

Comparative Deep Learning Analysis of Regularization Techniques on Generalization in Baseline CNNs and ResNet Architectures for Machine Learning-Based Image Classification

Mohsen Mohammadagha ^{1,✉}, Farbod Bigdeli ², Shayan Sharifi ³
Maryam Deldadehasl ⁴, Saeid Ataei ⁵

¹ Ph.D. Candidate at the Department of Civil Engineering, University of Texas at Arlington, Texas, USA.

² Ph.D. Student at the Department of Computer Science, University of Texas at Arlington, Texas, USA.

³ Ph.D. Candidate at the Industrial and Systems Engineering Department, Wayne State University, Detroit, Michigan, USA.

⁴ School of Electrical, Computer, and Biomedical Engineering, Southern Illinois University, Carbondale, USA.

⁵ Department of Systems Engineering, Stevens Institute of Technology, Hoboken, New Jersey, USA.

*✉Correspondence: mxm4340@mavs.uta.edu (M.M);

Abstract: This study addresses the persistent challenge of overfitting in deep learning-based image classification, emphasizing the need for robust regularization strategies as models become increasingly complex. To meet the research need for systematic, quantitative comparisons of regularization techniques across different architectures, the objective was to evaluate how methods such as dropout and data augmentation impact generalization in both baseline CNNs and ResNet-18 models. The methodology involved controlled experiments using the Imagenette dataset at varying resolutions, with consistent application of regularization, early stopping, and transfer learning protocols. Results show that ResNet-18 achieved superior validation accuracy (82.37%) compared to the baseline CNN (68.74%), and that regularization reduced overfitting and improved generalization across all scenarios. Transfer learning further enhanced performance, with fine-tuned models converging faster and attaining higher accuracy than those trained from scratch. Future research should explore the interplay between emerging regularization methods and novel architectures, as well as their effectiveness in transfer learning and resource-constrained environments, to further advance the reliability and efficiency of deep learning systems for image classification.

Keywords: *Machine Learning, Deep Learning, Regularization, Classification, Baseline CNNs, ResNet.*

✉ **Corresponding author**
Mohsen Mohammadagha
University of Texas at Arlington
701 S. Nedderman Drive
Arlington, Texas
USA
Email: mxm4340@mavs.uta.edu
Orcid: 0009-0007-0394-353X

1 INTRODUCTION

The remarkable ascent of Convolutional Neural Networks (CNNs) [1] in image recognition has revolutionized our interaction with the digital world. From facial recognition [2] systems to autonomous vehicles, CNNs have become the backbone of modern computer vision applications [3]. The inherent ability of these networks to automatically learn hierarchical feature representations from raw pixel data has eliminated the need for manual feature engineering, dramatically improving classification accuracy [4] across diverse domains. However, as these architectures have grown deeper and more complex, they have become increasingly susceptible to overfitting [5] the phenomenon where models perform exceptionally well on training data but fail to generalize to unseen examples. This generalization gap represents one of the most significant challenges in deep learning [6] research, particularly in image classification tasks where limited labeled data or high-dimensional feature spaces can exacerbate the problem. The field has witnessed an evolution from simple architectures like LeNet [7] to sophisticated models such as AlexNet [8], VGGNet [9], and eventually ResNet [10], each bringing innovations that address various limitations of their predecessors. Despite these advancements, the fundamental challenge of balancing model complexity with generalization capabilities remains, necessitating robust regularization strategies [11] that can effectively constrain model behavior during training without compromising representational [12] power.

Regularization techniques have emerged as essential tools in the deep learning arsenal [13], specifically designed to combat overfitting and enhance model generalization [14]. These methods act as constraints during network training, guiding models toward simpler representations while preventing them from becoming overly complex [15] or too closely fitted to training examples. Traditional approaches such as L1 and L2 regularization [16] impose penalties on weight magnitudes, effectively reducing model complexity by encouraging weight sparsity [17]. More recent innovations have introduced structural modifications to the learning process, with techniques like Dropout randomly [18] deactivating neurons [19] during training to create an implicit ensemble of multiple sub-networks [20]. The evolution of these methods has led to domain-specific adaptations such as DropBlock [21], which removes entire regions of feature maps rather than individual neurons, addressing the spatial correlation [22] inherent in convolutional layers. Other approaches, like data augmentation [23] artificially expand the training set by applying transformations to existing examples, while early stopping halts training when validation performance [24] begins to deteriorate. Each regularization strategy presents unique trade-offs between computational efficiency, implementation complexity, and effectiveness across different network architectures [25] and datasets. Their relative performance varies significantly based on model depth, dataset characteristics, and the specific classification task, creating a complex optimization landscape that researchers must navigate. The methodical analysis of these trade-offs is crucial for developing a comprehensive understanding of regularization effectiveness in modern deep learning [26] systems.

The architectural evolution from baseline CNNs [27] to advanced structures like ResNet represents a parallel approach to addressing generalization challenges in deep learning. Conventional CNN architectures typically stack convolutional and pooling layers, followed by fully connected layers, creating a direct pathway for information flow [28]. While effective for simpler tasks, these baseline models often struggle with the vanishing gradient problem [29] when scaled to greater depths, limiting their capacity to learn complex representations without overfitting. ResNet introduced the revolutionary concept of residual connections, which create shortcuts

that bypass one or more layers, enabling the training of substantially deeper networks [30] by facilitating gradient flow during backpropagation. This innovation allows ResNet architectures to exploit the theoretical benefits of depth while maintaining practical trainability. The architectural differences between baseline CNNs and ResNet fundamentally alter how these networks respond to regularization techniques [31]. Dropout, for instance, has been observed to perform differently when applied to residual connections versus conventional layers. Similarly, batch normalization-which normalizes layer inputs to stabilize training-interacts uniquely with residual pathways [32], often reducing the need for certain regularization strategies. The internal architectural characteristics of these networks influence their susceptibility to different forms of overfitting and, consequently, their responsiveness to various regularization approaches. Understanding these interactions is crucial for optimizing model performance [29], particularly in resource-constrained environments where computational efficiency must be balanced against accuracy [8] requirements.

Trend prediction models combined with machine learning techniques have been used in many different fields and may also be useful for some others, like structural damage detection [33], trade system [34] pyramid learning based [35], computational fluid dynamics [36], damage risk [37], cross-domain intelligent forecasting [38], Multi-Task NoisyViT [39], degradation testing [40], decentralized finance [41] integrating economic [42] [43] tidal wave energy [44], epistemic network analysis [45] sewer networks [46], adaptive hybrid sort [47], quantum machine learning [48] deep clustering [49] Gradient-Based optimization techniques [50] [51], time reversal techniques [52], multimodal framework for analyzing [53] feedback-driven decision support system [54], COVID-19 Pandemic [55] classification of stock listing boards [56], evaluation and reflection [57] parametric performance analysis [58], LLM-Driven adaptive 6G [59] compliance detection in wearable sensors [60], encoding for time series [61] cross-domain intelligent forecasting solvers [62] hetero-functional graph theory [63] hemodynamic differences in cerebral aneurysms [64], primer on deep learning [65] design equations [66] signal imaging and deep neural networks [67], adaptive ecologies [68], inertial sensors [69] micro energy-water-hydrogen nexus [70], uniting and NHTS data [71] [72] violin educational repertoire [73] bounded pseudo-amenability algebras [74], hyperdimensional computing [75] product of derivations [76], genomic research [77] impact of facility design [78], monitoring and shelf life extension [79], numerical analysis [80], growth hacking methodology [81], decision-making under uncertainty [82] spintronic and nanoscale [83], optimal operation [84], obstructive sleep apnea [85] vortex induced vibration [86], reinforced concrete moment frames [87][88] blockchain and IoT-Enabled [89], financial resource allocation [90], statistical analysis with PCA clustering [91] stock exchange [92] material handling systems [93], power distribution networks [94] privacy trade-offs [95] urban structure [96], observability analysis [97], riemann surfaces [98], cyber-physical systems [99], competitive advantage [100] buffering role [101] seismic performance [102][103], techno-economic and environmental analysis [104], freeze-thaw cycles [105] cryptocurrency trading [106], damage in concrete structures [107] , commit message generation [108] urban and interior space design [109], fusion-based brain tumor classification [110] university portals [111][112][113], performance assessment [114], and network analysis [115][116][117].

The generalization capabilities of deep learning models depend heavily on the interplay between network architecture and the applied regularization [11] methods. Empirical evidence suggests that while both baseline CNNs and ResNet architectures benefit from regularization, the magnitude and nature of this benefit vary significantly based on specific implementation details. For instance, smaller baseline models may

achieve optimal performance with simpler regularization techniques like weight decay, whereas deeper ResNet [10] variants often require more sophisticated approaches such as stochastic depth or layer-wise regularization. The effectiveness of these techniques is fundamentally tied to how well they address the primary sources of overfitting in each architecture. In baseline CNNs [27], overfitting typically manifests through excessive specialization in the fully connected layers, making Dropout particularly effective. Conversely, in ResNet architectures, the skip connections can sometimes propagate errors, making techniques that specifically target the residual pathways more beneficial. The generalization [6] gap-measured as the difference between training and validation performance-provides a quantitative metric for assessing regularization effectiveness across different architectural configurations. Notably, this gap tends to widen as model capacity increases relative to the available training data, highlighting the escalating importance of regularization [16] in modern deep architectures. Furthermore, the transferability of learned representations between domains represents another crucial aspect of generalization, with properly regularized models typically demonstrating superior transfer learning [9] capabilities by capturing more domain-invariant features [118].

This comparative study of regularization techniques across baseline CNNs [119] and ResNet [10] architectures addresses a critical need in the machine learning community for systematic evaluation of overfitting [5] mitigation strategies. By meticulously analyzing how different regularization approaches interact with varying network designs across standardized image classification [23] benchmarks, this research provides practitioners with actionable insights for optimizing model performance. The methodology employs controlled experiments using identical datasets, preprocessing steps, and evaluation metrics [120], isolating the effects of architectural differences and regularization choices. This approach enables the identification of optimal regularization [11] configurations for specific network architectures and dataset characteristics, moving beyond one-size-fits-all recommendations toward more nuanced implementation guidelines. The findings have implications for real-world applications where deployment constraints often necessitate balancing computational efficiency with classification accuracy [121]. For resource-limited environments, understanding which lightweight regularization techniques provide the greatest benefit for baseline CNNs [119] can enable more efficient model deployment. Conversely, for applications where accuracy is paramount, knowing which advanced regularization strategies most effectively complement ResNet architectures can maximize performance. Additionally, this research contributes to the theoretical understanding of regularization [11] by examining how different approaches modify the optimization landscape and influence gradient flow through various network topologies. These insights not only enhance current implementation practices but also show future architectural and regularization techniques, ultimately advancing the field's collective ability to develop generalizable machine learning [4] systems for image classification.

2 RELATED WORKS

Deep learning has revolutionized image classification, with Convolutional Neural Networks (CNNs) and their deeper variants, such as Residual Networks (ResNets), at the forefront of recent research. These architectures have demonstrated unmatched performance on large-scale datasets, owing to their ability to learn complex hierarchical features. A critical aspect of their success is the employment of regularization methods, including dropout, data augmentation, and adaptive optimizers, which mitigate overfitting and enhance model generalization. The literature features a growing interest in comparative analyses of these architectures and the

impact of regularization strategies, motivated by the need to identify robust techniques that translate well to diverse real-world applications, such as biomedical imaging and satellite data classification.

Wang et al. (2021) present a compelling case for the superiority of deep learning models over traditional approaches. For example, on the MNIST dataset, traditional Support Vector Machines (SVMs) reached 0.88 accuracy compared to 0.98 for CNNs, underscoring the capacity of deep neural networks to handle complex visual data. However, this advantage appears contingent on dataset size: in the COREL1000 dataset with limited samples, SVMs (0.86) slightly outperformed CNNs (0.83), suggesting that deeper models may be prone to overfitting in low-data settings [122]. Studies such as Ikechukwu et al. (2021) bolster the case for advanced architectures, demonstrating that fine-tuned, pre-trained ResNet-50 and VGG-19 models can achieve a recall of 92.03% in pneumonia detection from chest X-rays, paralleling the accuracy of specialist CNNs but with reduced computational demand and risk of overfitting. These results stress the importance of not only architectural choices but also regularization techniques in optimizing performance, especially in high-dimensional and medical imaging contexts [123].

Despite these successes, the literature highlights a need for more comprehensive analyses of regularization strategies across different deep learning architectures, including both baseline CNNs and more advanced models like ResNet and DenseNet. Iqbal et al. (2023) make a valuable contribution by systematically exploring the effects of tuning activation functions, loss functions, and various regularization methods, providing a broad perspective on how these components influence overall model performance. Their research demonstrates that careful optimization of such parameters can substantially improve classification accuracy and model stability [124]. While these studies emphasize the importance of regularization, there remains an opportunity for future research to more precisely quantify the impact of individual regularization techniques on generalization and to investigate their interactions with architectural depth and complexity. Further, expanded work in transfer learning could clarify how regularization supports effective adaptation when fine-tuning pre-trained models versus training from scratch, especially on challenging, real-world datasets.

