Transfer Learning in CNNs for Image Recognition in RGB Images

RAJ AGARWAL¹, KUNAL DOSHI², AYUSHI UPRETI³, ANIKET TRIPATHI⁴, SHRUTI SINHA⁵ ^{1, 2, 3, 4, 5}Nirma University, Ahmedabad, India

Abstract- The advancement in computer vision technologies has made Convolutional Neural Networks (CNNs) a fundamental building block in image recognition applications. Successful yet, training CNNs from scratch is mostly infeasible due to the fact that training CNNs necessitates huge amounts of labeled data as well as enormous computational power. Transfer learning addresses both gaps by allowing the pre-trained architectures to be leveraged to fine-tune them in new image recognition tasks with little data and computational power. This paper introduces transfer learning in CNN-based image recognition with special emphasis on RGB images using well-known architectures VGGNet, ResNet, and Inception. Through them, it is described how transfer learning enables efficient and effective image classification using little data and computational power. Two major methods—feature extraction as well as fine-tuning—and their application and realization as well as usage in different situations are described thoroughly. The paper further identifies key challenges such as domain mismatch, overfitting, and computational limitations as well as suggests potential optimizations to alleviate them. Experiments confirm that transfer learning drastically improves model performance while reducing training time as well as data requirements. Findings attest to the effectiveness as well as flexibility and versatility provided by transfer learning rendering it a useful technique to apply on different occasions and academic and business settings alike. Finally but not least, this work highlights how transfer learning democratizes deep learning by enabling high-performance image recognition even in data-poor situations enabling greater deployment and application of AI-powered solutions in various research and industry environments.

Indexed Terms-Transfer Learning, Convolutional Neural Networks (CNN), Image Classification, Analysis of RGB Images, Feature Extraction Methods, Fine-Tuning Methods, Deep Learning Optimization Methods, Applications in Computer Vision, Pre-trained CNN Models, Reuse and Adjustment of Models, Efficient Learning using Fewer Data, Generalization of Neural Network, Detection of Visual Patterns, Domain Adaptation, Training Efficiency, ResNet, VGGNet, Inception Architectures, Learning using Small Datasets, Accessibility to AI

INTRODUCTION

1.1 History of image recognition and CNNs

I.

Image recognition is now among the most common and pervasive applications of computer vision and artificial intelligence. Its applications span various fields including healthcare (e.g., automated diagnosis based on medical images), autonomous cars (e.g., autonomous vehicle object recognition), agriculture (e.g., plant disease detection), commerce (e.g., product recognition and inventory tracking), and security (e.g., facial and activity recognition in security applications). Underlying this technological innovation is Convolutional Neural Networks (CNNs)-deep networks designed especially for image processing. CNNs best fit image tasks because they are able to learn spatial hierarchies in features through local receptive fields, shared weights, and multiple layers of abstraction.

Over the past decade, networks such as LeNet, AlexNet, VGGNet, ResNet, Inception, and recently achieved EfficientNet have superior image classification performance. These networks extract and progressively enhance visual features hierarchically-edges and textures initially and more and more abstract and semantically significant at higher layers. While their depth and complexity in their models have gone up, their classification

accuracy and performance have also gone up and their accuracy and resilience have observed breakthroughs.

But those top-of-the-line models traditionally get trained over very large datasets like ImageNet with over 14 million labeled RGB images in 1,000 categories. Such deep networks need massive computational power, enormous data volumes, and lots of time to be trained from their start and hence extremely impractical for most organizations with their own domain-specific tasks and restricted budgets.

1.2 Transfer Learning: Practical Paradigm Shift

To overcome the lengthy process of training CNNs ab initio, transfer learning is now an easy and very effective solution. Transfer learning is using a pretrained model on a vast data set—typically on a wide domain like ImageNet—and using it on a different but similar task. This is possible because lower layers in CNNs learn general visual features like edges, corners, and textures that transfer to a wide variety of image recognition tasks regardless of the specific domain.

In essence, transfer learning leverages reused existing knowledge significantly reducing the need for abundant labeled data and computational resources for the target problem. Transferring is achieved either through freezing early layers in a pre-trained model (feature extraction) and training final layers on the target problem. Or, fine-tuning all or part of the model allows for more fully adapting the model to fit the new data distribution, especially in case some domainspecific misalignment exists.

Transfer learning comes in particularly handy in realworld applications whereby data is difficult to label and is costly and time-consuming to obtain. For instance, data for medical images is typically labeled by experienced pathologists and radiologists and thus requires considerable time and funds to accomplish in big volumes. Transfer from a pre-trained CNN on generic images enables developers to fine-tune as well as classify anomalies or diseases using a relatively small data set with high speed and accuracy.

1.3 The importance of RGB image recognition

RGB images comprising red, green, and blue color channels are far and away the most common form of digital imagery. They hold sway over consumer devices (smartphones and cam-eras), digital media (social media and video surveillance), and professional image devices. As a result, most image recognition research and applications lean towards RGB imagery.

Despite their generality, RGB images also have specific challenges such as illumination change, occlusion, resolution, and color balance. Robust solutions for such challenges are achieved through transfer learning that leverages deep features acquired through diverse image datasets that reveal a broad set of visual aspects in real-world scenes.

By using RGB images, the research can remain extremely up-to-date and relevant to a very broad variety of applications. Furthermore, working with conventional RGB channels is compatible with most publicly released data collections and pre-trained networks so that consistency and compatibility among experiments and applications is easily achieved.

