HIGH SCHOOL EDITION
Journal of Student Research
Volume 14 Issue 1 (2025)

Enhanced Wildfire Detection Using a Mixture of Experts Approach

Brent Kong[1] and Yu Zhang[#]

[1]Scotts Valley High School
[#]Advisor

ABSTRACT

With increasing global wildfire severity, effective fire detection methods are essential to mitigate widespread environmental and health impacts. A recent solution to this phenomena is the application of ensemble machine learning methods, which combine several models to create a more effective one. However, this raises several questions, notably whether an ensemble method is more effective than an individual model or if increasing the number of constituent models leads to overfitting. This paper conducts an ablation study on the Mixture of Experts (MoE) approach for forest fire detection via satellite imagery across a Canadian dataset. The model (MoE6) constitutes all six state-of-the-art architectures, including InceptionNet, ResNet, Vision Transformer (ViT), AlexNet, VGG-Net, and a baseline CNN. Experts of the MoE6 will be systematically removed to form MoE4 and MoE2, which constitute only the top four and top two performing constituent models respectively. We hypothesize that the MoE ensemble approach will outperform any constituent model (two heads are better than one). Furthermore, among the MoE architectures, we hypothesize MoE2 as the top model as it comprehensively integrates characteristics from top model architectures while mitigating overfitting. However, the results show that the original MoE6 was the top performer, achieving a peak accuracy of 93.13% and ROC-AUC of 0.9303. This work provides a promising solution for improving wildfire detection accuracy and response times, potentially reducing the devastation caused by wildfires globally.

Introduction

In recent years, forest fires have intensified and grown in size, with Canada experiencing a significant increase due to its dry and arid climate, especially during the 2023 fire season. By mid-June 2023, Canada had reported 2,619 wildfires, burning a total of 5.3 million hectares, an area 15 times larger than the 10-year average for that period (Government of Canada, 2023). Twelve provinces were heavily impacted, including British Columbia, Alberta, Ontario, Quebec, and Nova Scotia.

The Canadian wildfire season, running from April to mid-October, is influenced by weather and vegetation. Wildfire smoke poses a public health risk, containing toxic chemicals such as ammonia, carbon monoxide, and trace metals, which particularly affect children and the elderly. Quebec was one of the hardest-hit provinces in 2023, with Montreal experiencing the worst air quality globally at one point. By May 2024, Quebec had already reported 80 wildfires (Spector, 2024). Southern Quebec is the study area for this paper, due to its high wildfire susceptibility.

Multiple solutions have been developed in response to the escalating wildfire threat. An effective strategy involves employing machine learning (ML) techniques, notably Convolutional Neural Networks (CNNs), to swiftly and accurately detect forest fires from satellite imagery.

Various researchers have proposed and tested different CNN configurations for forest fire image detection. (You et al., 2024) applied a particle swarm optimization algorithm (PSO) to achieve 82.2% accuracy in forest fire prediction in China, surpassing logistic regression, random forest, support vector machine, k-nearest neighbors, and a standard CNN. Similarly, (Gaur et al., 2024) found CNN-SVM superior, achieving 96.84% testing accuracy compared to 95.79% for the standard CNN.

In (Reis and Turk, 2023), a variety of deep learning algorithms and transfer learning techniques, including InceptionV3, DenseNet121, ResNet50V2, and NASNetMobile, were explored. They achieved a peak accuracy of 99.32% using DenseNet121 pre-trained on ImageNet weights on the Fire Luminosity Airborne-based Machine Learning Evaluation dataset. Finally, (Sathishkumar et al., 2023) applied Learning without Forgetting (LwF) to CNN architectures like VGG16, InceptionV3, and Xception, aiming to reduce training time while enhancing or maintaining classification performance.

Despite achieving peak performance, current models suffer from high computational requirements and lengthy training times. This study investigates the efficacy of the ensemble method Mixture of Experts (MoE). The purpose of this work is to conduct an ablation study on the MoE approach by systematically reducing the number of experts it contains to understand the contributions of each model and the effects of having several or few experts. Variants of the MoE model, MoE6, MoE4, and MoE2, are compared to each other and to individual constituent models like InceptionNet, ResNet, Vision Transformer (ViT), AlexNet, Very Deep Convolutional Networks (VGG-Net), and a vanilla CNN. The comparison was conducted across the *Quebec Image Dataset*.