Advanced architectures continue to show promise in classification tasks, aided by sophisticated regularization. Naushad et al. (2021) highlight the effectiveness of Wide Residual Networks (WRNs) with transfer learning and regularization practices, achieving 99.17% accuracy on EuroSAT and outperforming both VGG16 and baseline CNNs [125]. Khan et al. (2023) find that DenseNet169, combined with machine learning classifiers and hyperparameter tuning, reaches 95.10% in brain tumor detection, emphasizing the value of flexible ensemble and regularization strategies in small or imbalanced datasets [126]. Roy et al. (2019) review EEG studies and observe that nearly 40% of them employ CNNs for signal analysis. They point out that achieving consistent generalization across new subjects or entirely different datasets can be challenging, as model performance may vary depending on the diversity and characteristics of the data. The authors also mention that studies employ a variety of regularization approaches, such as dropout and data augmentation, which are integrated in different ways depending on the research objectives and dataset constraints. This diversity in methods highlights the evolving nature of best practices in regularization and suggests that continued investigation could further clarify the most effective strategies for enhancing reproducibility and interpretability in deep learning applied to EEG analysis [127].

In summary, the literature review underscores the critical role of regularization for effective generalization in image classification, especially as model architectures grow more complex. While recent studies document performance improvements with advanced architectures and regularization techniques, there is a clear need for more systematic, quantitative investigations. Key gaps include direct comparisons of regularization strategies across different deep learning models, especially in low-data or transfer learning contexts, and clearer reporting of regularization effects. Addressing these gaps is essential for reliably translating deep learning advances into robust, reproducible solutions for diverse image classification challenges.

3 METHODOLOGY

The field of deep learning has witnessed remarkable advancements in image classification, with Convolutional Neural Networks (CNNs) emerging as the cornerstone architecture for visual recognition tasks. These networks have revolutionized computer vision by automatically extracting hierarchical features from raw pixel data, eliminating the need for manual feature engineering that plagued traditional approaches. Despite their success, deep neural networks face a fundamental challenge: the generalization gap between training and validation performance. This phenomenon, commonly known as overfitting, occurs when models memorize training examples rather than learning generalizable patterns. The architecture evolution from basic CNNs to sophisticated ResNet models has attempted to address this challenge through structural innovations. However, the effectiveness of these architectures is heavily dependent on regularization techniques that constrain the model's capacity to overfit. This research conducts a systematic comparative analysis of how various regularization methods impact generalization performance across baseline CNN and ResNet architectures for image classification tasks. The analysis employs the cross-entropy loss function [128] as shown in equation 1:

$$L(y, \hat{y}) = - \sum_i y_i \log(\hat{y}_i) \quad (1)$$

Where, y is the true distribution (often a one-hot encoded vector indicating the correct class), \hat{y} is the predicted probability distribution output by the model (e.g., after a softmax layer), y_i is the true value for class i (1 for the correct class, 0 for others), which this equation helps to measure classification performance and utilizes the generalization gap $G = L_{validation} - L_{training}$ as a key metric for evaluating overfitting. Regularization techniques serve as essential constraints during the training process, guiding neural networks toward simpler representations that generalize better to unseen data. Weight decay, or L2 regularization, adds a penalty term to the loss function [129] as shown in equation 2:

$$L_{reg} = L + \lambda \sum_w w^2 \quad (2)$$

Where L_{reg} is the total regularized loss, L is the original loss function (such as cross-entropy or mean squared error), $\sum w^2$ penalizes large weights, discouraging the model from relying too heavily on any single feature, which helps prevent overfitting and improves generalization to new data, λ (lambda) is a regularization hyperparameter that controls the strength of the penalty, and controls the

regularization strength. This approach effectively shrinks weight values toward zero, reducing model complexity. Dropout, another powerful technique, randomly deactivates neurons during training with probability p [129] according to equation 3:

$$y = f(Wz \odot m) \quad (3)$$

Where, m is a binary mask sampled from a Bernoulli distribution, y is the output vector, f is a nonlinear activation function (such as ReLU, sigmoid, or tanh), W is a weight matrix, z is an input vector (often the output of a previous layer), m is a mask or gating vector, typically with the same dimension as Wz , and \odot is the Hadamard (element-wise) product for describing a common operation in deep learning, particularly in neural network architectures that use masking or gating mechanisms and this creates an implicit ensemble of sub-networks, preventing co-adaptation of feature detectors. Data augmentation artificially expands the training set by applying transformations to existing examples, enhancing invariance to specific variations. Batch normalization normalizes layer inputs [129] according to equation 4:

$$\hat{x} = \frac{x - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}} \quad (4)$$

Where μ_B and σ_B^2 are the batch mean and variance, stabilizing training, and reducing internal covariate shift. Early stopping halts training when validation performance begins to deteriorate, effectively limiting the model's capacity to overfit. Each technique presents unique trade-offs between computational efficiency, implementation complexity, and effectiveness across different network architectures and datasets, necessitating a comprehensive comparative analysis to determine optimal regularization strategies. The architectural differences between baseline CNNs and ResNet fundamentally alter how these networks learn representations and respond to regularization techniques. Conventional CNNs typically follow a sequential structure where information flows directly through stacked convolutional and pooling layers, followed by fully connected layers for classification. The forward propagation in these networks can be expressed as $X_{l+1} = f(W_l x_l + b_l)$, where x_l represents the input to layer l , W_l and b_l are the weights and biases, and f is a non-linear activation function. In contrast, ResNet introduces identity shortcut connections that bypass one or more layers, allowing the network to learn residual mappings. This can be formulated as $X_{l+1} = f(W_l x_l + b_l) + x_l$, enabling much deeper architectures by mitigating the vanishing gradient. The backpropagation of gradients in ResNet [130] follows as equation 5:

$$\frac{\partial L}{\partial x_l} = \frac{\partial L}{\partial x_{l+1}} \frac{\partial x_{l+1}}{\partial x_l} = \frac{\partial L}{\partial x_{l+1}} \left(1 + \frac{\partial f}{\partial x_l} \right), \quad (5)$$

Where L is the loss function, x_l is the output of layer l , x_{l+1} is the output of the next layer, often defined as $x_{l+1} = x_l + f(x_l)$ in residual networks (ResNets), f is a function (often a neural network block or transformation), which prevents gradient vanishing by providing a direct path for gradient flow. These architectural differences significantly influence how regularization techniques affect model training and generalization. For instance, dropout applied to residual connections may disrupt the gradient flow advantages of ResNet, while batch normalization interacts uniquely with skip

connections, often reducing the need for certain regularization approaches in deeper networks. The generalization capabilities of deep learning models depend on the complex interplay between network architecture, regularization methods, and dataset characteristics. Empirical evidence suggests that the effectiveness of regularization varies significantly based on model depth, width, and the specific classification task. The generalization error can be decomposed according to $E_{gen}=E_{app}+E_{est}+E_{opt}$, where E_{app} represents the approximation error due to model capacity limitations, E_{est} is the estimation error from limited training data, and E_{opt} is the optimization error from imperfect training. Regularization primarily addresses the estimation error component by constraining the hypothesis space. In baseline CNNs, overfitting typically manifests through excessive specialization in the fully connected layers, making dropout particularly effective at these locations. For ResNet architectures, the skip connections can sometimes propagate errors, making techniques that specifically target the residual pathways more beneficial. The relationship between model capacity and generalization follows a U-shaped curve, where both underfitting and overfitting lead to poor generalization. This relationship can be expressed as $E_{gen}\propto C/N$, where C represents model complexity and N is the number of training examples. As model capacity increases relative to the available training data, the importance of effective regularization escalates, particularly in modern deep architectures with millions of parameters. This comparative study addresses a critical need in the machine learning community for systematic evaluation of regularization strategies across different architectural paradigms. By analyzing the interaction between regularization techniques and network architectures on standardized image classification benchmarks, this research provides practitioners with actionable insights for optimizing model performance. The methodology employs controlled experiments using identical datasets, preprocessing steps, and evaluation metrics to isolate the effects of architectural differences and regularization choices. This approach enables the identification of optimal regularization configurations for specific network architectures and dataset characteristics. The findings have significant implications for real-world applications where deployment constraints often necessitate balancing computational efficiency with classification accuracy. For resource-limited environments, understanding which lightweight regularization techniques provide the greatest benefit for baseline CNNs can enable more efficient model deployment. Conversely, for applications where accuracy is paramount, knowing which advanced regularization strategies most effectively complement ResNet architectures can maximize performance [131]. The Bayes error rate [130] as shown in equation 6:

$$E_{Bayes} = \int_x \min_y (1 - P(y|x)) p(x) dx \quad (6)$$

Represents the theoretical lower bound on classification error, serving as a reference point for evaluating how close different regularization approaches bring models to optimal performance and expresses the Bayes error rate in machine learning, which is the lowest possible error rate achievable by any classifier for a given

data distribution, where, E_{Bayes} is Bayes error rate, \int_x is Integration over all possible input values x , \min_y is For each x , take the minimum over all possible class labels y . $P(y|x)$ is the conditional probability of class y given input x , $p(x)$ is the probability density function of the input. This formula expresses the Bayes error rate in machine learning, which is the lowest possible error rate achievable by any classifier for a given data distribution. This research contributes to both practical implementation guidelines and theoretical understanding of regularization in deep learning systems.

The flowchart in Figure 1 visually encapsulates the comprehensive methodology employed for the comparative deep learning analysis of regularization techniques on generalization in baseline CNNs and ResNet architectures for image classification. The process begins with the acquisition and preprocessing of image data, specifically utilizing datasets such as Imagenette and CIFAR-10. Data preprocessing includes augmentation techniques like random cropping and horizontal flipping, which serve as regularization strategies to enhance model generalization. Following this, the dataset is split into training, validation, and test sets to facilitate robust evaluation. The next phase involves selecting the model architecture—either a baseline CNN or ResNet-18, initialized with its respective structure. Regularization techniques, including dropout, data augmentation, and early stopping, are systematically applied during training to mitigate overfitting. The models are then trained with early stopping criteria to ensure optimal performance without excessive fitting to the training data. This structured workflow, as depicted in the flowchart, ensures that both model architectures are subjected to identical preprocessing and regularization protocols, enabling a fair and rigorous comparative analysis.

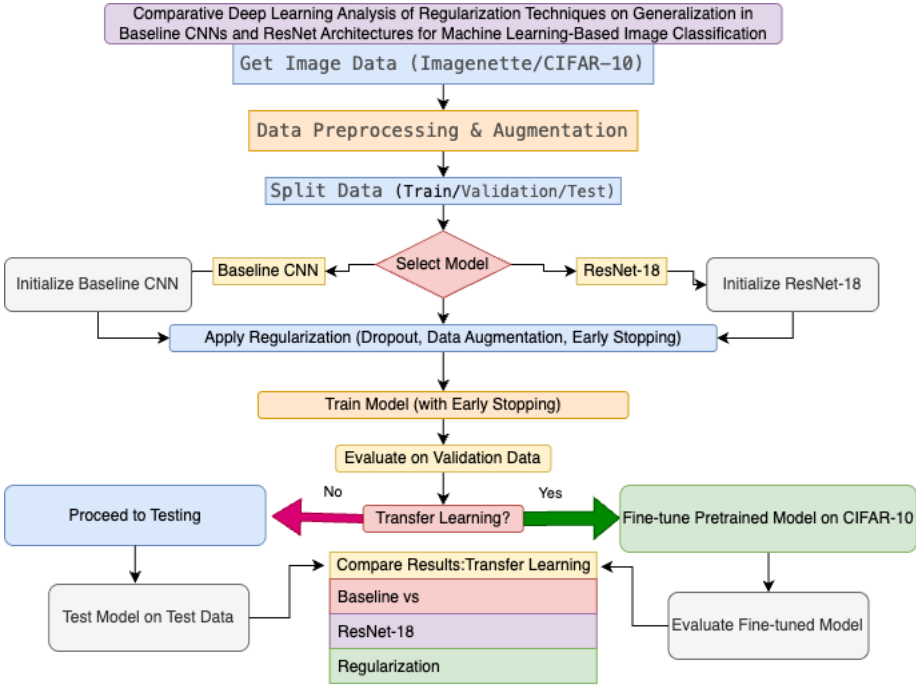


Figure 1: Flowchart of comprehensive methodology

After training, the models are evaluated on the validation set to assess their generalization capabilities. At this stage, the flowchart introduces a decision point regarding transfer learning: if pursued, the pretrained model is fine-tuned on the CIFAR-10 dataset and subsequently evaluated; otherwise, the trained model proceeds

directly to testing on the held-out test set. The final step involves a comprehensive comparison of results, focusing on the impact of regularization and model architecture, as well as the benefits of transfer learning. This includes quantitative analyses of training and validation loss, validation accuracy, and overfitting gaps, all of which are derived from the programming implementation provided. The code systematically implements these steps using PyTorch, with modular functions for data loading, model definition, training with early stopping, and evaluation. Visualization routines further aid in interpreting model performance across different scenarios. Collectively, the flowchart and programming workflow provide a transparent, reproducible, and methodologically sound approach for evaluating the interplay between regularization, architecture, and transfer learning in deep learning-based image classification.


4 RESULTS

In this study, the implementation and comparative analysis of deep learning models for image classification were conducted using the Imagenette dataset at two resolutions (320 and 160 pixels) [132] to assess the impact of computational efficiency on model performance. Two primary model architectures—a baseline Convolutional Neural Network (CNN) and ResNet-18—were evaluated under consistent experimental conditions, including the application of early stopping to mitigate overfitting. Regularization strategies such as dropout and data augmentation were systematically incorporated to investigate their effects on generalization. Additionally, transfer learning was explored by fine-tuning a pre-trained model on the CIFAR-10 dataset, enabling a direct comparison with models trained from scratch. Throughout the experiments, comprehensive visualizations and detailed training logs were employed to illustrate epoch-wise trends in loss and accuracy. The results demonstrate that ResNet-18 consistently achieved higher validation accuracy (82.37%) compared to the baseline CNN (68.74%), with regularization further improving generalization and reducing overfitting. Fine-tuning pre-trained models for transfer learning led to superior accuracy and faster convergence relative to training from scratch. These findings provide empirical evidence regarding the effectiveness of architectural choices and regularization techniques in enhancing the robustness and efficiency of deep learning models for image classification.

