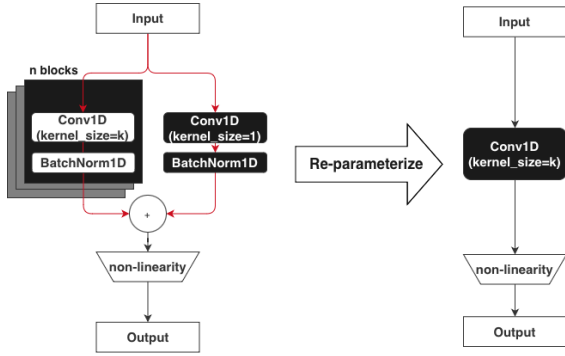


Figure 1: *RepConvBlock*: Re-parameterizable Convolutional Block

inative training work[6, 11, 14] with the approach of multi-branch topology for 1D depth-wise convolutions during training to over-parameterize and increase network capacity, and during inference we linearly combine the convolutional parameters to a trivial architecture that has lesser parameters as well as lower latency and memory usage.

### 3. Model Architecture

As discussed earlier, conventional single branch architectures have poor performance, primarily because of being difficult to train. Multi-branch architectures on the other hand are easier to train while being slower for on-device applications. This is true despite having lower theoretical FLOPs, primarily because of memory access cost [4]. Therefore, we use structural re-parameterization to have different training and inference graphs. This section explains our model architecture in detail that incorporates the benefits of multi-branch training graph and converting to a simple inference graph.

#### 3.1. RepConvBlock

Based on the principle of structural re-parameterization, we introduce Re-parameterizable Convolutional Block (*RepConvBlock*) as shown in Figure 1. This acts as the foundational element of our work with the training graph shown on the left and inference graph shown on the right. Since we are focusing on tiny ML models for speech processing, the *RepConvBlock* is designed for 1-D depth-wise separable convolution to be compute and memory efficient [8]. The *RepConvBlock* consists of 3 architectural elements during training as described below:

- Multiple convolutional kernels of the same size [8], followed by batch normalization layers in parallel in the training graph. This is done to over parameterize the training graph which has proven to be beneficial in prior works [8, 7].
- 1x1 convolutional kernel followed by a batch normalization layer parallel to the kernels above to emulate a skip connection.
- ReLU layers as non-linearity.

During inference the parallel convolutional and batch normalization layers are replaced by an equivalent fused convolutional layer. The fusion process is described in Section 3.2. We also employ non-linear layers that are replaced by a clip layer to keep the activations bounded.




Figure 2: *RepCNN* training architecture. The *RepConvBlocks* are parameterized as (input-channels, output-channels, kernel-size, num-branches). Here  $p=q=2$  i.e., *block1* and *block2* are stacked twice. *block1<sub>1</sub>*, *block1<sub>2</sub>*, *block2<sub>1</sub>*, *block2<sub>2</sub>* use kernel-size as  $k = 7, 9, 11, 13$  respectively. non linearity = ReLU.

#### 3.2. Reparameterization

A 1-D convolutional layer parameterized by  $w_k$ , followed by a batch normalization layer, during inference can be re-parameterized into another 1-D convolutional layer parameterized by  $(W_k, W_b)$  as shown in Equation (1).

$$\begin{aligned}
 y &= bn(\text{Conv}(x)) \\
 &= \frac{\alpha_{bn}}{\sigma_{bn}} \times \text{Conv}(x) - \frac{\mu_{bn}}{\sigma_{bn}} + \beta_{bn} \\
 &= \frac{\alpha_{bn}}{\sigma_{bn}} \times w_k * x - \frac{\mu_{bn}}{\sigma_{bn}} + \beta_{bn} \\
 &= \hat{W}_k * x + W_b
 \end{aligned} \tag{1}$$

Here  $k$  is the size of the kernel,  $W_b$  is the bias of the re-parameterized kernel.  $\alpha, \beta, \mu, \sigma$  correspond to the weight, bias, tracked mean and standard deviation for the batch normalization layer, respectively. We re-parameterize individual branches into equivalent convolutional kernels parameterized by  $W_k^i$ .