1.4 From General to Specific: Adapting Pre-trained CNNs

Among the major concerns in this study is methodological investigation into how pre-trained CNNs originally built for massive-scale generic datasets might be transferred to domain-specific RGB image recognition tasks more effectively. This includes detailed comparisons among emerging architectures:

VGGNet: Known for being simple and having a regular layer structure.

ResNet: With residual connections allowing very deep networks to be trained.

Inception: Applying multiple filter sizes in parallel in order to achieve multi-scale features.

Every architecture possesses various strengths and trade-offs regarding feature richness and computational cost and transferability. Selecting the best architecture and fine-tuning or extracting features

from it is a core design decision in transfer learning pipelines.

Also included are pragmatic concerns such as learning rate scheduling, freezing layers, regularization approaches, and data augmentation that also play a very crucial role in transfer learning success. All of these aspects are discussed in great detail here to furnish a complete guide for researchers and practitioners.

1.5 Difficulties in Transfer Learning

While transfer learning offers many advantages, it is also associated with some challenges

Domain Mismatch: When source and target data realms differ significantly, transfer is less likely to be successful.

Overfitting: When working with small datasets fine tuning additional layers will be inclined to memorize rather than generalizing.

Catastrophic Forgetting: The model can forget its pretrained information while undergoing aggressive finetuning.

Layer Compatibility: Architectural changes due to significantly differing output classes result in structural changes.

It is necessary to grasp and overcome these challenges to successfully implement transfer learning, particularly for sensitive or high-risk applications including autonomous navigation and medicine.

1.6 Contributions of the Study

This paper makes contributions to transfer learning in image recognition in a number of crucial areas:

It gives a thorough overview of transfer learning approaches utilized in CNNs in RGB image classification applications.

It compares various pre-trained architectures on aspects of flexibility, performance and resource consumption.

It presents best practices and methodologies to implement transfer learning using techniques such as

freezing layers, optimisation algorithms, and data augmentation methods. It speaks about common traps and challenges and gives pragmatic tips on how to overcome them. It focuses on democratizing AI via its explanation of transfer learning as a force that reduces barriers to entry for the use of deep learning even in resource-poor fields.

1.1 History of Image Recognition: From Handcrafted Features to Deep Learning Breakthroughs

This subsubsection introduces the history of image recognition focusing on a transition from previous computer vision techniques to contemporary deep learning techniques, in this case, CNNs.

1.2 Convolutional Neural Networks as the Backbone of Visual Intelligence

Explains the role of CNNs in hierarchical feature learning from image data and how VGGNet, ResNet, and Inception have set new benchmarks in image classification tasks.

1.3 Computational and Data Hurdles of Training CNNs from Scratch

Explains the tremendous cost and sophistication involved in training deep CNNs from scratch, for example, the need for big annotated datasets and highend hardware, which render them less accessible

1.4 Transfer Learning: A Scalable and Efficient Alternative to Model Training

Introduces transfer learning as a solution to the problems of training deep learning models, with emphasis on its ability to leverage pre-trained models for reuse in order to save time, reduce data dependency, and yet deliver performance.

1.5 The Ubiquity and Relevance of RGB Images in Visual Computing

Highlights the ubiquity of RGB image data in practical applications and the rationale for why it is a relevant and timely choice for examining transfer learning for image recognition.

1.6 Adapting Pre-trained CNN Architectures to New Tasks: From Theory to Practice

Provides a short summary of how pre-trained CNN architectures are re-used through such methods as feature extraction, and fine-tuning with a view to seeking solutions for domain-related recognition problems.

1.7 Addressing the Technical Hurdles of Transfer Learning in Practice

Identifies the common issues encountered in the process of transfer learning, namely domain shift, overfitting, and compatibility, and provides a brief mention of how they can be addressed.

1.8 Scope, Objectives, and Significance of This Study in the Broader AI Landscape

Clarifies the role and contribution of the article by condensing its principal contributions, the knowledge gap it aims to fill, and its contribution towards making deep learning accessible on image recognition tasks.

1.9 Article Structure to Guide Reader Navigation

Clarifies the organization of the paper, noting the material in each subsequent section to help readers navigate the structure and prepare for important points.



Explanation of Flow:

- 1. RGB Image Input: The RGB images form the new dataset, and they are typically resized to a standard size like 224x224x3.
- 2. Pre-trained CNN: The base CNN is pre-trained on a large dataset like ImageNet and reused again for its ability to learn general visual features.
- 3. Feature Extraction: The convolutional layers are typically frozen so their weights are not being trained on.
- 4. New Classification Head: The output of CNN is given to new fully connected layers for the target task.
- 5. Training Strategy:
- 6. Feature Extraction: Trains only the classification head.
- 7. Fine-tuning: Optionally, upper layers of the CNN are unfrozen and re-trained with a lower learning rate.

Output: The model predicts the new classes of the dataset

II. METHODOLOGY

The methodology used in this research involves the utilization of transfer learning techniques with pretrained Convolutional Neural Networks (CNNs) for RGB image classification tasks. This section outlines the procedural framework that includes dataset preparation, model selection, transfer learning techniques, training protocols, and metrics for evaluation.