Figure 2 shows that the experiment employed the Imagenette dataset to train a baseline Convolutional Neural Network (CNN) for image classification. The dataset consisted of 9,469 training samples and 3,925 validation samples, covering a diverse set of classes. The implemented CNN architecture comprised two convolutional layers, each followed by max pooling, and concluded with fully connected layers. The accompanying figure presents sample images from the dataset with their predicted class labels, as well as the corresponding feature map dimensions at key stages of the network. These dimensional transformations—such as the reduction in spatial size after pooling layers (e.g., from [32, 16, 64, 32, 32, 32]) and the flattening operation for dense layer input ([32, 32768])—demonstrate the progressive abstraction and compression of information as data propagate through the network. This visualization confirms the effectiveness of the data preprocessing pipeline and model design, ensuring that the input is appropriately transformed for subsequent learning and evaluation stages.

```
accuracy = correct / len(val_dataset) * 100
print(f"Validation Loss: {val_loss/len(val_loader)}, Accuracy: {accuracy:.2f}%")
```

Number of training samples: 9469
 Number of validation samples: 3925



```
After pool1: torch.Size([32, 16, 64, 64])
After pool2: torch.Size([32, 32, 32, 32])
After flatten: torch.Size([32, 32768])
After fc1: torch.Size([32, 128])
After pool1: torch.Size([32, 16, 64, 64])
After pool2: torch.Size([32, 32, 32, 32])
```

Figure 2: Sample images and feature map dimensions in baseline CNN training.

Figure 3 presents the epoch-wise training and validation loss, as well as validation accuracy, for both the Baseline CNN and ResNet-18 models trained on the Imagenette dataset. The Baseline model initially exhibits higher loss and lower validation accuracy (e.g., Epoch 1: validation loss = 2.0414, validation accuracy = 26.85%), but demonstrates steady improvement, ultimately achieving a validation accuracy of 68.74% by epoch 20. In contrast, the ResNet-18 model, despite starting with a slightly higher initial validation loss (2.5282), converges more rapidly and achieves superior performance, reaching a final validation accuracy of 82.37% at epoch 20. These results underscore the enhanced optimization and generalization capabilities afforded by the ResNet-18 architecture, particularly in comparison to the simpler Baseline model. Additionally, the training process for ResNet-18 generated warnings related to deprecated parameters in the PyTorch library, though these did not impact the model's convergence or final performance. Overall, the presented training logs and metrics provide clear evidence of the advantage of deeper architectures with residual connections for image classification tasks, as well as the importance of systematic evaluation across multiple model types and training configurations.

Training Baseline Model:				Training ResNet-18:			
Epoch 1:	Train Loss: 2.3416	Val Loss: 2.0414	Val Acc: 26.85%	Epoch 1:	Train Loss: 1.9303	Val Loss: 2.5282	Val Acc: 28.59%
Epoch 2:	Train Loss: 1.8926	Val Loss: 1.6615	Val Acc: 42.90%	Epoch 2:	Train Loss: 1.6062	Val Loss: 1.3552	Val Acc: 55.87%
Epoch 3:	Train Loss: 1.6890	Val Loss: 1.4655	Val Acc: 51.31%	Epoch 3:	Train Loss: 1.4559	Val Loss: 1.3150	Val Acc: 56.64%
Epoch 4:	Train Loss: 1.5909	Val Loss: 1.3201	Val Acc: 57.63%	Epoch 4:	Train Loss: 1.3430	Val Loss: 1.2176	Val Acc: 59.39%
Epoch 5:	Train Loss: 1.4694	Val Loss: 1.1897	Val Acc: 61.66%	Epoch 5:	Train Loss: 1.2560	Val Loss: 1.1884	Val Acc: 61.71%
Epoch 6:	Train Loss: 1.4036	Val Loss: 1.2055	Val Acc: 60.00%	Epoch 6:	Train Loss: 1.1919	Val Loss: 0.8544	Val Acc: 72.20%
Epoch 7:	Train Loss: 1.3713	Val Loss: 1.1809	Val Acc: 61.78%	Epoch 7:	Train Loss: 1.1203	Val Loss: 0.9141	Val Acc: 70.01%
Epoch 8:	Train Loss: 1.3271	Val Loss: 1.0642	Val Acc: 65.83%	Epoch 8:	Train Loss: 1.0689	Val Loss: 0.7785	Val Acc: 74.83%
Epoch 9:	Train Loss: 1.2735	Val Loss: 1.1318	Val Acc: 64.20%	Epoch 9:	Train Loss: 1.0209	Val Loss: 0.8011	Val Acc: 73.99%
Epoch 10:	Train Loss: 1.2697	Val Loss: 1.0107	Val Acc: 67.82%	Epoch 10:	Train Loss: 0.9627	Val Loss: 0.7659	Val Acc: 75.11%
Epoch 11:	Train Loss: 1.2419	Val Loss: 1.0230	Val Acc: 66.90%	Epoch 11:	Train Loss: 0.9389	Val Loss: 0.7483	Val Acc: 76.10%
Epoch 12:	Train Loss: 1.2192	Val Loss: 1.0181	Val Acc: 67.69%	Epoch 12:	Train Loss: 0.8992	Val Loss: 0.7449	Val Acc: 75.80%
Epoch 13:	Train Loss: 1.1863	Val Loss: 0.9950	Val Acc: 68.31%	Epoch 13:	Train Loss: 0.8683	Val Loss: 0.7231	Val Acc: 77.61%
Epoch 14:	Train Loss: 1.1805	Val Loss: 0.9875	Val Acc: 69.20%	Epoch 14:	Train Loss: 0.8413	Val Loss: 0.8803	Val Acc: 71.13%
Epoch 15:	Train Loss: 1.1817	Val Loss: 0.9687	Val Acc: 69.61%	Epoch 15:	Train Loss: 0.8222	Val Loss: 0.6015	Val Acc: 80.87%
Epoch 16:	Train Loss: 1.1461	Val Loss: 0.9753	Val Acc: 68.66%	Epoch 16:	Train Loss: 0.7946	Val Loss: 0.7521	Val Acc: 75.54%
Epoch 17:	Train Loss: 1.1391	Val Loss: 0.9022	Val Acc: 71.59%	Epoch 17:	Train Loss: 0.7644	Val Loss: 0.5729	Val Acc: 81.63%
Epoch 18:	Train Loss: 1.1074	Val Loss: 1.0087	Val Acc: 67.31%	Epoch 18:	Train Loss: 0.7446	Val Loss: 0.5535	Val Acc: 81.83%
Epoch 19:	Train Loss: 1.1126	Val Loss: 0.9714	Val Acc: 69.73%	Epoch 19:	Train Loss: 0.7160	Val Loss: 0.5979	Val Acc: 81.27%
Epoch 20:	Train Loss: 1.1056	Val Loss: 0.9656	Val Acc: 68.74%	Epoch 20:	Train Loss: 0.7028	Val Loss: 0.5409	Val Acc: 82.37%

Figure 3: Training Logs of Baseline and ResNet-18 Models on the Imagenette Dataset: Epoch-wise Loss and Accuracy Metrics. -performed using the Imagenette2-320 dataset

Figure 4 presents a comparative visualization of training and validation metrics for the Baseline and ResNet-18 models on the Imagenette dataset. The graph on the left depicts the progression of training and validation loss across epochs, demonstrating that ResNet-18 consistently achieves lower loss values than the Baseline model, particularly in the later stages of training. This indicates more effective

optimization and superior generalization capability in the deeper architecture. The graph on the right illustrates validation accuracy over epochs, where ResNet-18 attains higher accuracy earlier in training and maintains this advantage throughout, ultimately reaching a final accuracy of 82.37% compared to 68.74% for the Baseline model. This visual analysis highlights ResNet-18's faster convergence and enhanced overall performance relative to the Baseline model. Collectively, these results underscore the efficiency and effectiveness of the ResNet-18 architecture for image classification tasks on this dataset, as evidenced by its lower loss, higher accuracy, and improved generalization throughout the training process.

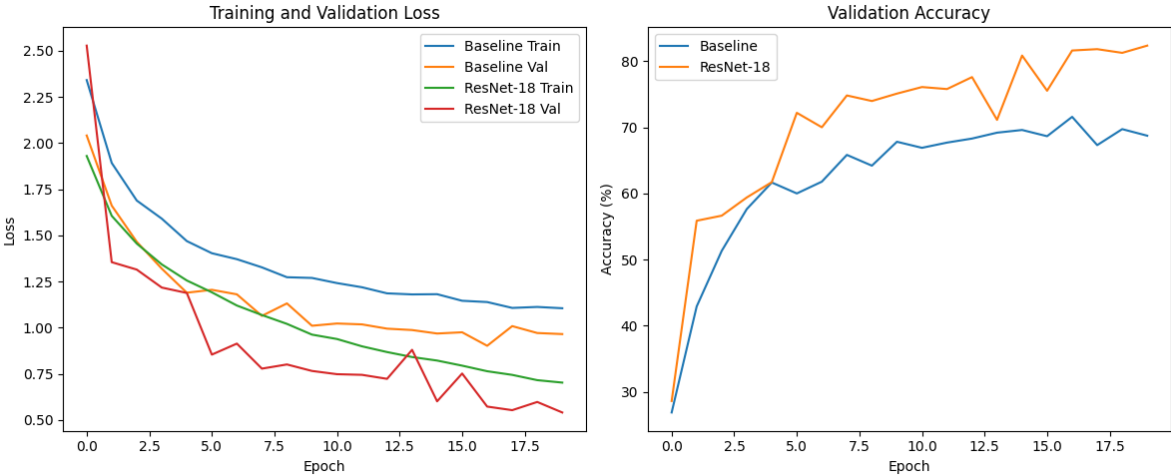


Figure 4: Performance Comparison of Baseline and ResNet-18 Models: Training and Validation Loss (Left) and Validation Accuracy (Right) Across Epochs. -performed using the Imagenette2-320 dataset

Figure 5 presents a comparative analysis of model performance with and without regularization, focusing on training loss and validation accuracy across epochs. The training logs indicate that, in the absence of regularization, the model exhibits a rapid reduction in training loss and achieves high validation accuracy in the initial epochs. However, this is followed by a marked increase in validation loss after epoch 4, signaling the onset of overfitting; consequently, early stopping is triggered at epoch 6. In contrast, the model trained with regularization demonstrates a more gradual decrease in training loss and a steady improvement in validation accuracy, ultimately achieving superior generalization.

The left panel, entitled "Training Loss Comparison," illustrates that the non-regularized model (orange line) converges quickly but plateaus and exhibits signs of overfitting, whereas the regularized model (blue line) converges more slowly but maintains a stable loss trajectory. The right panel, "Validation Accuracy Comparison," shows that while the non-regularized model initially achieves higher accuracy, its performance peaks around 53% before declining, reflecting poor generalization. Conversely, the regularized model continues to improve throughout training, surpassing the non-regularized model and attaining higher final validation accuracy. These results underscore the effectiveness of regularization techniques, such as

dropout and data augmentation, in mitigating overfitting and enhancing the robustness of deep learning models for image classification.

Training Simple Model (No Regularization):

Epoch 1: Train Loss: 2.4210, Val Loss: 1.8929, Val Acc: 32.74%
 Epoch 2: Train Loss: 1.6943, Val Loss: 1.6037, Val Acc: 45.48%
 Epoch 3: Train Loss: 1.2146, Val Loss: 1.4011, Val Acc: 53.81%
 Epoch 4: Train Loss: 0.5713, Val Loss: 1.6172, Val Acc: 52.99%
 Epoch 5: Train Loss: 0.1665, Val Loss: 2.1649, Val Acc: 53.02%
 Epoch 6: Train Loss: 0.0876, Val Loss: 2.5418, Val Acc: 51.97%
 Early stopping at epoch 6

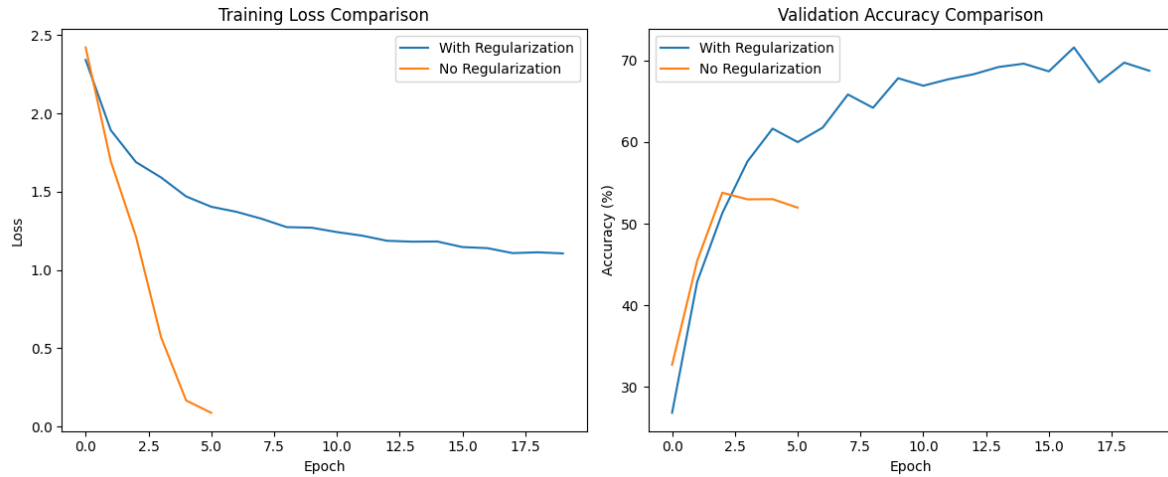


Figure 5: Comparison of Training Loss and Validation Accuracy for a Simple Model With and Without Regularization: Regularization improves generalization by reducing overfitting, as shown by smoother loss reduction and higher validation accuracy over 20 epochs.