The experiment is organized as follows: First, we assess the performance of the individual constituent models (InceptionNet, ResNet, ViT, AlexNet, VGG-Net, and vanilla CNN). Next, we select the top four (ResNet, CNN, VGG-Net, and Inception) and top two (ResNet and CNN) performing constituent models to form MoE4 and MoE2 respectively. MoE6 contains all six models. Finally, the performance of MoE2, MoE4, and MoE6 are compared to each other and to the six constituent models.

We hypothesize that MoE2, which contains only the peak performing models, would outperform any individual constituent model, as well as the MoE4 and MoE6 variants. This is because it comprehensively integrates characteristics from top-performing architectures while also reducing complexity of the model. Furthermore, previous ensemble tree-based methods, such as Random Forest and XGBoost, have worked effectively in generating highly accurate predictions, supporting our presumption.

In conclusion, we discovered that MoE6 was the overall top performing model, with a peak accuracy of 93.13%. It outperformed any individual constituent model as well as its MoE2 and MoE4 counterparts, supporting the effectiveness of integrating the characteristics from a wide range of models.

Materials and Methods

Deep Learning Models

We present a succinct overview of the deep learning models examined in this study. The detailed information on each architecture is summarized in Table 1.

Vanilla CNN, (Zhang et al., 1988): It is tailored for three-dimensional data, primarily for image classification and recognition. CNNs consist of convolutional, pooling, and fully-connected (FC) layers. They begin with convolutional layers to detect basic features, may include additional con- volutions or pooling to reduce input dimensions and complexity, and conclude with FC layers that classify inputs by producing class likelihood probabilities. Convolutional layers use filters to convert images into arrays while pooling layers apply functions like max or average pooling (IBM, 2021).

AlexNet, (Krizhevsky et al., 2012): It is a prominent deep convolutional neural network, which gained renown for its exceptional performance in the ILSVRC-2012 competition. The network comprises 6 million parameters and 650,000 neurons, featuring five convolutional layers followed by pooling layers and three FC layers.

VGG-Net, (Simonyan and Zisserman, 2015): It achieved stellar performance in the 2014 ImageNet Challenge. The network is composed of a series of convolution networks followed by pooling layers, ending with three FC layers. Typically, VGG-Net consists of 11 to 19 layers with trainable weights.

MoE, (Jacobs et al., 1991): Unlike traditional deep learning models, MoE utilizes a cost-effective mapping

function to select the most effective experts for processing specific inputs, optimizing model capacity while minimizing computational overhead. It achieves efficiency by dynamically selecting relevant experts through a gating network that uses softmax functions to assign weights based on each expert's influence. Predictions are then generated by aggregating weighted outputs from these experts. During training, the model iteratively refines both the experts and gate parameters. Each expert specializes in distinct data regions, while the gate learns to allocate appropriate weights to maximize predictive accuracy (Prasann, 2024).

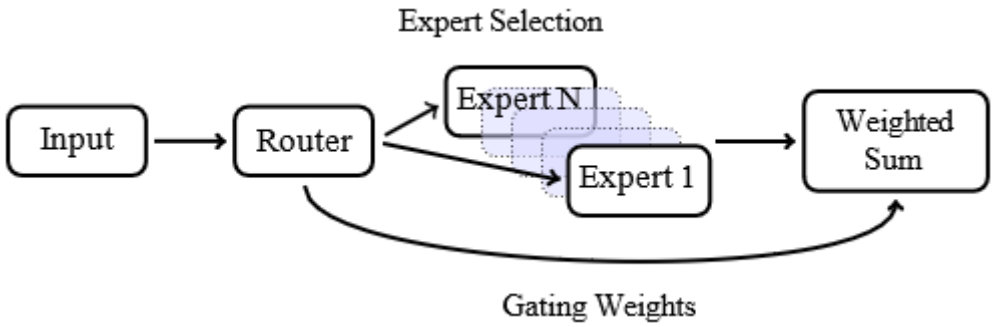

Figure 1. MoE partitions the learning model into sub-networks called experts, where each expert specializes in a distinct subset of input data. MoE architectures are distinguished by their capacity to significantly decrease computation costs and enhance training efficiency (IBM, 2024).