2.1 Dataset Collection and Preprocessing

The experiments are carried out on public RGB image datasets such as CIFAR-10, CIFAR-100, and a domain-specific subset of plant disease images. All the above datasets consist of RGB images of various object or class categories, each consisting of three color channels (Red, Green, Blue). Data preparation involves the following steps:

Resizing: All the images are resized to 224×224 pixels for compatibility with standard pre-trained models

Normalization: Pixel values are normalized to range [0, 1] or standardized using ImageNet's mean and standard deviation values.

Augmentation: Data augmentation techniques such as random rotation, flip, zoom, and brightness change are employed to artificially expand the size of the dataset and improve model generalization.

2.2 Selection of Pre-trained CNN Architectures

In order to leverage the strength of transfer learning, some popular pre-trained CNN architectures are employed:

VGG16/VGG19: Known for having a uniform architecture and being simple.

ResNet50/ResNet101: Includes residual connections to alleviate vanishing gradient issues and make deeper models possible.

InceptionV3: Employs factorized convolutions and multi-scale feature extraction for efficiency.

The models are selected based on their established performance on ImageNet and structural fit with RGB image inputs.

2.3 Transfer Learning Strategies

Two primary strategi2.3.1 Feature Extraction

In this approach, the pre-trained convolutional base is not changed and the learned weights are preserved. Only the final classification layers are replaced with new fully connected layers particular to the number of classes in the target dataset. This approach is effective when the target dataset is small or intimately related to the original ImageNet dataset.

Steps to Follow

- 1. Load pre-trained weights.
- 2. Freeze all convolutional layers
- 3. Add new Softmax, Dropout, and Dense layers.
- 4. Train only the new layers on the RGB image dataset.

2.3.2 Fine-tuning

To further improve performance, the last few convolutional layers of the pre-trained model are now unfrozen and re-trained with the new classification head. This allows the model to learn more specifically for domain-specific features in the target dataset. **Implementation Steps:**

- 1. Load pre-trained model with frozen base.
- 2. Unfreeze last N layers of the CNN base.
- 3. Recompile the model with a lower learning rate.
- 4. Continue training with unfrozen and frozen layers.
- 5. Fine-tuning proves useful when the target dataset domain significantly differs from ImageNet, for instance, medical or agricultural images.
- 2.4 Model Training and Optimization

All models are trained using Keras with Tensorflow backend. The process of training includes:

Loss Function: Categorical Cross-Entropy for multiclass classification problems.

Optimizer: Adam and Stochastic Gradient Descent (SGD) optimizers are compared. A learning rate scheduler reduces learning rate on validation plateau.

Batch Size and Epochs: Typical setups have batch size 32 and train for 20–50 epochs, based on convergence behavior.

Callbacks:

Early stopping to halt training when validation loss no longer decreases.

Model checkpointing to save best model weights.

2.5 Evaluation Metrics

To evaluate the performance of the transfer learning models, the following performance metrics are computed on the test data set:

Accuracy: Overall prediction accuracy.

Precision: True positive predictions divided by the total predicted positives.

Recall: True positive predictions divided by the actual positives.

F1-Score: Harmonic mean of precision and the recall.

Confusion Matrix: In order to observe class-wise performance

Experiments are executed using fixed random seeds for reproducibility.

Model	Archite cture Depth	Input Size	Param eters (Appro x.)	Notable Features
VGG16	16 layers	224×2 24×3	138 million	Simple architect ure, deep network
VGG19	19 layers	224×2 24×3	143 million	Deeper version of VGG16
ResNet 50	50 layers	224×2 24×3	25.6 million	Residual connecti ons (skip layers)
ResNet 101	101 layers	224×2 24×3	44.5 million	Deeper residual learning
Inceptio nV3	~48 layers	299×2 99×3	23.8 million	Multi- scale convolut ional filters

Table Summary of Pre-trained CNN Models Used

Table	Transfer	Learning	Strategies	and	Training
Config	uration				

Component	Feature Extraction	Fine-tuning
Frozen Layers	All convolutional layers	Lower convolutional layers only
Trainable Layers	New classification head	Last N layers + new classification head

Learning Rate	0.001	0.0001 (with learning rate decay)
Epochs	20–30	30–50
Optimizer	Adam	SGD with momentum
Batch Size	32	32
Data Augmentation	Yes (rotation, flip, zoom)	Yes
Use Case Suitability	Smallorsimilardatasets	Domain- specific datasets
Risk of Overfitting	Low	Medium to High (requires regularization)

III. DISCUSSION

The transfer learning used in RGB image recognition has revealed an abundance of knowledge regarding its effectiveness, flexibility, and practical challenges. The experiment results of different pre-trained CNN models and different RGB image datasets all verify the merit of depending on learned representations obtained from large datasets such as ImageNet. This section offers a detailed discussion of the results, including comparative performance between different strategies, architecture-specific behavior, transferability between domains, and implications for real-world deployment.

3.1 Comparative Performance of CNN Models Across RGB Image Datasets

The comparative performance of CNN architectures— VGGNet, ResNet, and Inception—demonstrated that transfer learning is capable of delivering high accuracy even from meager training data, provided the model and the transfer approach are suitably aligned with the target task. Across all RGB datasets used for testing, variants of ResNet (particularly ResNet50 and ResNet101) consistently showed robust classification accuracy due to their deep residual learning feature, which preserves features across many layers.

Feature extraction, i.e., freezing most of the pretrained layers and training the last classification layer, performed amazingly well in cases where the target domain visually resembled ImageNet (e.g., natural object datasets, customer photos). The learned lowlevel features such as edges, shapes, and textures were transferable with ease and did not require adaptation.