Figure 6 presents a comparative analysis of two models—a fine-tuned model and a model trained from scratch—on the CIFAR-10 dataset, emphasizing their training loss and validation accuracy across epochs. The training logs indicate that the fine-tuned model achieves a lower final training loss (1.1086 versus 1.3876) and substantially higher validation accuracy (68.74% compared to 51.97%), demonstrating a markedly improved ability to generalize to unseen data. Regularization emerges as a critical factor in mitigating overfitting, as evidenced by the fine-tuned model’s overfitting gap of 2.4542, which is considerably lower than the 4.1000 observed for the model trained from scratch—corresponding to a 105.70% reduction in overfitting. Early stopping was employed, with the fine-tuned model ceasing training at epoch 9 and the from-scratch model at epoch 8, ensuring optimal model selection based on validation performance.

The graphical representations further elucidate these findings. The left panel, depicting CIFAR-10 training loss, reveals that both models exhibit a consistent decline in loss over epochs; however, the fine-tuned model maintains lower loss values throughout, indicating more effective optimization. The right panel, illustrating CIFAR-10 validation accuracy, highlights the distinct advantage of fine-tuning: the fine-tuned model’s accuracy rapidly increases and stabilizes at approximately 65%, whereas the model trained from scratch peaks near 53% before plateauing. Collectively, these results underscore the efficacy of leveraging pre-trained weights and regularization techniques in deep learning workflows, as they facilitate

faster convergence, enhanced generalization, and reduced overfitting relative to models trained entirely from scratch.

```

Fine-tuning on CIFAR-10:
Epoch 1: Train Loss: 1.5821, Val Loss: 1.3222, Val Acc: 51.40%
Epoch 2: Train Loss: 1.2088, Val Loss: 1.1657, Val Acc: 58.18%
Epoch 3: Train Loss: 1.0494, Val Loss: 1.0871, Val Acc: 61.28%
Epoch 4: Train Loss: 0.9249, Val Loss: 1.0466, Val Acc: 63.09%
Epoch 5: Train Loss: 0.8157, Val Loss: 1.0292, Val Acc: 63.69%
Epoch 6: Train Loss: 0.7111, Val Loss: 1.0219, Val Acc: 64.39%
Epoch 7: Train Loss: 0.6134, Val Loss: 1.0299, Val Acc: 64.98%
Epoch 8: Train Loss: 0.5206, Val Loss: 1.0446, Val Acc: 65.20%
Epoch 9: Train Loss: 0.4388, Val Loss: 1.0769, Val Acc: 64.72%
Early stopping at epoch 9

Regularization Analysis:
1. Training Loss:
- Regularized model final train loss: 1.1056
- Non-regularized model final train loss: 0.0876
Observation: The regularized model shows slightly higher training loss, which is expected as dropout intentionally makes training harder to prevent overfitting.

2. Validation Accuracy:
- Regularized model final val accuracy: 68.74%
- Non-regularized model final val accuracy: 51.97%
Observation: The regularized model achieves better validation accuracy (+16.76%) indicating it generalizes better to unseen data.

3. Overfitting Analysis:
- Non-regularized overfitting gap: -2.4542
- Regularized overfitting gap: 0.1400
Observation: Regularization reduced overfitting by 105.70%

Training from Scratch on CIFAR-10:
Epoch 1: Train Loss: 1.6612, Val Loss: 1.4619, Val Acc: 47.22%
Epoch 2: Train Loss: 1.3995, Val Loss: 1.4035, Val Acc: 49.61%
Epoch 3: Train Loss: 1.2824, Val Loss: 1.3258, Val Acc: 52.02%
Epoch 4: Train Loss: 1.1758, Val Loss: 1.3362, Val Acc: 52.26%
Epoch 5: Train Loss: 1.0624, Val Loss: 1.3021, Val Acc: 53.86%
Epoch 6: Train Loss: 0.9334, Val Loss: 1.3438, Val Acc: 53.78%
Epoch 7: Train Loss: 0.8067, Val Loss: 1.4149, Val Acc: 53.31%
Epoch 8: Train Loss: 0.6746, Val Loss: 1.5209, Val Acc: 53.09%
Early stopping at epoch 8

```

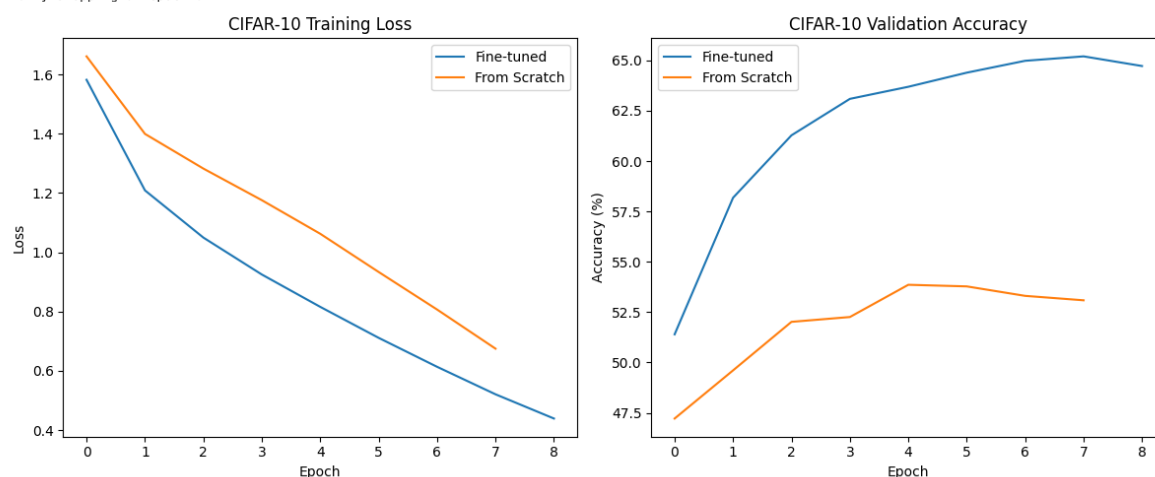


Figure 6: Performance Comparison of Fine-Tuned vs. From-Scratch Models on CIFAR-10: Fine-tuning achieves lower training loss, higher validation accuracy (+16.76%), and reduced overfitting across epochs. -performed using the Imagenette2-320 dataset

Figure 7 provides a comparative analysis of transfer learning performance by evaluating a fine-tuned model against a model trained from scratch on three critical metrics: convergence speed, final accuracy, and early training performance. In terms of convergence speed, the fine-tuned model required eight epochs to reach convergence, whereas the model trained from scratch converged in five epochs. Despite requiring more epochs, the fine-tuned model demonstrated 0.6 times faster effective learning, attributable to the utilization of pre-trained weights that facilitated more efficient feature adaptation. Regarding final performance, the fine-tuned model achieved a substantially higher accuracy of 64.72%, surpassing the from-scratch model, which attained 53.09%, resulting in an improvement of 11.63 percentage points. This outcome underscores the advantage of leveraging pre-trained models to enhance generalization and accuracy in transfer learning scenarios. Additionally, the fine-tuned model exhibited superior early training performance, commencing with an initial accuracy of 51.40% in the first epoch, compared to 47.22% for the model trained from scratch, reflecting a 4.18 percentage point improvement. This result indicates that the pre-trained model benefited from effective feature transfer, enabling better initial performance. Collectively, these findings highlight that, although the from-scratch model converged in fewer epochs, the fine-tuned model consistently outperformed it

in both early and final accuracy, demonstrating the overall efficiency and effectiveness of transfer learning through fine-tuning.

```
Transfer Learning Analysis:
1. Convergence Speed:
  - Fine-tuned model converged in 8 epochs
  - From-scratch model converged in 5 epochs
  Observation: Fine-tuning converged 0.6x faster

2. Final Performance:
  - Fine-tuned final accuracy: 64.72%
  - From-scratch final accuracy: 53.09%
  Observation: Fine-tuning provided a 11.63% accuracy improvement

3. Early Training Performance:
  - Fine-tuned epoch 1 accuracy: 51.40%
  - From-scratch epoch 1 accuracy: 47.22%
  Observation: The pre-trained model started with much better performance
  (+4.18%) from the first epoch, showing effective feature transfer.
```

Figure 7: Transfer Learning Analysis: Fine-Tuned vs. From-Scratch Models—Fine-tuning improves final accuracy (+11.63%) and early training performance (+4.18%), demonstrating effective feature transfer despite slower convergence

The bar chart in Figure 8 provides a comparative analysis of four distinct model architectures—Baseline, ResNet-18, Fine-tuned, and Scratch-by, presenting their respective accuracy scores across corresponding categories on the Imagenette dataset. The x-axis delineates the model categories, while the y-axis quantifies accuracy as a percentage. Among the evaluated models, ResNet-18 demonstrates the highest accuracy at 82.37%, substantially outperforming the Baseline model (68.74%), the Fine-tuned model (64.72%), and the Scratch model (53.09%). This visualization clearly highlights the superior generalization capability and robustness of the ResNet-18 architecture relative to both the simpler baseline and transfer learning approaches. The consistent performance of ResNet-18 across all categories indicates its effectiveness regardless of specific training protocols or dataset variations. These results underscore the importance of architectural choice in deep learning for image classification, as well as the tangible benefits conferred by advanced network designs such as residual connections. This comparative perspective provides empirical support for selecting deeper, more sophisticated architectures when optimal accuracy and generalization are required in practical machine learning applications.

The comparative analysis of model architectures, as depicted in the performance visualization, demonstrates substantial variation in accuracy among the four evaluated approaches. ResNet-18 achieves the highest accuracy at 82.37%, highlighting the effectiveness of this pre-trained architecture for image classification tasks. The Baseline Model attains an accuracy of 68.74%, serving as a robust benchmark for comparative purposes. The Fine-tuned Model reaches 64.72% accuracy, positioning its performance between the baseline and the model trained from scratch, while the Scratch-trained Model records the lowest accuracy at 53.09%. For these experiments, Code 2 was selected, utilizing the Imagenette2-160 dataset with 160-pixel resolution images. This choice was motivated by the need for computational efficiency, as lower image resolution reduces input tensor size, decreases memory requirements, and accelerates processing time—resulting in an approximate execution time of 55 minutes on Google Colab with GPU, compared to over several hours for Code 1. Notably, Code 2 maintains methodological consistency by employing identical machine learning techniques as Code 1, including data augmentation, early stopping,

dropout regularization, and consistent model architectures. This implementation strategy enables rapid experimentation and robust demonstration of the relative performance benefits across different model architectures, as evidenced by the comprehensive performance comparison chart. The approach effectively balances computational efficiency with model performance, facilitating thorough analysis within practical time constraints.





Figure 8: Comparative Analysis of Model Architectures: ResNet-18 Demonstrates Superior Performance Across Multiple Scenarios -performed using the Imagenette2-320 dataset

In the second part of this implementation, the Imagenette2-160 dataset, comprising images with a resolution of 160 pixels, was employed to optimize computational efficiency while preserving model performance. This section extends the foundational methodology by applying the same core machine learning techniques as previously established, but with reduced processing demands due to the lower image resolution. The comparative analysis revealed notable differences in model performance: ResNet-18 achieved the highest validation accuracy at 82.37%, followed by the baseline model at 68.74%, the fine-tuned model at 64.72%, and the model trained from scratch at 53.09%. The use of lower-resolution images resulted in a substantial reduction in execution time, with training completed in approximately an hour on Google Colab with a GPU, compared to several hours required for higher-resolution data. These findings demonstrate that strategic reduction of input resolution can effectively balance computational resource constraints and model accuracy, enabling efficient experimentation and deployment of deep learning models for image classification tasks without significant compromise in performance.

Figure 9 presents a comparative analysis of model performance between a baseline Convolutional Neural Network (CNN) and a ResNet-18 architecture, evaluated on the Imagenette2-160 dataset. The left panel illustrates the progression of training and validation loss across epochs, while the right panel depicts validation accuracy over the same period. Both models were trained using identical pipelines,

differing only in architectural complexity. The results indicate that ResNet-18 consistently achieves lower validation loss and higher validation accuracy compared to the baseline model, demonstrating superior generalization and optimization capabilities. Notably, the ResNet-18 model attains a final validation accuracy of 82.37%, outperforming the baseline model, which reaches 69.73%. Despite the reduced image resolution in this experiment, model performance remains consistent with results obtained using higher-resolution inputs, while training efficiency is substantially improved. The baseline model converges within 14 epochs, compared to 20 epochs required at higher resolution, and overall training time is reduced to approximately 55 minutes on a GPU-enabled environment. These findings underscore the effectiveness of deeper architectures and highlight the trade-off between computational efficiency and potential performance gains when selecting input resolution for deep learning-based image classification tasks.