For this paper, there are three variants of the MoE: MoE2, MoE4, and MoE6. MoE2 contains ResNet and CNN as experts while MoE4 leverages ResNet, CNN, VGG-Net, and InceptionNet as predictors. MoE6 consists of all six models.

Table 1. Various architectures of CNN-based models.

CNN		AlexNet		VGG-Net	
Layer	Shape	Layer	Shape	Layer	Shape
Conv2D	(, 32, 32, 32)	Conv2D	(, 8, 8, 96)	Conv2D (×2)	(, 32, 32, 64)
MaxPooling2D	(, 16, 16, 32)	MaxPooling2D	(, 4, 4, 96)	MaxPooling2D	(, 16, 16, 64)
Dropout	(, 16, 16, 32)	Conv2D	(, 4, 4, 256)	Conv2D (×2)	(, 16, 16,128)
Flatten	(, 8,192)	MaxPooling2D	(, 2, 2, 256)	MaxPooling2D	(, 8, 8, 128)
Dense (×4)	(, 190)	Conv2D (×2)	(, 2, 2, 384)	Conv2D (×2)	(, 8, 8, 256)
Dense	(, 2)	Conv2D	(, 2, 2, 256)	MaxPooling2D	(, 4, 4, 256)
		MaxPooling2D	(, 1, 1, 256)	Conv2D (×3)	(, 4, 4, 512)
		Flatten	(, 256)	MaxPooling2D	(, 2, 2, 512)
		Dropout	(, 256)	Conv2D (×3)	(, 2, 2, 512)
		Dense	(, 4,096)	MaxPooling2D	(, 1, 1, 512)
		Dropout	(, 4,096)	Flatten	(, 512)
		Dense	(, 4,096)	Dense (×3)	(, 4,096)
		Dense	(, 2)	Dense	(, 2)

InceptionNet		ResNet	
Layer	Shape	Layer	Shape
Conv2D	(, 16, 16, 64)	ZeroPadding2D	(, 38, 38, 3)
MaxPooling2D	(, 8, 8, 64)	Conv2D	(, 19, 19, 64)

Conv2D	(, 8, 8, 64)	BatchNormalization	(, 19, 19, 64)
Conv2D	(, 8, 8, 192)	Activation	(, 19, 19, 64)
MaxPooling2D	(, 4, 4, 192)	MaxPooling2D	(, 10, 10, 64)
Inception Block	(, 4, 4, 256)	Convolutional Block	(, 5, 5, 256)
Inception Block	(, 4, 4, 480)	Identity Block (×2)	(, 5, 5, 256)
MaxPooling2D	(, 2, 2, 480)	Convolutional Block	(, 3, 3, 512)
Inception Block (×3)	(, 2, 2, 512)	Identity Block (×3)	(, 3, 3, 512)
Inception Block	(, 2, 2, 528)	Convolutional Block	(, 2, 2, 1024)
Inception Block	(, 2, 2, 832)	Identity Block (×5)	(, 2, 2, 1024)
MaxPooling2D	(, 1, 1, 832)	Convolutional Block	(, 1, 1, 2048)
Inception Block	(, 1, 1, 832)	Identity Block (×2)	(, 1, 1, 2048)
Inception Block	(, 1, 1, 1024)	AveragePooling2D	(, 1, 1, 2048)
AveragePooling2D	(, 1, 1, 1024)	Flatten	(, 2048)
Flatten	(, 1024)	Dense	(, 512)
Dropout	(, 1024)	Dense	(, 2)
Dense	(, 2)		

Inception, (Szegedy et al., 2014): The Inception Network (GoogLeNet) was responsible for the state of the art classification and detection in the ImageNet Large-Scale Visual Recognition Challenge 2014 (ILSVRC 2014). It employs a sparsely connected network architecture to avoid overfitting and reduce computational resources.