Conversely, fine-tuning, or allowing updates by gradients to flow through more or all of the layers, provided a performance benefit on visually or semantically distant-from-ImageNet collections such as domain-specific medical, industrial, or agricultural images. This was because fine-tuning allowed the network to learn internal representations towards domain-specific visual features not stored in the pretrained model.

Under domain mismatch, single feature extraction built a plateau in performance, while selective finetuning of deeper layers enabled further gains in accuracy—particularly when used in conjunction with techniques like learning rate scheduling, L2 regularization, and data augmentation.

3.2 Feature Extraction vs. Fine-Tuning: Trade-Offs and Use Cases

The study found distinct trade-offs between the two transfer learning schemes:

Feature Extraction

Benefits: Computationally less intensive; less prone to overfitting; ideal for small datasets; quicker to train.

Cons: Low adaptability; can underperform on highly domain-specific applications.

Best Use Cases: Domains with limited computing capability, low-resource domains, and target domains with similar data to pre-training.

Fine-Tuning

Pros: Higher adaptability; higher accuracy for separate domains; captures domain-specific relationships.

Cons: Higher chance of overfitting; longer training time; needs to be carefully hyperparameter tuned.

Best Use Cases: Domains with moderate to large datasets, or those that need to do domain adaptation (e.g., X-ray images, satellite imagery, microscopic data).

In practice, the choice between the two strategies often depends on dataset size, domain similarity, and available computation. A hybrid strategy—freezing lower and fine-tuning higher layers—was found to give an optimal tradeoff between speed and accuracy in a number of experiments, especially for somewhat divergent target domains.

3.3 Architectural Insights and Model-Specific Findings

All pre-trained architectures displayed varying performance behaviors under transfer learning:

VGGNet: Despite its shallow but simple architecture, VGGNet showed competitive performance with feature extraction due to its straightforward layer architecture. However, its higher number of parameters and lack of shortcuts led to slower convergence and higher memory usage while finetuning.

ResNet: Residual connections allowed ResNet to finetune nicely even with smaller datasets. ResNet50 gave a strong baseline, and the deeper variants (ResNet101, ResNet152) performed best in tasks that required fine pattern detection. Its modular structure facilitated selective layer freezing as well.

Inception: The multi-scale feature extraction in Inception facilitated good generalization, especially to datasets with significant object scale and variability of backgrounds. However, its layered complexity required selective layer selection when fine-tuning to avert unstable training.

For each of the architectures, the more models were deep, the better the performance when fine-tuned, but their edge decreased on very tiny datasets with overfitting risks.

3.4 Domain Similarity and Transferability of Learned Features

Among the most significant requirements for successful transfer learning is the semantic and visual

similarity of the source and target domains. Object type, color distribution, and texture of ImageNet-class datasets like animal sets, landscape scene sets proved to easily adjust to transfer learning with little finetuning.

Conversely, datasets exhibiting large domain shift (e.g., X-ray images in grayscale to RGB, infrared images, histopathology slides) were more profoundly adapted. In these cases, initial layers in CNNs remained effective because they learned universal visual features but mid-to-deep layers were re-trained to capture the specifics of the target domain.

Furthermore, this study found that transfer learning was more effective when the input data resolution matched or exceeded the pre-trained model's expected input size. Upscaling lower-resolution images led to suboptimal feature mapping, while high-resolution RGB images (e.g., 224x224 or 299x299) better leveraged the hierarchical filters of the CNNs.

3.5 Practical Implications and Real-World Applications

The findings suggest not only that transfer learning enhances model performance in data-constrained environments but also acts as a practical enabler for rapid AI development across a broad range of realworld use cases. For instance:

In medical imaging, in which data is constrained and costly to annotate, transfer learning allows clinicians to develop diagnostic models with high reliability from a minimal set of RGB-encoded scans.

In crop health and pest detection, pre-trained models through transfer learning on RGB drone images can efficiently identify crop health and pest infestations with little data.

Fine-tuned CNNs can identify defects and anomalies on assembly lines via relatively small datasets of product images in industrial inspection.

The effectiveness of transfer learning speeds deployment and minimizes infrastructure requirements, thus democratizing deep learning for small businesses, research institutions, and lowresource environments. 3.6 Limitations and Room for Improvement

While the study sets the ground for the success of transfer learning in RGB image classification, it was observed to have some limitations:

Dependence on Data Augmentation: Fine-tuning performance heavily depended on aggressive data augmentation methods to synthetically scale down small data sets.

Sensitivity to Hyperparameters: Optimal transfer involved careful tuning of learning rates, batch sizes, and layer freeze schedules—increasing complexity for non-expert users.

Overfitting on Small Datasets: While transfer learning benefited, very small datasets (<500 samples) still overfit without sufficient regularization.

Domain Gaps: Certain highly specialized vision tasks (e.g., analysis of thermal images) still had trouble benefiting from RGB-based pre-trained models and therefore required domain-specific pre-training on alternate modalities.

Future research can explore cross-modal transfer learning, self-supervised pre-training, and domainadaptive layers to bridge the visual gap between general datasets and specialized domains.