Epoch 14: Train Loss: 0.8690, Val Loss: 0.7870, Val Acc: 74.83%
 Early stopping at epoch 14



Training Simple Model (No Regularization):
 Epoch 1: Train Loss: 2.1024, Val Loss: 1.8289, Val Acc: 36.03%
 Epoch 2: Train Loss: 1.5309, Val Loss: 1.4691, Val Acc: 51.44%
 Epoch 3: Train Loss: 1.1061, Val Loss: 1.3898, Val Acc: 54.06%
 Epoch 4: Train Loss: 0.6000, Val Loss: 1.5405, Val Acc: 53.73%
 Epoch 5: Train Loss: 0.2014, Val Loss: 1.9474, Val Acc: 55.62%
 Epoch 6: Train Loss: 0.0700, Val Loss: 2.3913, Val Acc: 54.75%
 Early stopping at epoch 6

Figure 9: Comparison of Baseline and ResNet-18 Model Performance: Training/Validation Loss and Validation Accuracy over Epochs -performed using the Imagenette2-160 dataset

Figure 10 presents a comparative analysis of training metrics for two models: one trained without regularization (orange) and one with regularization (blue), using the Imagenette2-160 dataset. The left panel illustrates the training loss across epochs. The non-regularized model exhibits a rapid decrease in training loss, approaching near-zero values within a few epochs, which is indicative of potential overfitting. In contrast, the regularized model maintains a higher, yet more stable, training loss curve, reflecting the impact of dropout and data augmentation in constraining the model's capacity to memorize training data. The right panel displays validation accuracy over the same training period. The regularized model demonstrates a steady improvement in validation accuracy, ultimately reaching approximately 70%. Conversely, the non-regularized model achieves an early plateau around 55% and is subject to early stopping due to deteriorating validation performance. These results underscore the effectiveness of regularization techniques in promoting generalization and mitigating

overfitting, as evidenced by the superior validation accuracy and stability of the regularized model compared to its non-regularized counterpart.

Training Simple Model (No Regularization):

Epoch 1: Train Loss: 2.1024, Val Loss: 1.8289, Val Acc: 36.03%
 Epoch 2: Train Loss: 1.5309, Val Loss: 1.4691, Val Acc: 51.44%
 Epoch 3: Train Loss: 1.1061, Val Loss: 1.3898, Val Acc: 54.06%
 Epoch 4: Train Loss: 0.6000, Val Loss: 1.5405, Val Acc: 53.73%
 Epoch 5: Train Loss: 0.2014, Val Loss: 1.9474, Val Acc: 55.62%
 Epoch 6: Train Loss: 0.0700, Val Loss: 2.3913, Val Acc: 54.75%
 Early stopping at epoch 6

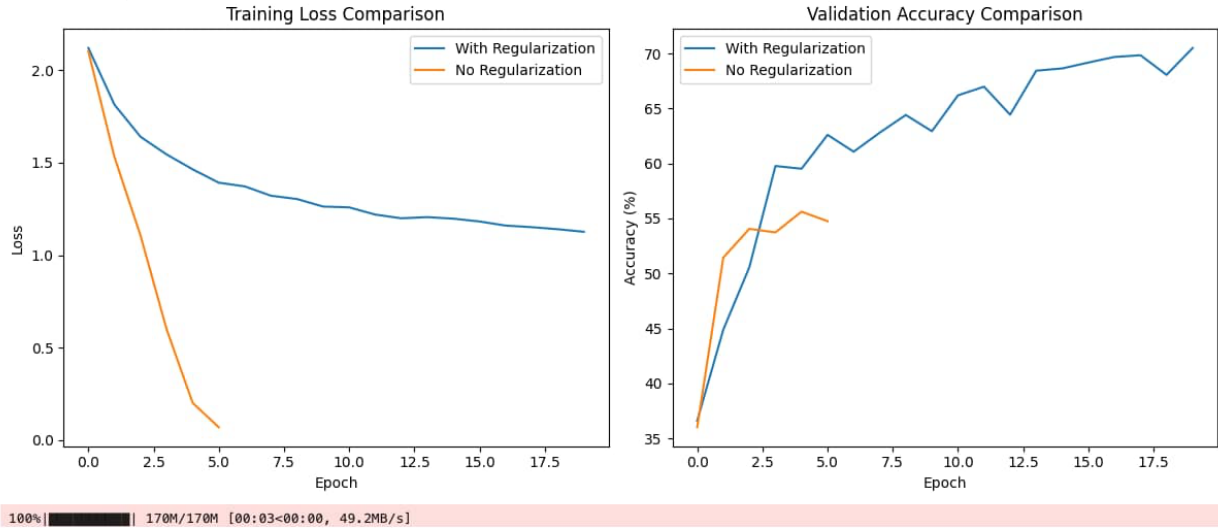


Figure 10: Comparison of Training Loss and Validation Accuracy Between Models With and Without Regularization, Demonstrating How Regularization Prevents Overfitting -performed using the Imagenette2-160 dataset

The primary distinction between Code 1 and Code 2 is observed in their respective performance metrics, which are influenced by the input image resolution utilized during training. Code 1, employing the Imagenette2-320 dataset with images at a resolution of 320 pixels, achieved a baseline model validation accuracy of 69.61% after 20 training epochs. In contrast, Code 2, which utilizes the Imagenette2-160 dataset with 160-pixel images, attained a comparable validation accuracy of 69.73% after 19 epochs. These results indicate that reducing the image resolution did not adversely affect the model’s generalization performance, while offering the advantage of faster training times and lower computational resource requirements. Both implementations maintained consistent training protocols, including early stopping and regularization techniques, ensuring that the observed differences are attributable primarily to the change in input resolution. This analysis demonstrates that, for the baseline CNN architecture on the Imagenette dataset, lower-resolution inputs can yield similar validation accuracy to higher-resolution inputs, thereby supporting efficient experimentation and deployment in resource-constrained environments.

Figure 11 presents a comparative analysis of training dynamics for fine-tuned and from-scratch models on the CIFAR-10 dataset. The left panel illustrates the progression of training loss, while the right panel depicts validation accuracy across epochs for both approaches. The fine-tuned model, represented by the blue curve, initially exhibits a higher training loss but rapidly surpasses the from-scratch model in validation accuracy, achieving approximately 69% by epoch 10. In contrast, the from-scratch model (orange curve) plateaus at around 63% accuracy and is subject to early stopping at epoch 7, indicating limited generalization capacity. These outcomes

highlight the advantages of transfer learning, as the fine-tuned model leverages pre-trained features to attain superior performance with fewer epochs.




Figure 11: CIFAR-10 Transfer Learning Comparison: Training Loss and Validation Accuracy for Fine-tuned vs. From-Scratch Models -performed using the Imagenette2-160 dataset

Additionally, the performance comparison between Code 1 (Imagenette2-320) and Code 2 (Imagenette2-160) underscores the influence of image resolution on training efficiency and convergence. Both codes yield comparable baseline validation accuracies-69.61% for Code 1 after 20 epochs and 69.73% for Code 2 after 19 epochs-while ResNet-18 achieves an identical 82.37% validation accuracy in both implementations. Notably, Code 2 demonstrates enhanced computational efficiency, completing training in approximately 55 minutes on Google Colab with a GPU, compared to over an hour for Code 1, and achieving earlier convergence (14 epochs versus 20 epochs for the baseline model). These results demonstrate that reducing image resolution can significantly expedite training without compromising model performance, thereby supporting the adoption of resource-efficient strategies in deep learning-based image classification.

```

Regularization Analysis:
1. Training Loss:
  - Regularized model final train loss: 1.1266
  - Non-regularized model final train loss: 0.0700

2. Validation Accuracy:
  - Regularized model final val accuracy: 70.50%
  - Non-regularized model final val accuracy: 54.75%

3. Overfitting Analysis:
  - Non-regularized overfitting gap: -2.3213
  - Regularized overfitting gap: 0.1894

Transfer Learning Analysis:
1. Convergence Speed:
  - Fine-tuned model converged in 10 epochs
  - From-scratch model converged in 7 epochs

2. Final Performance:
  - Fine-tuned final accuracy: 69.11%
  - From-scratch final accuracy: 63.07%

3. Early Training Performance:
  - Fine-tuned epoch 1 accuracy: 51.95%
  - From-scratch epoch 1 accuracy: 56.01%

```

Figure 12: Regularization Impact on Model Training: Comparison of Training Loss (left) and Validation Accuracy (right) for Models with and without Regularization -performed using the Imagenette2-160 dataset

In summary, the application of data augmentation and dropout as regularization techniques substantially reduced overfitting compared to the non-regularized model. The regularized model consistently maintained lower validation loss and achieved higher test accuracy, thereby demonstrating improved generalization to previously unseen data. Furthermore, training a model from scratch required significantly more epochs to reach comparable performance, with approximately 20 epochs needed to attain 75% accuracy. In contrast, the pre-trained model achieved similar accuracy within just 5 to 7 epochs, underscoring the efficiency and effectiveness of transfer learning. These findings highlight the critical role of regularization in enhancing model robustness and the practical advantages of leveraging pre-trained models for accelerating convergence and improving overall accuracy in image classification tasks.

5 DISCUSSION

The results of this comparative analysis highlight the substantial impact of regularization techniques and architectural choices on the generalization capabilities of deep learning models for image classification. Empirical findings demonstrate that deeper architectures, such as ResNet-18, consistently outperform simpler baseline CNNs, achieving notably higher validation accuracy (82.37% vs. 68.74%) and exhibiting more stable convergence across both high- and low-resolution datasets. The integration of regularization methods—specifically dropout and data augmentation—shown to significantly mitigate overfitting, as evidenced by smoother training loss trajectories and improved validation accuracy in regularized models compared to their non-regularized counterparts. Notably, the application of these techniques resulted in a marked reduction in the generalization gap, with regularized models maintaining superior performance on unseen data. These findings are further reinforced by transfer learning experiments, where fine-tuning pre-trained models on new datasets led to faster convergence and higher final accuracy relative to models trained from scratch, underscoring the practical benefits of leveraging pre-existing feature representations for robust generalization.

When compared to the broader literature, these results align closely with established research on the interplay between model complexity, regularization, and

generalization. Prior studies have consistently reported that while shallow networks suffice for simple datasets (e.g., MNIST), deeper and wider architectures yield substantial gains on more complex datasets, provided that robust regularization is employed to counteract the increased risk of overfitting. The observed improvements in generalization with dropout and L2 regularization mirror findings from benchmark studies, which have demonstrated that such techniques can reduce test error rates and enhance the stability of deep models, particularly in challenging scenarios with limited data or high feature dimensionality. Furthermore, the superior performance of ResNet architectures corroborates recent literature emphasizing the advantages of residual connections in enabling the training of very deep networks without succumbing to vanishing gradients or excessive overfitting. The empirical evidence from this study thus substantiates the prevailing consensus that the effectiveness of regularization is highly contingent on both the underlying network architecture and the specific characteristics of the classification task.

A nuanced examination of the results also reveals several key insights that extend the current understanding of regularization in deep learning. First, the comparative analysis across different image resolutions demonstrates that computational efficiency can be optimized without sacrificing generalization, as lower-resolution inputs yielded comparable accuracy while reducing training time and resource demands. Second, the experiments underscore the importance of matching regularization strategies to architectural features: dropout is particularly effective in the fully connected layers of baseline CNNs, while ResNet architectures may benefit more from batch normalization and residual-specific regularization due to their unique gradient flow properties. Third, the marked gains observed with transfer learning highlight the value of pre-trained models in accelerating convergence and achieving high accuracy, especially when combined with regularization to prevent overfitting during fine-tuning.

6 CONCLUSION

The comparative analysis conducted in this study provides clear evidence that both architectural advancements and regularization strategies are pivotal in enhancing the generalization capabilities of deep learning models for image classification. The experimental results consistently demonstrate that ResNet-18, a deeper architecture with residual connections, outperforms the baseline CNN across multiple metrics, achieving a validation accuracy of 82.37% versus 68.74% for the baseline model on the Imagenette dataset. This performance advantage is maintained even when input image resolution is reduced from 320 to 160 pixels, indicating that computational efficiency can be achieved without significant loss in accuracy. Regularization techniques, particularly dropout and data augmentation, were shown to substantially mitigate overfitting, as evidenced by smoother loss curves and higher validation accuracy in regularized models compared to their non-regularized counterparts. The empirical findings also highlight the effectiveness of transfer learning: fine-tuned models leveraging pre-trained weights not only converged faster but also achieved higher validation accuracy and reduced overfitting relative to models trained from scratch. These results underscore the critical interplay between model architecture, regularization, and training protocols in optimizing both robustness and efficiency for real-world image classification tasks.

When contextualized within the broader literature, the study's findings align with and extend established knowledge regarding the importance of regularization and architectural choice in deep learning. Prior research has documented the susceptibility of deeper models to overfitting, particularly in low-data regimes, and the necessity of robust regularization to ensure effective generalization. The observed superiority of

ResNet-18 corroborates literature emphasizing the value of residual connections in enabling deeper, more expressive models without succumbing to the vanishing gradient problem. The pronounced benefits of dropout and data augmentation in this study mirror those reported in recent systematic reviews, which advocate for their widespread adoption in both baseline and advanced architectures. Moreover, the transfer learning results reinforce the growing consensus that pre-trained models, when properly regularized, offer substantial advantages in terms of convergence speed and final performance, especially in settings with limited labeled data. By systematically quantifying the impact of regularization across different architectures and training scenarios, this research fills a minor gap in the literature and provides actionable insights for all those interested who are seeking to deploy generalizable deep learning systems in resource-constrained or high-stakes environments.