The Inception Block forms the basis for much of the model. This module is a combination of 1 × 1 Convolutional, 3 × 3 Convolutional, and 5 × 5 Convolutional layers, with their outputs concatenated to form the output for the next stage (Shaikh, 2023).

ResNet, (He et al., 2015): The Residual Network (ResNet) is an upgraded version of the VGGNet. It employs skip connections to minimize the vanishing gradient problem and avoid layers that reduce model performace.

Two types of skip connections are utilized for ResNet, namely the identity block and the convolutional block. Both blocks are composed of two 3 × 3 Convolutional and two Batch Normalization layers, as well as a ReLU function. However, the convolutional block performs a convolution followed by a batch normalization to the residue before adding it to the output (Shinde, 2021).

Vision Transformer (ViT), (Dosovitskiy et al., 2021): Vision Transformers have recently emerged as a competitive alternative to CNN-based architectures. It employs techniques from natural language processing to image classification and detection, achieving higher accuracy and requiring fewer computational resources. Its architecture is outlined in Figure 2.

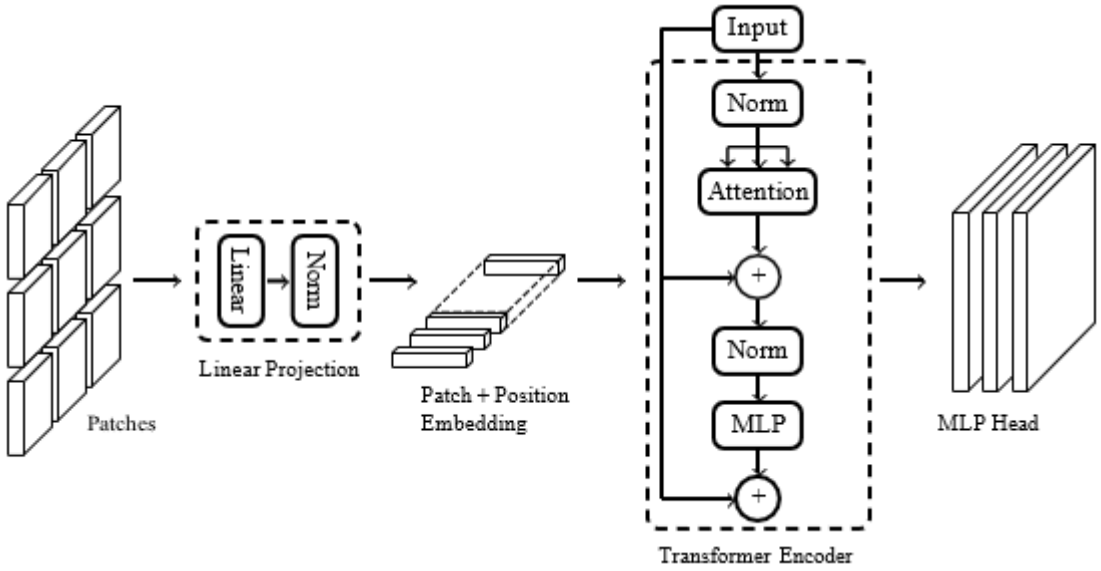

Figure 2. ViT divides an image into patches, which are flattened to become lower-dimensional linear embeddings. The sequence is inputted into a transformer encoder, and the multilayer perceptron determines the output.

Dataset Information and Processing

The dataset used for this research, referred to as *Quebec Image Dataset*, contains 42,850 satellite images, each sized at 350 × 350 pixels (Aaba, 2023). The images primarily cover Southern Quebec, Canada, spanning from April 30, 1972, to October 31, 2021. The dataset is balanced, comprising 22,710 instances of wildfire and 20,140 instances of no wildfire. The latitude and longitude coordi- nates for instances with fires (>0.01 acres burned) and without fires were sourced from (Government and Municipalities of Québec), using data collected from satellite images, aerial photographs, surveys, fire scar dating, and archival records.