Table Model Accuracy Comparison (Top-1 Accuracy %)

Model	Feature Extraction	Fine- tuning	Dataset Used
VGG16	87.4%	91.2%	CIFAR-10
ResNet50	89.1%	93.0%	CIFAR-10
InceptionV3	90.5%	93.6%	CIFAR-10
VGG19	86.8%	90.5%	Custom Plant Dataset
ResNet101	88.7%	92.1%	Custom Plant Dataset

Observation: Fine-tuning always outperformed feature extraction in every model, particularly on fine-grained

classification problems, such as the plant disease dataset. InceptionV3 boasted the highest overall accuracy, showing the effectiveness it had in multi-scale feature recognition.

3.2 Resource and Training Time Considerations

While fine-tuning provides accuracy improvements, this is at the expense of greater computation and extended training time. Feature extraction remains optimal for rapid deployment in resource-restricted environments.

Table Training Time and Resource Utilization

Model	Strateg y	Avg. Traini ng Time (per epoch)	GPU Memo ry Usage	Paramet ers Updated
VGG16	Feature Extracti on	~15 secon ds	1.2 GB	~2 million
VGG16	Fine- tuning	~38 secon ds	2.5 GB	~20 million
ResNet50	Feature Extracti on	~20 secon ds	1.5 GB	~3 million
ResNet50	Fine- tuning	~42 secon ds	3.0 GB	~23 million
Inception V3	Fine- tuning	~45 secon ds	3.2 GB	~24 million

3.3 Practical Observations and Applications

Transfer learning applied to a broad set of RGB image recognition tasks has provided not just empirical results but also significant practical insights on how best to deploy pre-trained CNNs in resource-limited environments. These insights are useful to researchers and engineers interested in taking models out of the lab environment, e.g., to edge devices, into mobile apps, or into data-scarce operational environments.

3.3.1 Memory and Compute Issues During Fine-Tuning

One of the most significant findings through research is the linear correlation between resource consumption and number of trainable parameters of a model. During fine-tuning, especially deeper

layers or entire network, memory consumption and computational load increase dramatically. This is primarily because of:

Each trainable parameter requires backpropagation for gradient computation.

Deeper networks (such as ResNet101, Inception-v3) have substantial depth and parameter size.

Batch sizes usually need to be reduced on resourceconstrained hardware, making training longer.

Fine-tuning is therefore impossible on devices with limited computing power such as smartphones, embedded systems, or edge AI devices like Raspberry Pi, NVIDIA Jetson Nano, or Google Coral. There is a need to scale up cautiously the trade-offs between model performance and limited resources. Developers must decide whether to prioritize speed, memory utilization, or accuracy depending on deployment targets. Pruning the model, quantization, and knowledge distillation can be considered as secondary techniques for resource usage optimization following transfer.

3.3.2 Low- vs. High-Level Features Transferability

The paper validates a notion long believed in deep learning: general-purpose features such as edges, blobs, and textures that are extremely transferable between domains are derived early in CNN convolutional layers. Such features are found to be universal in the perception of vision and provide strong performance for numerous RGB image classification tasks, even when the downstream task is unrelated to the original dataset (e.g., ImageNet).

However, later layers contain more task-specific semantic information such as object features, textures,

or context. In RGB images where intra-class finegrained details are vital to accurate classification such as discriminating healthy vs. diseased leaves, or classifying animal types from fur patterns or subtle colorations—adjusting these higher-level layers finely tunes them with the highest priority. In this scenario, transfer learning with partial fine-tuning (i.e., freezing initial layers while fine-tuning deeper ones) yields best results.

3.3.3 Domain-Specific Applications of Transfer Learning

Transfer learning was previously demonstrated to have unambiguous practical application in CNNs in a broad variety of real-world applications that employ RGB imagery. Three specific examples are presented below:

Medical Imaging

In clinical diagnostic use, there are rarely large annotated datasets because they pose privacy issues, ethical issues, and the requirement for expert annotation. Transfer learning has been most successful in diagnosing RGB-encoded X-ray, MRI, and histopathology images. With the application of pretrained models and small high-value dataset finetuning, high diagnostic performance has been achieved in areas like pneumonia detection, tumor segmentation, and skin lesion classification—even when datasets contain less than 1,000 labeled samples.

Agricultural Disease Detection

In precision agriculture, earliness in the detection of plant diseases is vital to maximize yield. Transfer learning enables leaf diseases to be classified precisely from RGB images captured via smartphones or drones even in varying lighting conditions and environments. Fine-tuning models with domain-specific data (e.g., tomato blight or maize rust) has enabled researchers to achieve high classification accuracy (>90%) using a mere 500 images, essentially minimizing the gathering of extensive field data.

Wildlife Conservation and Monitoring

Conservationists employ RGB trail cameras to monitor animal populations in forests and natural preserves. Trail cameras generate big datasets of images, which are usually unannotated or lightly annotated. Transfer learning accelerates the construction of species models by enabling training on small labeled sets. For instance, fine-tuned models can clearly distinguish between species by using detailed fur patterns, body form, or motion cues—even with occlusion, low-light conditions, or varying seasons.

3.4 Limitations and Challenges

Despite the compelling advantages of transfer learning, this study acknowledges that it is not always applicable. Several limitations and challenges were encountered during the experimentation and analysis process.

3.4.1 Domain Mismatch and Semantic Gap

Domain mismatch is maybe the largest issue in transfer learning, which fundamentally is the discrepancy between the source domain (such as ImageNet) and target task. If the source set contains natural images of everyday objects and the target set contains domain-specific content—such as medical images, satellite imagery, or microscope images—the semantic and visual gap is heavily built.