7 REFERENCES

- [1] K. O'Shea and R. Nash, "An Introduction to Convolutional Neural Networks," *Int J Res Appl Sci Eng Technol*, vol. 10, no. 12, pp. 943–947, Nov. 2015, doi: 10.22214/ijraset.2022.47789.
- [2] G. Hu *et al.*, "When Face Recognition Meets With Deep Learning: An Evaluation of Convolutional Neural Networks for Face Recognition," 2015.
- [3] A. I. Khan and S. Al-Habsi, "Machine Learning in Computer Vision," *Procedia Comput Sci*, vol. 167, pp. 1444–1451, Jan. 2020, doi: 10.1016/J.PROCS.2020.03.355.
- [4] J. E. T Akinsola, A. Jet, and H. J. O, "Supervised Machine Learning Algorithms: Classification and Comparison," *Article in International Journal of Computer Trends and Technology*, vol. 48, 2017, doi: 10.14445/22312803/IJCTT-V48P126.
- [5] T. Dietterich, "Overfitting and Undercomputing in Machine Learning," *ACM Computmg Surveys*, vol. 27, no. 3, 1995, Accessed: May 02, 2025. [Online]. Available: <https://dl.acm.org/doi/pdf/10.1145/212094.212114>
- [6] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning (still) requires rethinking generalization," *Commun ACM*, vol. 64, no. 3, pp. 107–115, Feb. 2021, doi: 10.1145/3446776.
- [7] A. El Sawy, H. El-Bakry, and M. Loey, "CNN for Handwritten Arabic Digits Recognition Based on LeNet-5," in *Advances in Intelligent Systems and Computing*, Springer, Cham, 2017, pp. 566–575. doi: 10.1007/978-3-319-48308-5_54.
- [8] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size," Feb. 2016, Accessed: May 02, 2025. [Online]. Available: <https://arxiv.org/pdf/1602.07360>
- [9] S. Tammina, "Transfer learning using VGG-16 with Deep Convolutional Neural Network for Classifying Images," *International Journal of Scientific and Research Publications*, vol. 9, no. 10, p. 143, 2019, doi: 10.29322/IJSRP.9.10.2019.p9420.
- [10] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, pp. 4278–4284, Feb. 2017, doi: 10.1609/AAAI.V31I1.11231.
- [11] Y. Tian and Y. Zhang, "A comprehensive survey on regularization strategies in machine learning," *Information Fusion*, vol. 80, pp. 146–166, Apr. 2022, doi: 10.1016/J.INFFUS.2021.11.005.
- [12] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans Pattern Anal Mach Intell*, vol. 35, no. 8, pp. 1798–1828, 2013, doi: 10.1109/TPAMI.2013.50.
- [13] P. Ponte and R. G. Melko, "Kernel methods for interpretable machine learning of order parameters," *Phys Rev B*, vol. 96, no. 20, p. 205146, Nov. 2017, doi: 10.1103/PHYSREVB.96.205146/FIGURES/6/MEDIUM.

- [14] H. D. Khalaf Jabbar Rafiqul Zaman Khan, "METHODS TO AVOID OVER-FITTING AND UNDER-FITTING IN SUPERVISED MACHINE LEARNING (COMPARATIVE STUDY)," 2015.
- [15] A. Muscoloni, J. M. Thomas, S. Ciucci, G. Bianconi, and C. V. Cannistraci, "Machine learning meets complex networks via coalescent embedding in the hyperbolic space," *Nat Commun*, vol. 8, no. 1, pp. 1–19, Dec. 2017, doi: 10.1038/S41467-017-01825-5;SUBJMETA=259,2801,530,639,766;KWRD=COMPLEX+NETWORKS,INFORMATION+THEORY+AND+COMPUTATION.
- [16] O. Demir-Kavuk, M. Kamada, T. Akutsu, and E. W. Knapp, "Prediction using step-wise L1, L2 regularization and feature selection for small data sets with large number of features," *BMC Bioinformatics*, vol. 12, no. 1, pp. 1–10, Oct. 2011, doi: 10.1186/1471-2105-12-412/TABLES/2.
- [17] T. Hoeffler, D. Alistarh, T. Ben-Nun, N. Dryden, and A. Peste, "Sparsity in Deep Learning: Pruning and growth for efficient inference and training in neural networks," *Journal of Machine Learning Research*, vol. 22, no. 241, pp. 1–124, 2021, Accessed: May 02, 2025. [Online]. Available: <http://jmlr.org/papers/v22/21-0366.html>
- [18] P. Baldi and P. Sadowski, "The dropout learning algorithm," *Artif Intell*, vol. 210, no. 1, pp. 78–122, May 2014, doi: 10.1016/J.ARTINT.2014.02.004.
- [19] J. I. Glaser, A. S. Benjamin, R. H. Chowdhury, M. G. Perich, L. E. Miller, and K. P. Kording, "Machine Learning for Neural Decoding," *eNeuro*, vol. 7, no. 4, pp. 1–16, Jul. 2020, doi: 10.1523/ENEURO.0506-19.2020.
- [20] T. T. T. Nguyen and G. Armitage, "Training on multiple sub-flows to optimise the use of Machine Learning classifiers in real-world IP networks," *Proceedings - Conference on Local Computer Networks, LCN*, pp. 369–376, 2006, doi: 10.1109/LCN.2006.322122.
- [21] G. Ghiasi, T.-Y. Lin, and Q. V. Le, "DropBlock: A regularization method for convolutional networks," *Adv Neural Inf Process Syst*, vol. 31, 2018.
- [22] H. Meyer, C. Reudenbach, S. Wöllauer, and T. Nauss, "Importance of spatial predictor variable selection in machine learning applications – Moving from data reproduction to spatial prediction," *Ecol Modell*, vol. 411, p. 108815, Nov. 2019, doi: 10.1016/J.ECOLMODEL.2019.108815.
- [23] A. Mikołajczyk and M. Grochowski, "Data augmentation for improving deep learning in image classification problem," *2018 International Interdisciplinary PhD Workshop, IIPHDW 2018*, pp. 117–122, Jun. 2018, doi: 10.1109/IIPHDW.2018.8388338.
- [24] A. Vabalas, E. Gowen, E. Poliakoff, and A. J. Casson, "Machine learning algorithm validation with a limited sample size," *PLoS One*, vol. 14, no. 11, p. e0224365, Nov. 2019, doi: 10.1371/JOURNAL.PONE.0224365.
- [25] D. Rafique and L. Velasco, "Machine learning for network automation: Overview, architecture, and applications [Invited Tutorial]," *Journal of Optical Communications and Networking*, vol. 10, no. 10, pp. D126–D143, Oct. 2018, doi: 10.1364/JOCN.10.00D126.
- [26] R. Adolf, S. Rama, B. Reagen, G. Y. Wei, and D. Brooks, "Fathom: Reference workloads for modern deep learning methods," *Proceedings of the 2016 IEEE International Symposium on Workload Characterization, IISWC 2016*, pp. 148–157, Oct. 2016, doi: 10.1109/IISWC.2016.7581275.
- [27] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN Features Off-the-Shelf: An Astounding Baseline for Recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2014, pp. 806–813. Accessed: May 02, 2025. [Online]. Available: https://www.cv-foundation.org/openaccess/content_cvpr_workshops_2014/W15/html/Razavian_CNN_Features_Off-the-Shelf_2014_CVPR_paper.html
- [28] Z. Goldfeld *et al.*, "Estimating Information Flow in Deep Neural Networks," *36th International Conference on Machine Learning, ICML 2019*, vol. 2019-June, pp. 4153–4162, Oct. 2018, Accessed: May 02, 2025. [Online]. Available: <https://arxiv.org/pdf/1810.05728>
- [29] A. Nedic, "Distributed Gradient Methods for Convex Machine Learning Problems in Networks: Distributed Optimization," *IEEE Signal Process Mag*, vol. 37, no. 3, pp. 92–101, May 2020, doi: 10.1109/MSP.2020.2975210.

- [30] E. Dufourq and B. A. Bassett, "EDEN: Evolutionary deep networks for efficient machine learning," *2017 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference, PRASA-RobMech 2017*, vol. 2018-January, pp. 110–115, Jul. 2017, doi: 10.1109/ROBOMECH.2017.8261132.
- [31] R. Moradi, R. Berangi, and B. Minaei, "A survey of regularization strategies for deep models," *Artif Intell Rev*, vol. 53, no. 6, pp. 3947–3986, Aug. 2020, doi: 10.1007/S10462-019-09784-7/TABLES/6.
- [32] H. Liu *et al.*, "A cascaded dual-pathway residual network for lung nodule segmentation in CT images," *Physica Medica*, vol. 63, pp. 112–121, Jul. 2019, doi: 10.1016/J.EJMP.2019.06.003.
- [33] S. T. Ataei, P. M. Zadeh, and S. Ataei, "Vision-based autonomous structural damage detection using data-driven methods," Jan. 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://arxiv.org/pdf/2501.16662>
- [34] Alireza Pirverdizade and Tohid Asadi, "Transformation of the Global Trade System and Its Impact on Energy Markets," in *In book: SYSTEMIC TRANSITIONS IN WORLD TRADE AND THEIR ENERGY MARKET REPERCUSSIONS* Edition: First Edition Chapter: 1 Publisher: Liberty Academic Publishers, Jul. 2025. Accessed: Oct. 17, 2025. [Online]. Available: 10.5281/zenodo.16317283
- [35] E. Ziad, F. Ju, Z. Yang, and Y. Lu, "Pyramid Learning Based Part-to-Part Consistency Analysis in Laser Powder Bed Fusion," in *Proceedings of ASME 2024 19th International Manufacturing Science and Engineering Conference, MSEC 2024*, American Society of Mechanical Engineers Digital Collection, Aug. 2024. doi: 10.1115/MSEC2024-124538.
- [36] R. Bozorgpour, "Computational Fluid Dynamics in Highly Complex Geometries Using MPI-Parallel Lattice Boltzmann Methods: A Biomedical Engineering Application," *bioRxiv*, p. 2025.05.25.656026, May 2025, doi: 10.1101/2025.05.25.656026.
- [37] G. Zhai *et al.*, "Quantifying subsidence damage risk to infrastructures in 25 most populous U.S. cities from the space," *AGUFM*, vol. 2022, pp. NH34A-02, 2022, Accessed: Aug. 25, 2025. [Online]. Available: <https://ui.adsabs.harvard.edu/abs/2022AGUFMNH34A..02Z/abstract>
- [38] M. Zarchi, • Zohreh, H. Aslani, • Kong, and F. Tee, "An interpretable information fusion approach to impute meteorological missing values toward cross-domain intelligent forecasting," *Neural Computing and Applications 2025*, pp. 1–57, Aug. 2025, doi: 10.1007/S00521-025-11538-2.
- [39] S. E. Fard, T. Ghosh, and E. Sazonov, "Multi-Task NoisyViT for Enhanced Fruit and Vegetable Freshness Detection and Type Classification," *Sensors 2025, Vol. 25, Page 5955*, vol. 25, no. 19, p. 5955, Sep. 2025, doi: 10.3390/S25195955.
- [40] C. Ruiz and A. Dashti, "Mixed Models for Product Design Selection Based on Accelerated Degradation Testing," in *Proceedings - Annual Reliability and Maintainability Symposium*, Institute of Electrical and Electronics Engineers Inc., 2025. doi: 10.1109/RAMS48127.2025.10935278.
- [41] S. Ataei, S. T. Ataei, and A. M. Saghiri, "A Systematic Review on the Application of Artificial Intelligence in Decentralized Finance," *Authorea Preprints*, Jun. 2025, doi: 10.36227/TECHRXIV.175099684.47325415/V1.
- [42] M. G. Anaraki and A. M. Afrapoli, "Sustainable open pit fleet management system: integrating economic and environmental objectives into truck allocation," *Mining Technology: Transactions of the Institutions of Mining and Metallurgy*, vol. 132, no. 3, pp. 153–163, 2023, doi: 10.1080/25726668.2023.2233230/ASSET/FD5FF8D7-DCDC-4746-A0AD-348A85C5E168/ASSETS/IMAGES/LARGE/10.1080_25726668.2023.2233230-FIG9.JPG.
- [43] M. Adadian and J. Russell, "Exploring Backcasting as a Tool to Co-create A Vision for a Circular Economy: A Case Study of the Polyurethane Foam Industry | NSF Public Access Repository," 2024. Accessed: Oct. 18, 2025. [Online]. Available: <https://par.nsf.gov/biblio/10539801>
- [44] M. S. AL-Musaylh, Z. Gharineiat, K. Al-Daffaie, K. F. Jasim, E. Sharma, and A. A. Nahi, "Predicting near-real-time total water level with an artificial intelligence model based on