To prepare the images for model processing, several steps were taken. First, the images were resized to 32x32 pixels to reduce file size. Next, they were converted to arrays and normalized from a [0,255] scale to a [0,1] scale for standardization (GeeksforGeeks, 2024). The dataset was then split into training, validation, and testing sets in a ratio of 7:1:2, with stratification and shuffling to maintain class proportions and ensure randomness. Finally, LabelEncoder converted categorical strings to binary labels (1 or 0), and to_categorical transformed these labels into binary vectors.

Model Training

The training subsets were fed into their respective models, with training loss monitored using the validation set. The vanilla CNN underwent hyperparameter tuning via Keras_Tuner to optimize performance, producing an effective CNN architecture. The proposed structures of Inception, ResNet, ViT, AlexNet, and VGG-Net were implemented. MoE6, MoE4, and MoE2 integrate subsets of the elements from Inception, ResNet, ViT, AlexNet, and VGG-Net architectures. To mitigate overfitting, EarlyStopping with a min_delta of 0.1 and patience of 20 was employed. The best performing weights during training were restored.

Various tools were used for extensive simulations: Google Colab as the primary web-based com- puting environment, NumPy and Pandas for array and dataframe management, Scikit-learn for data preprocessing and met- rics, Matplotlib for data visualization, and TensorFlow for building neural networks including CNN, AlexNet, VGG-Net, and MoE. We acknowledge (Prasann, 2024), (Patil, 2024), (Tripathi, 2024), (Varshney, 2020a), (Varshney, 2020b), Shaikh (2023), Shinde (2021), Elgazar (2023), and Salama (2021) for valuable code references that formed the

computational foundation for this work.

Results

Performance Metrics

We assess model performance using accuracy, precision, recall, F1 score, and ROC-AUC due to the task's nature. Accuracy measures the fraction of correct predictions, which is effective given that the dataset is balanced. Precision focuses on the accuracy of positive predictions, while recall evaluates the completeness of positive predictions. The F1 score, as the harmonic mean of precision and recall, captures aspects of both metrics.

- $\text{Accuracy} = \frac{\text{Number of correct predictions}}{\text{Total number of predictions}}$,
- $\text{Precision} = \frac{\text{True Positive}}{\text{True Positive+False Positive}}$,
- $\text{Recall} = \frac{\text{True Positive}}{\text{True Positive+False Negative}}$,
- $\text{F\_1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision+Recall}}$.

In addition, ROC-AUC assesses model discrimination by measuring the area under the ROC curve, where a score of 1.0 indicates perfect prediction, and scores below 0.5 indicate worse-than-random performance. This metric complements accuracy, especially in imbalanced datasets, offering insights into overall model performance (Brownlee, 2020).

Effective models should demonstrate high accuracy, precision, recall, F1 score, and ROC-AUC values, reflecting a thorough comprehension of the dataset and robust predictive capabilities.

Model Performance

Table 2. Individual model performance on Quebec Image Dataset.

Model	Accuracy	Precision	Recall	F_1	ROC-AUC
ResNet	0.9313	0.9049	0.9725	0.9375	0.9286
CNN	0.9308	0.9331	0.9366	0.9348	0.9304
VGG-Net	0.9235	0.9110	0.9483	0.9292	0.9219
Inception	0.9226	0.9129	0.9441	0.9282	0.9213
AlexNet	0.9062	0.8607	0.9819	0.9173	0.9014
ViT	0.8231	0.7619	0.9692	0.8531	0.8138

Table 3. MoE performance on Quebec Image Dataset.

Model	Accuracy	Precision	Recall	F_1	ROC-AUC
MoE6	0.9313	0.9250	0.9472	0.9359	0.9303
MoE2	0.9240	0.9199	0.9384	0.9290	0.9231
MoE4	0.9062	0.9206	0.9007	0.9105	0.9065

Discussion

As summarized in Table 2 and 3, all nine individual models exhibited exceptional performance on the *Quebect Image Dataset*, with most models achieving accuracies exceeding 90%.