Consequently, parallel convolutional kernels of size  $k$  can be fused together into a single kernel by simply adding the corresponding weights. The branch with kernel size = 1 can be represented as a kernel of size  $k$  where all weights except the central element are zeros. The re-parameterization of the parallel branches  $\{W_1, W_k^1, W_k^2 \dots\}$  is shown in Equation (2). We get a single convolutional kernel ( $\hat{W}_k$ ) to represent the combination of multiple parallel branches as shown in Figure 1.

$$\begin{aligned}
 y &= W_1 * x + \sum_i (W_k^i * x) \\
 &= W_1 * x + \sum_i (W_k^i) * x \\
 &= \underbrace{[0, \dots, W_1, \dots, 0]}_{k\text{-elements}} + (\sum_i W_k^i) * x \\
 &= \hat{W}_k * x
 \end{aligned} \tag{2}$$

#### 3.3. RepCNN

The design of the *RepCNN* model during training is illustrated in Figure 2. The model comprises of a regular 1D-convolutional layer with kernel size 5 and a stride of 2, followed by a 1D-batch normalization layer and a ReLU layer. The *RepConvBlocks* are used in *block1* and *block2*. In both of these modules we stack 2 *RepConvBlocks* next to one another with the same kernel size followed by a point-wise convolution, batch normalization and ReLU. We use relatively larger kernels in these models with modules of *block1* type using kernels of  $k = 5$  &  $k = 7$ , and modules of *block2* type using kernels of  $k = 11$  &  $k = 13$

respectively. We can use these larger kernels without significant increase in runtime latency and memory usage because of absence of branching during inference, which avoid data movement in between nodes of the model as in [8].

## 4. Experiments and Results

We apply our method for detection of a given wakeword like "Alexa" or "Hey Google" in streaming speech. The following sections will provide further details on our dataset, training strategies, results, and ablation studies that delve into the impact of various architectural choices made in this paper.

### 4.1. Datasets

We have constructed a robust training dataset of 1M audio utterances, each of which includes the wake-word followed by a user query. These audio samples have been meticulously enhanced to mimic real-world acoustic conditions using room-impulse responses (RIRs) along with echo residuals. To augment the training data and improve model robustness, we have included ambient noise samples from various acoustic environments. Furthermore, we applied a diverse range of training-time gain augmentations, varying from 10dB to -40dB, to simulate various volume levels in real-world scenarios. For each utterance in training data, we use an acoustic model to get per-frame forced alignments. This provides us with the indices of start and end frame phonemes of the wake-word. We define a positive sample as a window containing the entire wake-word while ending with indices of the end frame. We also harvest random windows with partial or no wake word in them as negative samples. For every positive sample, we pick 20 random negative windows to mimic the real world probability of wake word being present in streaming speech.

The positive test data encompass both near-field and far-field utterances, clearly capturing the wake-word at 3ft and 6ft distances, providing an extensive range of testing scenarios. The negative test data consists of approximately 2000 hours of dense general speech, ensuring a thorough evaluation of the model's ability to distinguish between relevant and irrelevant cues. These datasets have been collected through internal user studies with informed consent approvals.

### 4.2. Training and Evaluation

In order to conduct a comparative study, we trained three models with different architectural designs. The first model is BC-ResNet[3], which serves as an example of a multi-branch architecture and has the same training and inference graph. The second model is a basic stacked 1D convolutional model[11] which is used to represent a single branch architecture. The third model is our proposed model, which utilizes stacked *RepConvBlocks* in place of simple 1-D depth-wise convolutional layers during training and is re-parameterized during evaluation to single-branch architecture. All these models have roughly 15k parameters and a receptive field of 1.5 seconds to ensure fair comparison.

For input featurization, we apply 26 Mel-Spectral Filter-Banks and extract 16 Mel-frequency cepstral coefficients (MFCCs) at a 10 ms frame rate for every 25 ms of audio, a commonly used approach in speech processing. We use focal-loss [27] as the classification loss to handle the skewed distribution of positive to negative samples in every batch. To make the best use of limited model capacity, for every batch iteration we sort the loss values of the negative samples and only pick the top-

Model	FRR (%) @ 3FA/hr
BC-ResNet[3]	1.66
Single branch Arch.	2.88
<b>RepCNN</b>	<b>1.66</b>

Table 1: Comparative results of BC-ResNet, single branch architecture, and RepCNN on wake-word detection at 3 FA/hr operating point

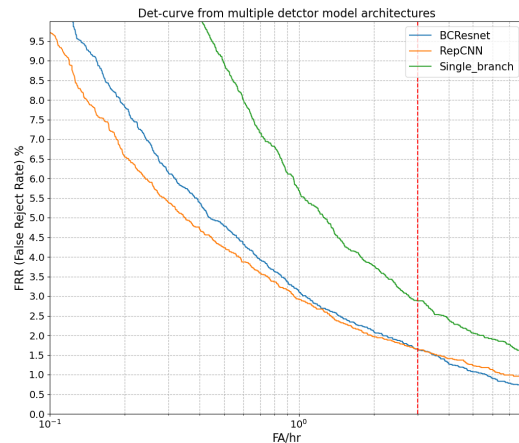


Figure 3: DET curves for model architectures as compared to RepCNN (Orange). We choose 3FA/hr as an ideal operating point for the detection.

$K$  ( $K = 50$ ) hardest negative samples along with all positive sample losses for back-propagation. We also repeat the training with 3 different seeds and the results reported are the average of them.