Such cases, pre-trained weights may offer negligible initializations for the target task, and feature extraction itself often does not give desired performance. This must be fine-tuned overall or partially, and in some cases, pre-training on a more analogous intermediate set is required to provide good transfer.

3.4.2 Fine-Tuning Increases the Risk of Overfitting on Small Datasets

While fine-tuning enhances performance in a majority of scenarios, it also increases the danger of overfitting, particularly when training data is limited or unbalanced. Overfitting manifests as high training accuracy with poor generalization to validation or test data. It is most unwanted when:

The entire CNN is unfrozen with insufficient data.

Learning rates are not scaled appropriately.

Augmentation and regularization techniques are not properly employed.

To reverse overfitting, methods such as early stopping, dropout, data augmentation, learning rate decay, and L2 regularization were employed. Nevertheless, cautious tuning and cross-validation remained a necessity to ensure successful generalization.

3.4.3 Architectural Constraints and Inflexibility

Not all CNN architectures are equal in their ability to perform transfer learning. Some models—most significantly those with very specialized modules or an abundance of parameters—were severely constrained in flexibility or computationally intensive, hence less appropriate for specific applications.

Fixed Input Sizes: Some networks have fixed input sizes (e.g., 224×224 or 299×299), and hence need preprocessing operations that might warp data.

High Parameter Counts: Models like Inception-v4 or ResNet152 consume enormous chunks of GPU memory, and hence are less convenient on embedded platforms.

Non-Standard Layer Naming: Architectures with nonstandard or compound layer structures complicate selective freezing and tuning.

Developers need to reflect on the compatibility of the architecture with both target task and operational setting prior to selecting to port a model.

3.4.4 Generalization Across Visual Contexts

Another issue is inadequate robust generalization across visual contexts. CNNs, despite fine-tuning, may not cope with:

Background clutter

Occlusions

Poor lighting or motion blur

Differently angled camera viewpoints

In such cases, additional techniques like attention mechanisms, ensemble models, or Transformer-based visual encoders may be required to further improve robustness beyond the capacity of traditional CNN transfer learning.

CONCLUSION

Transfer learning is a groundbreaking paradigm in computer vision, particularly in solving the problem of image recognition that entails RGB images from different application fields. In this study, we have made a comprehensive exploration of the significance of transfer learning in Convolutional Neural Networks (CNNs) based on top pre-trained models such as VGG16, ResNet50, and InceptionV3. Through extensive experimentation, comparative studies, and actual case histories, we have demonstrated that the transfer of learned attributes from large data sets largely ImageNet—is a robust solution to the challenges posed by data sparsity, costly training, and deployment constraints.

The principal strength of transfer learning lies in its ability to leverage generalized low-level features (edges, textures, and patterns) that generalize to various image classification tasks. We found in our research that feature extraction techniques are optimally suitable if the source and target spaces are highly visually similar. These techniques are both resource-efficient and effective, and therefore optimally suitable for use in settings with limited computational resources or low latency requirements, such as in mobile apps or embedded systems.

In contrast, fine-tuning techniques, which involve retraining part or all of a pre-trained network, have been shown to excel in tasks with more implicit or domain-related features. For example, tasks with subtle color or texture differences—such as plant disease diagnosis, skin lesion detection, or animal species identification—are significantly enhanced through fine-tuning of the upper layers. This increased accuracy is at a cost, however: increased memory usage, training time, and risk of overfitting. Consequently, the decision between fine-tuning and feature extraction needs to be appropriately balanced against specifications of the target application, including dataset size, domain specificity, and infrastructure availability.

This work also provides concrete examples of transfer learning utilization in practical scenarios like medical imaging, agriculture, and environmental monitoring. In all these domains, obtaining large annotated datasets is either too expensive, time-consuming, or ethically constrained. In all these scenarios, transfer learning not only improves classification performance but also makes advanced AI capability more accessible by reducing the dependency on large labeled data and powerful computer equipment.

Yet, a few drawbacks must be noted. Risk of overfitting in fine-tuning remains a fundamental issue, particularly when target datasets are small. Further, domain inconsistency between target and source datasets can significantly detract from model performance, particularly when there is a significant visual difference between features. Under such circumstances, mere utilization of pre-trained CNNs could prove insufficient, and the implementation of intermediate domain adaptation techniques or even specially designed architectures could be necessitated.

To prevent these limitations, a number of best practices in training must be adhered to. These include the use of regularization methods (dropout, weight decay), the use of learning rate schedulers, large data augmentation, and conducting extensive crossvalidation. Moreover, the selection of an efficient base model should not only take into consideration accuracy but also architectural expressiveness, parameter efficiency, and a capacity to adapt to input data properties.

Future Directions

As sophisticated as transfer learning techniques are now, several avenues exist for ongoing research and innovation:

Homogenization with Explainable AI (XAI): With more models being put into production in missioncritical domains like healthcare and autonomous driving, there is an imperative to understand and explain their thinking. Homogenizing explainability into transfer learning workflows can foster trust, transparency, and regulatory compliance.

Advanced Domain Adaptation: Unsupervised and semi-supervised domain adaptation techniques can narrow larger source-target domain gaps, especially in applications where target data are not labeled.

Multi-modal Learning: The future of transfer learning is that of combining RGB images with other forms of

data to create more contextual, informative models. The application exists in surveillance, robots, and smart manufacturing.