Among the individual models, ResNet emerged as the highest performer at 93.13%. CNN, VGG-Net, and Inception closely followed with accuracies of 93.08%, 92.35%, and 92.26% respectively. Most models showed balanced performance across precision, recall, and F1 score metrics. However, ViT was the lowest performer by a large margin. Its recall was near perfect; however, its precision and ROC-AUC were lacking, perhaps underscoring a potential limitation of this method. Moreover, each ROC-AUC value closely mirrored accuracy, highlighting most models' robust discriminatory capabilities.

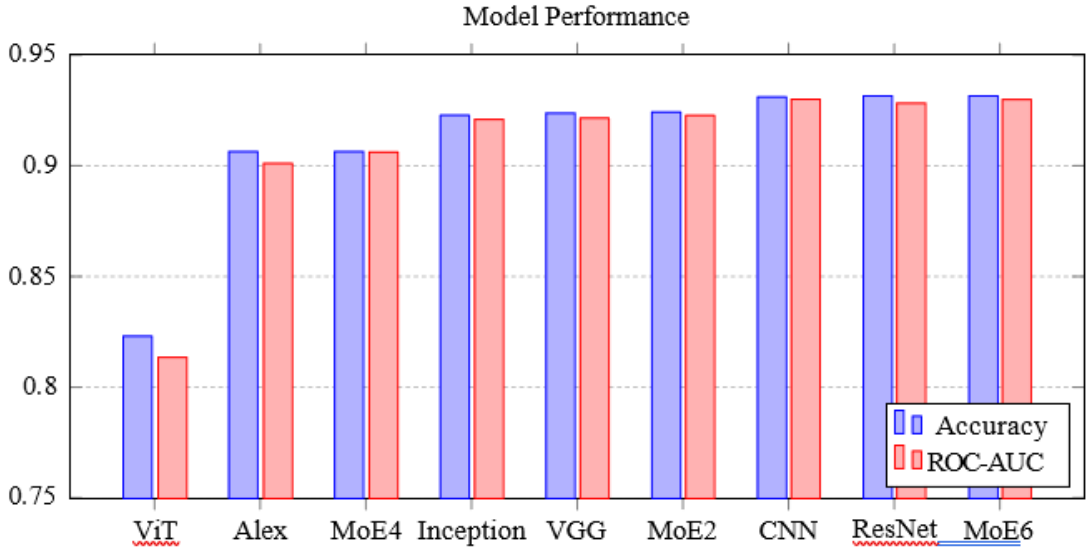

Figure 3. Individual and ensemble model results

In regards to the ensemble methods, MoE6 emerged as the highest performer at 93.13% accuracy. MoE2 followed closely at 92.40% accuracy, exhibiting a competitive ROC-AUC. Among the MoE architectures, MoE4 performed worse at 90.62% accuracy. However, despite its shortcomings, it demonstrated a balanced precision, recall, and F_1 score. Overall, MoE6 was the peak performer when compared to all methods, outscoring ResNet in ROC-AUC while matching its accuracy. As shown in Fig. 3, MoE methods were effective in general, placing 1st, 4th, and 7th when compared to the individual methods.

Considering the training and validation learning curves of the two top-performing models in Fig. 4, each model appeared to overfit both datasets. MoE6 exhibited substantial overfitting when compared to ResNet, wherein training accuracy approached 1 while validation accuracy stagnated close to 0.94. Despite this drawback, MoE6 was able to match and excel against ResNet.