In our evaluation framework, a detected wake-word is considered a true-positive (TP) if it has any overlap with the ground-truth segment. If there is no overlap, it is counted as a false accept (FA). On the other hand, all positive segments containing wake-word where the model fails to detect are classified as false rejects (FR). We measure the False-Reject Rate (FRR) which is proportion of samples containing the wake-word that were rejected from the positive set. And, we measure the number of False-Activation per hour (FA/hr) using the negative set. To compare between different model performances, we report the FRR at a chosen operating point of 3 FA/hr. We also compute the DET-curve characteristic of the different models.

### 4.3. Results

For comparison the results for all the candidate models are reported in Table 1. These results demonstrate that complex multi-branch architecture like BC-ResNet[3] are approximately 43% more accurate than the single branched stacked convolutional model. RepCNN is able to meet or beat this performance accuracy at 3FA/hr operating point or lower, and the DET-curve shown in Figure 3 shows the differences in accuracy at various operating points. We can see that single branch architecture is worse than multi-branch architecture approaches.

To compare the overall performance of BC-ResNet vs RepCNN, we compute the ROC-AUC of the models. Both BC-ResNet and RepCNN have 97.9% AUC score, however BC-

Model	Latency (ms)	Peak memory (MB)
BC-ResNet[3]	4.5	1.2
RepCNN Training graph (n=2)	1.8	0.8
<b>Re-parameterized RepCNN</b>	<b>0.4</b>	<b>0.5</b>

Table 2: Latency and runtime memory usage to produce one output from various model architectures on iPhone Xs CPU.




Figure 4: Plot of training and validation loss of RepCNN as compared to a no branch architecture.

ResNet is better than RepCNN at higher FA/hr operating points. To validate that our model has a faster inference speed, we have conducted a on-device latency and memory profile of the BC-ResNet multi-branch architecture against the RepCNN training graph, and the RepCNN re-parameterized model on the CPU of an iPhone Xs. The results are presented in Table 2. The latency results show that the complex inference graph of RepCNN re-parameterized is 10x faster than BC-ResNet and consumes half the peak runtime memory usage.

This validates our approach that a multi-branch re-parameterizable model can be designed such that it is as accurate as a complicated BC-ResNet model while being more efficient to run on a device for always-on power constrained applications like wake-word triggering. To test its extensibility, we verified that on Google Speech commands [28], RepCNN achieves the same accuracy as BCResNet with 50% faster runtime and 40% lesser memory.

#### 4.4. Ablation study

To further validate that our approach is easier to train, we plot the training and validation loss curves of RepCNN in Figure 4. For exactly same training setup and hyper-parameters, it can be observed that RepCNN models have a better training and validation loss curve trend compared to the no branch case (a vanilla stacked 1D-CNN similar to [11]). This is realized across different seeds of the training as well.

We also explore the effect of different number of branches in the multi-branch topology during training and compare it against the single-branch as part of the ablation study. Primarily, we are interested in comparing the validation loss curves which is an indicator of the training health and also the final performance metrics. In Figure 5, we verify that over-parameterization does helps to train a better model by looking into the validation loss of various branched versions of RepCNN versus the single-branch architecture. The variations in the loss is due to randomness in sampling hard negative sub-windows and varying augmentations across epochs during training. These models have different degrees of over-




Figure 5: Plot of validation loss of RepCNN with various degrees of over-parameterization. Branches corresponds to the number of parallel 1D depth-wise convolutional kernels apart from the kernel with size 1.

Num. branches	1	2	3	4	5
FRR (%) @ 3 FA/hr	2.16	<b>1.66</b>	1.93	1.98	2.03

Table 3: Ablation study on the effect of accuracy with various levels of over-parameterization

parameterization introduced by the number of parallel convolutional kernels (denoted as  $n$  in Figure 1) in RepConvBlock. The final inference graph and parameter count for all these versions is the same after training due to re-parameterization. Additionally, on evaluation we find that using 2 branches gives the best FRR at chosen operating point while using 3,4 or 5 branches is always better than using 1 branch (single-branch) during training as shown in Table 3. It can be observed that allowing multiple branched paths helps to converge to a better model but beyond a limit such redundancy does not tend to improve the metrics.

## 5. Conclusion

In this work, we have introduced structural re-parameterization of complex multi-branch models to detect wake words accurately and fast. We have shown that we can take advantage of multi-branch architectures during training without impacting inference latency and runtime memory. We achieve it by designing RepCNN which is a multi-branch architecture during training and is re-parameterized into a single branch architecture during inference. We demonstrate that RepCNN is 10x faster than BC-ResNet. It runs as fast as a regular stacked 1D-CNN while being 43% better in accuracy. In our further work in this direction, we plan to improve the gap in accuracy with complex architectures, like BC-ResNet at higher FA operating points, through additions like squeeze and excite components while ensuring simplistic inference graph. Furthermore, we aim to explore compression techniques that help in better re-parameterization between multi-branch convolutions.