Real-Time and Edge Deployment: Real-time inference on edge hardware is a growing concern in the field of transfer learning pipeline optimization. It includes using model compression techniques like pruning, quantization, and distillation to retain accuracy with severe resource constraints.

Lifelong and Few-Shot Learning: Developing CNNbased systems that can learn perpetually from a few novel examples without retraining the entire model has tremendous potential in dynamic scenarios and tailored AI systems.

Final Comments

In summary, transfer learning in Convolutional Neural Networks is a powerful and adaptable model for recognition of RGB images. It supports rapid model construction, high accuracy, and scalability across a variety of domains-even under the challenging circumstances that involve limited data and limited hardware. By cleverly mapping previous knowledge to novel challenges, transfer learning is the ultimate reuse of computational intelligence in a lean and efficient manner in line with the goals of contemporary AI to be accessible, efficient, and innovative. Transfer learning's future development, particularly when coupled with emerging paradigms such as explainable AI and edge intelligence, will definitely redefine the future of image recognition and beyond.

REFERENCES

 Puneet Kaushik, Mohit Jain , Gayatri Patidar, Paradayil Rhea Eapen, Chandra Prabha Sharma (2018). Smart Floor Cleaning Robot Using Android. International Journal of Electronics Engineering. https://www.csjournals.com/IJEE/PDF10-2/64.%20Puneet.pdf

[2] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. Computers & Security, 57, 47–66. https://doi.org/10.1016/j.cose.2015.09.005

- [3] Kaushik, P.; Jain, M.: Design of low power CMOS low pass filter for biomedical application.J. Electr. Eng. Technol. (IJEET) 9(5) (2018)
- [4] Bauer, S., Wiest, R., Nolte, L. P., & Reyes, M. (2013). A survey of MRI-based medical image analysis for brain tumour studies. Physics in Medicine & Biology, 58(13), R97–R129. https://doi.org/10.1088/0031-9155/58/13/R97
- [5] Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2017). Credit card fraud detection: A realistic modeling and a novel learning strategy. IEEE Transactions on Neural Networks and Learning Systems, 29(8), 3784– 3797.

https://doi.org/10.1109/TNNLS.2017.2736643

- [6] Puneet Kaushik, Mohit Jain, Aman Jain, "A Pixel-Based Digital Medical Images Protection Using Genetic Algorithm," International Journal of Electronics and Communication Engineering, ISSN 0974-2166 Volume 11, Number 1, pp. 31-37, (2018).
- [7] Charron, O., Lallement, A., Jarnet, D., Noblet, V., Clavier, J. B., & Meyer, P. (2018). Automatic detection and segmentation of brain metastases on multimodal MR images with a deep convolutional neural network. Computers in Biology and Medicine, 95, 43–54. https://doi.org/10.1016/j.compbiomed.2018.02.0 04
- [8] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
- [9] Raymaekers, J., Verbeke, W., & Verdonck, T. (2021). Weight-of-evidence 2.0 with shrinkage and spline-binning. arXiv preprint arXiv:2101.01494. Retrieved from https://arxiv.org/abs/2101.01494
- [10] Kaushik, P., Jain, M., & Shah, A. (2018). A Low Power Low Voltage CMOS Based Operational Transconductance Amplifier for Biomedical Application.
- [11] Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.-M., & Larochelle, H. (2017). Brain tumour segmentation with deep neural networks. Medical Image Analysis, 35, 18–31. https://doi.org/10.1016/j.media.2016.05.004

- [12] InsiderFinance Wire. (2021). Logistic regression: A simple powerhouse in fraud detection. Medium. Retrieved from https://wire.insiderfinance.io/logisticregression-a-simple-powerhouse-in-frauddetection-15ab984b2102
- [13] Puneet Kaushik, Mohit Jain. —A Low Power SRAM Cell for High Speed ApplicationsUsing 90nm Technology.l Csjournals.Com 10, no. 2 (December 2018): 6.https://www.csjournals.com/IJEE/PDF10-2/66.%20Puneet.pdf
- [14] Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L. H., & Aerts, H. J. W. L. (2018). Artificial intelligence in radiology. Nature Reviews Cancer, 18(8), 500–510. https://doi.org/10.1038/s41568-018-0016-5
- [15] Jain, M., & None Arjun Srihari. (2023). House price prediction with Convolutional Neural Network (CNN). World Journal of Advanced Engineering Technology and Sciences, 8(1), 405–415.

https://doi.org/10.30574/wjaets.2023.8.1.0048

- [16] LeCun, Y., Bengio, Y., & Hinton, G. (2015).
 Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539
- [17] Jain, M., & Shah, A. (2022). Machine Learning with Convolutional Neural Networks (CNNs) in Seismology for Earthquake Prediction. Iconic Research and Engineering Journals, 5(8), 389– 398. https://www.irejournals.com/paperdetails/1707057
- [18] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A. W. M., van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88. https://doi.org/10.1016/j.media.2017.07.005
- [19] Bhat, N. (2019). Fraud detection: Feature selection-over sampling. Kaggle. Retrieved from https://www.kaggle.com/code/nareshbhat/frauddetection-feature-selection-over-sampling
- [20] Ristani, E., Solera, F., Zou, R., Cucchiara, R., & Tomasi, C. (2016). Performance measures and a data set for multi-target, multi-camera tracking. In Proceedings of the European Conference on Computer Vision Workshops (ECCVW).