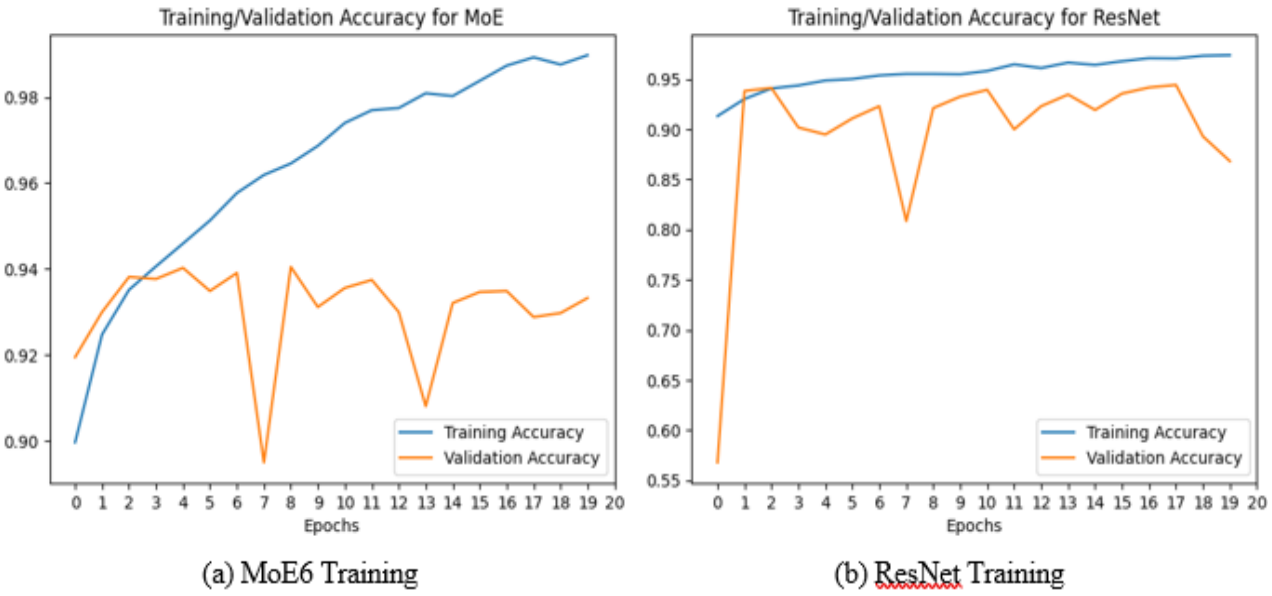

Figure 4. The comparison of learning curves across ResNet and MoE6.

Conclusion

In conclusion, this paper evaluated the effectiveness of a Mixture of Experts (MoE) approach for forest fire detection using satellite imagery, comparing it to state-of-the-art architectures like ViT, InceptionNet, ResNet, AlexNet, VGG-Net, and a baseline CNN. Furthermore, variants of the approach, notably MoE6, MoE4, and MoE2, were compared to reveal the effects of increasing the number of experts. All models performed strongly on the *Quebec Image Dataset*, with most models achieving accuracy above 90%. MoE6 stood out as the top performer with a peak 93.13% accuracy and ROC-AUC of 0.9303.

These findings showcase the potential of the ensemble approaches but also highlight the capabilities of deep CNN-based methods for wildfire detection, with the MoE approach showing competitive performance alongside established architectures. Future research directions could explore additional transformer-based models and integrate diverse types of satellite imagery. The models developed in this study hold promise for practical applications in forest fire monitoring, offering potential contributions to more effective wildfire management and mitigation strategies globally.

Limitations

Several limitations were observed during research. First, numerous models exhibited overfitting during the training process, notably CNN and the MoE variants. This issue can be circumvented by introducing data augmentation, regularization, or additional dropout layers to aid model learning. Furthermore, a majority of the models analyzed in the ablation study were CNN-based architectures, with the exception of ViT. To fully understand the capabilities of the MoE approach, additional transformer-based models, such as Swin Transformer, could be considered as experts.

Acknowledgments

The author sincerely appreciates Shourya Bose at UC Santa Cruz for the helpful discussions.

References

A. Aaba. Wildfire prediction dataset (satellite images), Feb 2023. URL https://www.kaggle.com/datasets/abdelghaniaaba/wildfire-prediction-dataset/data.

J. Brownlee. ROCcurvesandprecision-recallcurvesfor imbalancedclassi- fication, Sep 2020. URLhttps://machinelearningmastery.com/roc-curves-and-precision-recall-curves-for-imbalanced-classification/.

A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. 2021. URL https://arxiv.org/abs/2010. 11929.

E. Elgazar.Vision transformer (vit) + keras pretrained models, Feb 2023. URL https://www.kaggle.com/code/ebrahimelgazar/ vision-transformer-vit-keras-pretrained-models.