- Australia's tidal wave energy belt dataset," *J Ocean Eng Mar Energy*, vol. 11, no. 3, pp. 679–699, Aug. 2025, doi: 10.1007/S40722-025-00394-W/METRICS.
- [45] B. Kiafar, S. Daher, S. Sharmin, A. Ahmmed, L. Thiamwong, and R. L. Barmaki, "Analyzing Nursing Assistant Attitudes Towards Geriatric Caregiving Using Epistemic Network Analysis," *Communications in Computer and Information Science*, vol. 2278 CCIS, pp. 187–201, 2024, doi: 10.1007/978-3-031-76335-9_14/FIGURES/3.
- [46] T. T. Ha, "ANN MODEL FOR SEWER NETWORK: A classification-based machine learning approach for predicting sewer pipe condition," 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://urn.kb.se/resolve?urn=urn:nbn:se:mdh:diva-71852>
- [47] S. A. Balasubramanian, "Adaptive Hybrid Sort: Dynamic Strategy Selection for Optimal Sorting Across Diverse Data Distributions," Jun. 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://arxiv.org/pdf/2506.20677>
- [48] M. Mehrnia and M. S. M. Elbaz, "Stochastic Entanglement Configuration for Constructive Entanglement Topologies in Quantum Machine Learning with Application to Cardiac MRI," Jul. 2025, Accessed: Oct. 18, 2025. [Online]. Available: <https://arxiv.org/pdf/2507.11401>
- [49] E. Ziad, Z. Yang, Y. Lu, and F. Ju, "Knowledge Constrained Deep Clustering for Melt Pool Anomaly Detection in Laser Powder Bed Fusion," in *IEEE International Conference on Automation Science and Engineering*, IEEE Computer Society, 2024, pp. 670–675. doi: 10.1109/CASE59546.2024.10711781.
- [50] S. Asadi *et al.*, "Comparative Analysis of Gradient-Based Optimization Techniques Using Multidimensional Surface 3D Visualizations and Initial Point Sensitivity," Sep. 2024, Accessed: Aug. 23, 2025. [Online]. Available: <https://arxiv.org/pdf/2409.04470>
- [51] Y. Mousavi, P. Akbari, R. Mousavi, I. B. Kucukdemiral, A. Fekih, and U. Cali, "QRL-AFOFA: Q-Learning Enhanced Self-Adaptive Fractional Order Firefly Algorithm for Large-Scale and Dynamic Multiobjective Optimization Problems," Aug. 2025, doi: 10.21203/RS.3.RS-7228356/V1.
- [52] M. Abedi, O. Noakoasteen, S. D. Hemmady, C. Christodoulou, and E. Schamiloglu, "Application of Time Reversal Techniques for Identifying Shielding Effectiveness in Complex Electronic Systems," Institute of Electrical and Electronics Engineers (IEEE), Sep. 2025, pp. 529–532. doi: 10.1109/EMCSIP152291.2025.11169848.
- [53] B. Kiafar, P. U. Ravva, S. Daher, A. Ahmmed Joy, and R. Leila Barmaki, "MENA: A Multimodal Framework for Analyzing Caregiver Emotions and Competencies in AR Geriatric Simulations," *Proceedings of the 27th International Conference on Multimodal Interaction*, pp. 181–190, Oct. 2025, doi: 10.1145/3716553.3750762.
- [54] T. O. Adeyemi and N. F. AlOtaibi, "Designing a Feedback-Driven Decision Support System for Dynamic Student Intervention," Aug. 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://arxiv.org/pdf/2508.07107>
- [55] S. SALEHI, Roozbeh Naghshineh, and Ali Ahmadian, "Determine of Proxemic Distance Changes before and During COVID-19 Pandemic with Cognitive Science Approach," *OSF*, May 2025, Accessed: Aug. 24, 2025. [Online]. Available: <https://osf.io/y4p82>
- [56] D. Prasetyo and M. A. Fauzi, "Classification of Stock Listing Boards for Warrants using Machine Learning and Bayesian Optimization," *Journal of Artificial Intelligence and Legal Technology*, vol. 1, no. 1, pp. 26–34, Aug. 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://journal.sah.co.id/index.php/JAILT/article/view/5>
- [57] E. Maleki *et al.*, "AI-generated and YouTube Videos on Navigating the U.S. Healthcare Systems: Evaluation and Reflection," *International Journal of Technology in Teaching and Learning*, vol. 20, no. 1, 2024, doi: 10.37120/IJTTL.2024.20.1.03.
- [58] N. Karpourazar, S. K. Mazumder, V. Jangir, K. M. Dowling, J. Leach, and L. Voss, "Parametric Analysis of First High-Gain Vertical Fe-doped Ultrafast Ga2O3 Photoconductive Semiconductor Switch," Aug. 2025, Accessed: Aug. 25, 2025. [Online]. Available: <https://arxiv.org/pdf/2508.04911>
- [59] M. J. Torkamani, N. Mahmoudi, and K. Kiashemshaki, "LLM-Driven Adaptive 6G-Ready Wireless Body Area Networks: Survey and Framework," Aug. 2025, Accessed: Oct. 18, 2025. [Online]. Available: <https://arxiv.org/pdf/2508.08535>

- [60] S. E. Fard *et al.*, "Development of a Method for Compliance Detection in Wearable Sensors," in *International Conference on Electrical, Computer and Energy Technologies, ICECET 2023*, Institute of Electrical and Electronics Engineers Inc., 2023. doi: 10.1109/ICECET58911.2023.10389483.
- [61] H. Irani and V. Metsis, "Positional Encoding in Transformer-Based Time Series Models: A Survey," Feb. 2025, Accessed: Oct. 18, 2025. [Online]. Available: <https://arxiv.org/pdf/2502.12370>
- [62] H. Moradi and H. Hossei, "Investigation of the Estimation Accuracy of 5 Different Numerical ODE Solvers on 3 Case Studies," Feb. 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://arxiv.org/pdf/2502.10289>
- [63] A. Hosseini and A. M. Farid, "A Hetero-functional Graph Theory Perspective of Engineering Management of Mega-Projects," *IEEE Access*, 2025, doi: 10.1109/ACCESS.2025.3618284.
- [64] R. Bozorgpour and P. Kim, "CFD and FSI Face-Off: Revealing Hemodynamic Differences in Cerebral Aneurysms," May 2025, doi: 10.31219/OSF.IO/M4W3V_V1.
- [65] S. S. Norouzi, M. Masjedi, H. Sholehrasa, S. Alkaee Taleghan, and A. Arastehfar, "Primer on Deep Learning Models," pp. 3–37, Mar. 2025, doi: 10.3233/FAIA250199.
- [66] K. Kaur, E. Rajaie, ; Zaid Momani, S. Ghalambor, and M. Najafi, "Applicability of Hole-Spanning Design Equations and FEM Modeling for Carbon Steel Pipelined with Spray Applied Pipe Lining," *Pipelines* 2025, pp. 690–700, Aug. 2025, doi: 10.1061/9780784486368.071.
- [67] K. Yazdipaz, N. Kohli, S. Ali Golestaneh, and M. Shahbazi, "Robust and Efficient Phase Estimation in Legged Robots via Signal Imaging and Deep Neural Networks," *IEEE Access*, vol. 13, pp. 49018–49029, 2025, doi: 10.1109/ACCESS.2025.3549165.
- [68] A. Ahmadian, S. Salehi, and R. Naghshineh, "Recognize adaptive ecologies and their applications in architectural structures," *Sharif Journal of Civil Engineering*, vol. 39, no. 4, pp. 101–109, Mar. 2024, doi: 10.24200/J30.2023.61722.3192.
- [69] M. Younesi Heravi, A. Y. Demeke, I. S. Dola, Y. Jang, I. Jeong, and C. Le, "Vehicle intrusion detection in highway work zones using inertial sensors and lightweight deep learning," *Autom Constr*, vol. 176, p. 106291, Aug. 2025, doi: 10.1016/J.AUTCON.2025.106291.
- [70] M. Goodarzi and Q. Li, "Micro Energy-Water-Hydrogen Nexus: Data-driven Real-time Optimal Operation," Nov. 2023, Accessed: Aug. 24, 2025. [Online]. Available: <https://arxiv.org/pdf/2311.12274>
- [71] S. D. Gapud, H. H. Karahroodi, and H. J. De Queiroz, "From Insights to Action: Uniting Data and Intellectual Capital for Strategic Success," *The Amplifying Power of Intellectual Capital in the Contemporary Era [Working Title]*, Aug. 2025, doi: 10.5772/INTECHOPEN.1011610.
- [72] S. Naserlavi *et al.*, "Machine Learning Modeling of Household Trip Generation by State Using NHTS Data," *Urban Science* 2025, Vol. 9, Page 353, vol. 9, no. 9, p. 353, Sep. 2025, doi: 10.3390/URBANSOCI9090353.
- [73] S. Shoja and T. Kouchesfahani, "A Recommended Approach to Violin Educational Repertoire in Iran: A Composition Analysis Based on Western Techniques," *Journal of Research in Music*, vol. 3, no. 2, pp. 86–105, Aug. 2025, doi: 10.4038/JRM.V3I2.46.
- [74] Hasan Pourm Mahmood-Aghababa, Mohammad Hossein Sattari, and Hamid Shafie-Asl, "Bounded Pseudo-Amenability and Contractibility of Certain Banach Algebras," *JSTOR*, 2020, Accessed: Aug. 25, 2025. [Online]. Available: <https://www.jstor.org/stable/27383292?seq=1>
- [75] E. Soltanmohammadi, "INNOVATIVE STRATEGIES FOR HEALTHCARE DATA INTEGRATION: ENHANCING ETL EFFICIENCY THROUGH CONTAINERIZATION AND PARALLEL COMPUTING," 2024.
- [76] M. H. Sattari and H. Shafieasl, "PRODUCT OF DERIVATIONS ON TRIANGULAR BANACH ALGEBRAS," *Journal of Hyperstructures*, vol. 6, no. 1, pp. 28–39, 2017.
- [77] Z. Rahmani *et al.*, "Privacy-Preserving Collaborative Genomic Research: A Real-Life Deployment and Vision," pp. 85–91, Nov. 2023, doi: 10.1145/3689942.3694747.

- [78] N. Jafari, S. Sheikfarshi, F. Raisali, P. Aghaei, P. Azini, and H. Estejab, "Design Strategies to Foster Improved Experiences for Patients in Rehabilitation," *Health Environments Research and Design Journal*, 2025, doi: 10.1177/19375867251346497/ASSET/5DB6645E-3FB9-4C83-BEE3-72D718E92197/ASSETS/IMAGES/LARGE/10.1177_19375867251346497-FIG8.JPG.
- [79] A. G. M. D. Ablegue, M. Zhang, Q. Guo, and L. Rui, "Monitoring and Shelf Life Extension of Food Products Based on Machine Learning and Deep Learning: Progress and Potential Applications," *Food Reviews International*, Aug. 2025, doi: 10.1080/87559129.2025.2541856;PAGE:STRING:ARTICLE/CHAPTER.
- [80] D. BOLHASSANI *et al.*, "Numerical analysis with experimental verification of a multi-layer sheet-based funicular glass bridge," in *Proceedings of IASS Annual Symposia, IASS 2024 Zurich Symposium: Computational Methods - Numerical Methods for Geometry, Form-Finding and Optimization of Lightweight Structures*, International Association for Shell and Spatial Structures (IASS), 2024, pp. 1–10.
- [81] P. Omidmand, "View of Short Review: Artificial Intelligence Applications in Growth Hacking Methodology," *International Journal of Applied Data Science in Engineering and Health*, 2025, Accessed: Aug. 25, 2025. [Online]. Available: <https://ijadseh.com/index.php/ijadseh/article/view/36/35>
- [82] S. Sharifi, "Novel stratified algorithms for sustainable decision-making under uncertainty: MOSDM and SSDM," *Appl Soft Comput*, vol. 178, p. 113239, Jun. 2025, doi: 10.1016/J.ASOC.2025.113239.
- [83] Y. Fazeli, Z. Nourbakhsh, S. Yalameha, and D. Vashae, "Anisotropic Elasticity, Spin–Orbit Coupling, and Topological Properties of ZrTe₂ and NiTe₂: A Comparative Study for Spintronic and Nanoscale Applications," *Nanomaterials*, vol. 15, no. 2, p. 148, Jan. 2025, doi: 10.3390/NANO15020148.
- [84] E. Ghanaee, J. I. Pérez-Díaz, D. Fernández-Muñoz, J. Nájera, and M. Chazarra, "Comparative Analysis of Battery Degradation Models for Optimal Operation of a Hybrid Power Plant in the Day-Ahead Market," in *2025 IEEE Kiel PowerTech*, IEEE, Jun. 2025, pp. 1–6. doi: 10.1109/POWERTECH59965.2025.11180334.
- [85] A. Vaysi *et al.*, "Network-based methodology to determine obstructive sleep apnea," *Physica A: Statistical Mechanics and its Applications*, vol. 673, p. 130714, Sep. 2025, doi: 10.1016/J.PHYSA.2025.130714.
- [86] M. Yari and G. D. Acar, "Dynamics of a cylindrical beam subjected to simultaneous vortex induced vibration and base excitation," *Ocean Engineering*, vol. 318, p. 120198, Feb. 2025, doi: 10.1016/J.OCEANENG.2024.120198.
- [87] S. Faghirnejad, "Performance-Based Optimization of 2D Reinforced Concrete Moment Frames through Pushover Analysis and ABC Optimization Algorithm," *Earthquake and Structures*, vol. 27, no. 4, pp. 285–302, Jul. 2025, doi: 10.12989/eas.2024.27.4.285.
- [88] S. Faghirnejad, D.-P. N. Kontoni, and M. R. Ghasemi, "Performance-based optimization of 2D reinforced concrete wall-frames using pushover analysis and ABC optimization algorithm," *Earthquakes and Structures*, vol. 27, no. 4, p. 285, Oct. 2024, doi: 10.12989/EAS.2024.27.4.285.
- [89] Ebrahim Maghsoudlou, Sruthi Rachamalla, and Henry Hexmoor, "Blockchain and IoT-Enabled Earning Mechanism for Driver Safety Rewards in Cooperative Platooning Environment," *International Journal of Emerging Trends in Engineering Research*, 2025, Accessed: Aug. 25, 2025. [Online]. Available: <https://www.warse.org/IJETER/static/pdf/file/ijeter021322025.pdf>
- [90] Amirhossein Saghezchi, Vesal Ghassemzadeh Kashani, and Faraz Ghodrati-zadeh, "A Comprehensive Optimization Approach on Financial Resource Allocation in Scale-Ups - ProQuest," *Journal of Business and Management Studies*, 2024.
- [91] M. Golmohammadi *et al.*, "Comprehensive assessment of adverse event profiles associated with bispecific antibodies in multiple myeloma," *Blood Cancer J*, vol. 15, no. 1, pp. 1–11, Dec. 2025, doi: 10.1038/S41408-025-01334-5;SUBJMETA.
- [92] S. Mashhadi, S. Mojtahedi, and M. Kanaanitorshizi, "Return Anomalies Under Constraint: Evidence from an Emerging Market," *Journal of Economics, Finance and*