## 6. References

- [1] 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*, vol. abs/1704.04861, 2017.
- [2] A. Howard, M. Sandler, G. Chu, L. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, "Searching for mobilenetv3," *CoRR*, vol. abs/1905.02244, 2019.
- [3] B. Kim, S. Chang, J. Lee, and D. Sung, "Broadcasted Residual Learning for Efficient Keyword Spotting," in *Proc. Interspeech 2021*, 2021, pp. 4538–4542.
- [4] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 116–131.
- [5] J. Hu, L. Shen, S. Albanie, G. Sun, and E. Wu, "Squeeze-and-excitation networks," 2019.
- [6] A. Kundu, M. Samragh, M. Cho, P. Padmanabhan, and D. Naik, "Heimdall: Highly efficient method for detection and localization of wake-words," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [7] X. Ding, X. Zhang, N. Ma, J. Han, G. Ding, and J. Sun, "Repvgg: Making vgg-style convnets great again," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 13 733–13 742.
- [8] P. K. A. Vasu, J. Gabriel, J. Zhu, O. Tuzel, and A. Ranjan, "Mobileone: An improved one millisecond mobile backbone," in *CVPR*, 2023. [Online]. Available: <https://arxiv.org/abs/2206.04040>
- [9] S. Sigtia, R. Haynes, H. B. Richards, E. Marchi, and J. S. Bridle, "Efficient voice trigger detection for low resource hardware," in *INTERSPEECH*, 2018.
- [10] T. Sainath and C. Parada, "Convolutional neural networks for small-footprint keyword spotting," in *Interspeech*, 2015.
- [11] T. Higuchi, M. Ghasemzadeh, K. You, and C. Dhir, "Stacked 1D Convolutional Networks for End-to-End Small Footprint Voice Trigger Detection," in *Proc. Interspeech 2020*, 2020, pp. 2592–2596.
- [12] A. Veit, M. J. Wilber, and S. Belongie, "Residual networks behave like ensembles of relatively shallow networks," *Advances in neural information processing systems*, vol. 29, 2016.
- [13] I.-F. Chen and C.-H. Lee, "A hybrid hmm/dnn approach to keyword spotting of short words," in *INTERSPEECH*, 2013, pp. 1574–1578.
- [14] R. Alvarez and H.-J. Park, "End-to-end streaming keyword spotting," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6336–6340.
- [15] R. Prabhavalkar, R. Alvarez, C. Parada, P. Nakkiran, and T. N. Sainath, "Automatic gain control and multi-style training for robust small-footprint keyword spotting with deep neural networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4704–4708.
- [16] A. Shrivastava, A. Kundu, C. Dhir, D. Naik, and O. Tuzel, "Optimize what matters: Training dnn-hmm keyword spotting model using end metric," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 4000–4004.
- [17] S. Fernández, A. Graves, and J. Schmidhuber, "An application of recurrent neural networks to discriminative keyword spotting," in *International Conference on Artificial Neural Networks*. Springer, 2007, pp. 220–229.
- [18] C. Lengerich and A. Hannun, "An end-to-end architecture for keyword spotting and voice activity detection," *arXiv preprint arXiv:1611.09405*, 2016.
- [19] K. Hwang, M. Lee, and W. Sung, "Online keyword spotting with a character-level recurrent neural network," *arXiv preprint arXiv:1512.08903*, 2015.
- [20] M. Sandler, A. G. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation," *CoRR*, vol. abs/1801.04381, 2018.
- [21] R. Tang and J. Lin, "Deep residual learning for small-footprint keyword spotting," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5484–5488.
- [22] S. Choi, S. Seo, B. Shin, H. Byun, M. Kersner, B. Kim, D. Kim, and S. Ha, "Temporal convolution for real-time keyword spotting on mobile devices," *CoRR*, vol. abs/1904.03814, 2019.
- [23] X. Li, X. Wei, and X. Qin, "Small-footprint keyword spotting with multi-scale temporal convolution," 2020.
- [24] S. Majumdar and B. Ginsburg, "MatchboxNet: 1D Time-Channel Separable Convolutional Neural Network Architecture for Speech Commands Recognition," in *Proc. Interspeech 2020*, 2020, pp. 3356–3360.
- [25] C. Chen, Z. Guo, H. Zeng, P. Xiong, and J. Dong, "Repghost: a hardware-efficient ghost module via re-parameterization," *arXiv preprint arXiv:2211.06088*, 2022.
- [26] A. Wang, H. Chen, Z. Lin, H. Pu, and G. Ding, "Repvit: Revisiting mobile cnn from vit perspective," *arXiv preprint arXiv:2307.09283*, 2023.
- [27] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [28] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *ArXiv*, vol. abs/1804.03209, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:4719239>