S. Gaur, J. S. Kumar, and S. Shukla. A comparative assessment of CNN-Sigmoid and CNN-SVM model for forest fire detection. In *2024 IEEE 9th Intl. Conf. for Convergence in Technology (I2CT)*, pages 1–6, 2024.

GeeksforGeeks. Normalize an image in OpenCV Python, May 2024. URL https://www.geeksforgeeks.org/normalize-an-image-in-opencv-python/.

Government and Municipalities of Québec. URL https://open.canada.ca/data/en/ dataset/9d8f219c-4df0-4481-926f-8a2a532ca003.

Government of Canada, Jul 2023. URL https://www.canada.ca/en/ public-health/services/emergency-preparedness-response/rapid-risk-assessments-public-health-professionals/ risk-profile-wildfires-2023.html.

K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL http://arxiv.org/abs/1512.03385.

IBM,Oct2021. URL https://www.ibm.com/topics/ convolutional-neural-networks.

IBM, Apr 2024. URL https://www.ibm.com/topics/mixture-of-experts.

R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1):79–87, 03 1991. ISSN 0899-7667.

A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolu- tional neural networks. In F. Pereira, C. Burges, L. Bottou, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper_files/paper/2012/ file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.

V. Patil. Wildfire prediction CNN, May 2024. URL https://www.kaggle.com/code/ vaishnavipatil4848/wildfire-prediction-cnn/notebook.

Prasann. Mixture of experts (MoE) explained, Jan 2024. URL https://www.kaggle.com/code/newtonbaba12345/mixture-of-experts-moe-explained.

H. C. Reis and V. Turk. Detection of forest fire using deep convolutional neural networks with transfer learning approach. *Applied Soft Computing*, 143:110362, 2023. ISSN 1568-4946. URL https://www.sciencedirect.com/science/article/pii/S1568494623003800.

K. Salama. Keras documentation: Image classification with vision transformer, 2021. URL https://keras.io/examples/vision/image_classification_with_ vision_transformer/.

V. E. Sathishkumar, J. Cho, M. Subramanian, and O. S. Naren. Forest fire and smoke detection using deep learning-based learning without forgetting. *Fire Ecology*, 19(1), Feb 2023.

J. Shaikh. Deep learning in the trenches: Understanding inception network from scratch, Nov 2023. URL https://www.analyticsvidhya.com/blog/2018/10/ understanding-inception-network-from-scratch/.

Y. Shinde.How to code your resnet from scratch in tensorflow?, Sep 2021. URL https://www.analyticsvidhya.com/blog/2021/08/ how-to-code-your-resnet-from-scratch-in-tensorflow/.

K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Intl. Conf. on Learning Representations*, 2015.

D. Spector.Will Quebec's forest fire season be as bad as it was last year?, May 2024. URL https://globalnews.ca/news/10497658/ quebec-forest-wild-fire-outlook-summer-2024/.

C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions, 2014. URL https://arxiv.org/abs/ 1409.4842.

M. Tripathi.Image processing using CNN: A beginners guide, May 2024. URL https://www.analyticsvidhya.com/blog/2021/06/ image-processing-using-cnn-a-beginners-guide/.

P. Varshney. Alexnet architecture: A complete guide, Jul 2020a. URL https://www.kaggle. com/code/blurredmachine/alexnet-architecture-a-complete-guide.

P. Varshney.VGGNet-16 architecture:A complete guide, Jul 2020b. URL https://www.kaggle.com/code/blurredmachine/ vggnet-16-architecture-a-complete-guide.

X. You, Z. Zheng, K. Yang, L. Yu, J. Liu, J. Chen, X. Lu, and S. Guo. A PSO-CNN-based deep learning model for predicting forest fire risk on a national scale. *Forests*, 15(1), 2024. ISSN 1999-4907. URL https://www.mdpi.com/1999-4907/15/1/86.

W. Zhang, J. Tanida, K. Itoh, and Y. Ichioka. Shift invariant pattern recognition neural network and its optical architecture. Proceedings of Annual Conference of the Japan Society of Applied Physics, 1988.