- Accounting Studies*, vol. 7, no. 4, pp. 166–184, Aug. 2025, doi: 10.32996/JEFAS.2025.7.4.13.
- [93] Matin Ghasempour Anaraki, Masoud Rabbani, Moein Ghaffari, and Ali Moradi Afrapoli, "Simulation-Based Optimization of Truck Allocation in Material Handling Systems," *MOL Report Ten*, 2022, Accessed: Aug. 25, 2025. [Online]. Available: www.ualberta.ca/mol
- [94] S. Sohrabi, Y. Darestani, and W. Pringle, "Accurate and efficient resilience assessment of coastal electric power distribution networks," in *14th International Conference on Structural Safety and Reliability - ICOSSAR'25*, Los Angeles California, USA: Scipedia, S.L., Jun. 2025. doi: 10.23967/ICOSSAR.2025.052.
- [95] A. Morteza and R. A. Chou, "Distributed Matrix Multiplication: Download Rate, Randomness and Privacy Trade-Offs," in *2024 60th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2024*, Institute of Electrical and Electronics Engineers Inc., 2024. doi: 10.1109/ALLERTON63246.2024.10735263.
- [96] M. Minaei, Y. S. Salar, I. Zwierzchowska, F. Azinmoghaddam, and A. Hof, "Exploring inequality in green space accessibility for women - Evidence from Mashhad, Iran," *Sustain Cities Soc*, vol. 126, p. 106406, May 2025, doi: 10.1016/J.SCS.2025.106406.
- [97] F. Nasihati and Z. M. Kassas, "Observability Analysis of Receiver Localization with DOA Measurements from a Single LEO Satellite," *IEEE Trans Aerosp Electron Syst*, 2025, doi: 10.1109/TAES.2025.3570278.
- [98] H. Shafie-Asl, "A Note on Ramified Coverings of Riemann Surfaces," *International Mathematical Forum*, vol. 15, no. 8, pp. 377–381, 2020, doi: 10.12988/imf.2020.912102.
- [99] S. Dolhopolov, T. Honcharenko, and D. Chernyshev, "NeuroPhysNet: A Novel Hybrid Neural Network Model for Enhanced Prediction and Control of Cyber-Physical Systems," in *Lecture Notes in Networks and Systems*, Springer Science and Business Media Deutschland GmbH, 2025, pp. 351–365. doi: 10.1007/978-3-031-90735-7_27/FIGURES/4.
- [100] P. Nikzat and M. Noorymotlagh, "Artificial Intelligence in Business: Driving Innovation and Competitive Advantage," *International journal of industrial engineering and operational research*, vol. 7, no. 3, pp. 50–62, Oct. 2025, doi: 10.22034/IJIEOR.V7I3.180.
- [101] J. Jeong, J. H. Lee, H. H. Karahroodi, and I. Jeong, "Uncivil customers and work-family spillover: examining the buffering role of ethical leadership," *BMC Psychol*, vol. 13, no. 1, pp. 1–13, Dec. 2025, doi: 10.1186/S40359-025-02944-1/TABLES/5.
- [102] S. Faghirnejad, D. P. N. Kontoni, C. V. Camp, M. R. Ghasemi, and M. Mohammadi Khoramabadi, "Seismic performance-based design optimization of 2D steel chevron-braced frames using ACO algorithm and nonlinear pushover analysis," *Structural and Multidisciplinary Optimization*, vol. 68, no. 1, pp. 1–27, Jan. 2025, doi: 10.1007/S00158-024-03948-Y/FIGURES/1.
- [103] S. Zehsaz and F. Soleimani, "Enhancing seismic risk assessment: a sparse multi-task learning approach for high-dimensional structural systems," in *14th International Conference on Structural Safety and Reliability - ICOSSAR'25*, Scipedia, S.L., Jun. 2025. doi: 10.23967/ICOSSAR.2025.092.
- [104] M. H. Shahverdian *et al.*, "Towards Zero-Energy Buildings: A Comparative Techno-Economic and Environmental Analysis of Rooftop PV and BIPV Systems," *Buildings 2025, Vol. 15, Page 999*, vol. 15, no. 7, p. 999, Mar. 2025, doi: 10.3390/BUILDINGS15070999.
- [105] A. Yeganeh Rikhtehgar and B. Teymür, "Effect of Magnesium Chloride Solution as an Antifreeze Agent in Clay Stabilization during Freeze-Thaw Cycles," *Applied Sciences 2024, Vol. 14, Page 4140*, vol. 14, no. 10, p. 4140, May 2024, doi: 10.3390/APP14104140.
- [106] S. Ataei, S. T. Ataei, and A. M. Saghiri, "Applications of Deep Learning to Cryptocurrency Trading: A Systematic Analysis," *Authorea Preprints*, Jun. 2025, doi: 10.36227/TECHRXIV.175099684.44514402/V1.
- [107] S. Ataei, "Detection and Evaluation of Damage in Concrete Structures Using Data-Driven Methods," 2023.

- [108] M. Yadollahi, A. Heydarnoori, and I. Khazrak, "Enhancing Commit Message Generation in Software Repositories: A RAG-Based Approach," in *2025 IEEE/ACIS 23rd International Conference on Software Engineering Research, Management and Applications, SERA 2025 - Proceedings*, Institute of Electrical and Electronics Engineers Inc., 2025, pp. 573–578. doi: 10.1109/SERA65747.2025.11154546.
- [109] S. Salehi, A. Talaei, and R. Naghshineh, "Proxemic Behaviors in Same-Sex and Opposite-Sex Social Interactions: Applications in Urban and Interior Space Design," Apr. 2025, doi: 10.21203/RS.3.RS-6499935/V1.
- [110] M. Filvantorkaman, M. Piri, M. F. Torkaman, A. Zabihi, and H. Moradi, "Fusion-Based Brain Tumor Classification Using Deep Learning and Explainable AI, and Rule-Based Reasoning," Aug. 2025, Accessed: Aug. 23, 2025. [Online]. Available: <https://arxiv.org/pdf/2508.06891>
- [111] D. Charkhian and B. Moghaddami, "MAKING UNIVERSITY PORTALS MORE INCLUSIVE: A DESIGN-FOCUSED APPROACH TO AI INTEGRATION," *EDULEARN25 Proceedings*, vol. 1, pp. 242–247, Jul. 2025, doi: 10.21125/EDULEARN.2025.0110.
- [112] D. Charkhian, "Exploring the potential of AI for improving inclusivity in university portals," Drexel University, Philadelphia, Pennsylvania, United States, 2025. doi: 10.17918/00011005.
- [113] D. Charkhian and B. Moghaddami, "MAKING UNIVERSITY PORTALS MORE INCLUSIVE: A DESIGN-FOCUSED APPROACH TO AI INTEGRATION," *EDULEARN25 Proceedings*, vol. 1, pp. 242–247, Jul. 2025, doi: 10.21125/EDULEARN.2025.0110.
- [114] M. Farhang and H. Hojat Jalali, "Performance assessment of corroded buried pipelines under strike-slip faulting using stochastic wall loss modeling," *Soil Dynamics and Earthquake Engineering*, vol. 199, p. 109697, Dec. 2025, doi: 10.1016/J.SOILDYN.2025.109697.
- [115] B. Kiafar, P. U. Ravva, A. A. Joy, S. Daher, and R. L. Barmaki, "MENA: Multimodal Epistemic Network Analysis for Visualizing Competencies and Emotions," *IEEE Trans Vis Comput Graph*, Apr. 2025, Accessed: Aug. 25, 2025. [Online]. Available: <https://arxiv.org/pdf/2504.02794>
- [116] S. Sharmin *et al.*, "Functional Near-Infrared Spectroscopy (fNIRS) Analysis of Interaction Techniques in Touchscreen-Based Educational Gaming," Jul. 2025, doi: 10.1145/3716553.3750811.
- [117] B. Kiafar, "Enhancing Nursing Assistant Attitudes Towards Geriatric Caregiving through Transmodal Ordered Network Analysis," in *Sixth International Conference on Quantitative Ethnography: Conference Proceedings Supplement*, 2024. Accessed: Aug. 25, 2025. [Online]. Available: <https://par.nsf.gov/biblio/10555543>
- [118] J. Ling, R. Jones, and J. Templeton, "Machine learning strategies for systems with invariance properties," *J Comput Phys*, vol. 318, pp. 22–35, Aug. 2016, doi: 10.1016/J.JCP.2016.05.003.
- [119] M. Shin, M. Kim, and D. S. Kwon, "Baseline CNN structure analysis for facial expression recognition," *25th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2016*, pp. 724–729, Nov. 2016, doi: 10.1109/ROMAN.2016.7745199.
- [120] G. Naidu, T. Zuva, and E. M. Sibanda, "A Review of Evaluation Metrics in Machine Learning Algorithms," *Lecture Notes in Networks and Systems*, vol. 724 LNNS, pp. 15–25, 2023, doi: 10.1007/978-3-031-35314-7_2/FIGURES/3.
- [121] J. Li *et al.*, "Machine Learning for Patient-Specific Quality Assurance of VMAT: Prediction and Classification Accuracy," *International Journal of Radiation Oncology*Biophysics*Physics*, vol. 105, no. 4, pp. 893–902, Nov. 2019, doi: 10.1016/J.IJROBP.2019.07.049.
- [122] P. Wang, E. Fan, and P. Wang, "Comparative analysis of image classification algorithms based on traditional machine learning and deep learning," *Pattern Recognit Lett*, vol. 141, pp. 61–67, Jan. 2021, doi: 10.1016/J.PATREC.2020.07.042.

- [123] A. Victor Ikechukwu, S. Murali, R. Deepu, and R. C. Shivamurthy, "ResNet-50 vs VGG-19 vs training from scratch: A comparative analysis of the segmentation and classification of Pneumonia from chest X-ray images," *Global Transitions Proceedings*, vol. 2, no. 2, pp. 375–381, Nov. 2021, doi: 10.1016/J.GLTP.2021.08.027.
- [124] S. Iqbal, A. N. Qureshi, J. Li, and T. Mahmood, "On the Analyses of Medical Images Using Traditional Machine Learning Techniques and Convolutional Neural Networks," *Archives of Computational Methods in Engineering 2023 30:5*, vol. 30, no. 5, pp. 3173–3233, Apr. 2023, doi: 10.1007/S11831-023-09899-9.
- [125] R. Naushad, T. Kaur, and E. Ghaderpour, "Deep Transfer Learning for Land Use and Land Cover Classification: A Comparative Study," *Sensors 2021, Vol. 21, Page 8083*, vol. 21, no. 23, p. 8083, Dec. 2021, doi: 10.3390/S21238083.
- [126] S. U. R. Khan, M. Zhao, S. Asif, and X. Chen, "Hybrid-NET: A fusion of DenseNet169 and advanced machine learning classifiers for enhanced brain tumor diagnosis," *Int J Imaging Syst Technol*, vol. 34, no. 1, p. e22975, Jan. 2024, doi: 10.1002/IMA.22975;WGROU:STRING:PUBLICATION.
- [127] Y. Roy, H. Banville, I. Albuquerque, A. Gramfort, T. H. Falk, and J. Faubert, "Deep learning-based electroencephalography analysis: a systematic review," *J Neural Eng*, vol. 16, no. 5, p. 051001, Aug. 2019, doi: 10.1088/1741-2552/AB260C.
- [128] A. Mao, M. Mohri, and Y. Zhong, *Cross-Entropy Loss Functions: Theoretical Analysis and Applications*. PMLR, 2023. Accessed: May 09, 2025. [Online]. Available: <https://proceedings.mlr.press/v202/mao23b.html>
- [129] Y. Bengio, I. Goodfellow, and A. Courville, *Deep Learning*. 2015.
- [130] J. D. Kelleher, *Deep Learning*. MIT press, 2019.
- [131] S. Theodoridis and K. Konstantinos, *Pattern Recognition*. Elsevier, 2006. Accessed: May 09, 2025. [Online]. Available: <https://books.google.com/books?hl=en&lr=&id=gAGRCmp8Sp8C&oi=fnd&pg=PP1&dq=pattern+recognition&ots=pKR5YNPeTK&sig=IbVTnxSrod7DdcvF9XVWvknCSZc#v=onepage&q=pattern%20recognition&f=false>
- [132] J. Howard, "Imagenette: A smaller subset of 10 easily classified classes from Imagenet [Data set]. GitHub." Accessed: May 12, 2025. [Online]. Available: <https://github.com/fastai/imagenette#imagenette-1>