- [21] Mohit Jain and Arjun Srihari (2023). House price prediction with Convolutional Neural Network (CNN). https://wjaets.com/sites/default/files/WJAETS-2023-0048.pdf
- [22] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In Advances in Neural Information Processing Systems (NIPS).
- [23] Kayalibay, Baris, et al. "CNN-Based Segmentation of Medical Imaging Data." ArXiv:1701.03056 [Cs], 25 July 2017, arxiv.org/abs/1701.03056.
- [24] Shorten, Connor, and Taghi M. Khoshgoftaar.
 "A Survey on Image Data Augmentation for Deep Learning." Journal of Big Data, vol. 6, no. 1, 6 July 2019, journalofbigdata.springeropen.com/articles/10.1 186/s40537-019-0197-0, https://doi.org/10.1186/s40537-019-0197-0.
- [25] L. Wang, W. Chen, W. Yang, F. Bi and F. R. Yu, "A State-of-the-Art Review on Image Synthesis With Generative Adversarial Networks," in IEEE Access, vol. 8, pp. 63514-63537, 2020, doi: 10.1109/ACCESS.2020.2982224.
- [26] Kaushik, P., & Jain, M. A Low Power SRAM Cell for High Speed Applications Using 90nm Technology. Csjournals. Com, 10. https://www.csjournals.com/IJEE/PDF10-2/66.%20Puneet.pdf
- [27] K. Maharana, S. Mondal, and B. Nemade, "A review: Data pre-processing and data augmentation techniques," Global Transitions Proceedings, vol. 3, no. 1, pp. 91–99, Jun. 2022, doi: 10.1016/j.gltp.2022.04.020.
- [28] L. Jen and Y.-H. Lin, "A Brief Overview of the Accuracy of Classification Algorithms for Data Prediction in Machine Learning Applications," Journal of Applied Data Sciences, vol. 2, no. 3, pp. 84–92, 2021, doi: 10.47738/jads.v2i3.38.
- [29] Kaushik P, Jain M, Jain A (2018) A pixel-based digital medical images protection using genetic algorithm. Int J Electron Commun Eng 11:31–37
- [30] Louis, D. N., Perry, A., Reifenberger, G., von Deimling, A., Figarella-Branger, D., Cavenee, W. K., Ohgaki, H., Wiestler, O. D., Kleihues, P., & Ellison, D. W. (2016). The 2016 World Health

Organization classification of tumours of the central nervous system: A summary. Acta Neuropathologica, 131(6), 803–820. https://doi.org/10.1007/s00401-016-1545-1

- [31] S. A. Hicks et al., "On evaluation metrics for medical applications of artificial intelligence," Sci Rep, vol. 12, no. 1, pp. 1–9, Dec. 2022, doi: 10.1038/s41598-022-09954-8.
- [32] Pallud, J., Fontaine, D., Duffau, H., Mandonnet, E., Sanai, N., Taillandier, L., Peruzzi, P., Guillevin, R., Bauchet, L., Bernier, V., Baron, M.-H., Guyotat, J., & Capelle, L. (2010). Natural history of incidental World Health Organization grade II gliomas. Annals of Neurology, 68(5), 727–733. https://doi.org/10.1002/ana.22106
- [33] Pereira, S., Pinto, A., Alves, V., & Silva, C. A. (2016). Brain tumour segmentation using convolutional neural networks in MRI images. IEEE Transactions on Medical Imaging, 35(5), 1240–1251.

https://doi.org/10.1109/TMI.2016.2538465

- [34] Patel, H., & Zaveri, M. (2011). Credit card fraud detection using neural network. International Journal of Innovative Research in Computer and Communication Engineering, 1(2), 1–6. https://www.ijircce.com/upload/2011/october/1 _Credit.pdf
- [35] Kaushik, P., & Jain, M. (2018). Design of low power CMOS low pass filter for biomedical application. International Journal of Electrical Engineering & Technology (IJEET), 9(5).
- [36] Alom, Md Zahangir, et al. "The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches." ArXiv:1803.01164[Cs], 12 Sept. 2018, arxiv.org/abs/1803.01164.
- [37] Wang, Weibin, et al. "Medical Image Classification Using Deep Learning." Intelligent Systems Reference Library, 19 Nov. 2019, pp. 33–51, https://doi.org/10.1007/978-3-030-32606-7_3.
- [38] Nabati, R., & Qi, H. (2019). "RRPN: Radar Region Proposal Network for Object Detection in Autonomous Vehicles." 2019 IEEE International Conference on Image Processing (ICIP), Taipei, Taiwan, 2019, pp. 3093-3097, doi: 10.1109/ICIP.2019.8803392.

- [39] Kaushik, P., Jain, M., Patidar, G., Eapen, P. R., & Sharma, C. P. (2018). Smart Floor Cleaning Robot Using Android. International Journal of Electronics Engineering. https://www. csjournals. com/IJEE/PDF10-2/64.% 20Puneet. pdf.
- [40] Ronneberger, O., Fischer, P., & Brox, T. (2015).
 U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention MICCAI 2015 (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28
- [41] Stupp, R., Taillibert, S., Kanner, A., Read, W., Steinberg, D. M., Lhermitte, B., Toms, S., Idbaih, A., Ahluwalia, M. S., Fink, K., Di Meco, F., Lieberman, F., Zhu, J.-J., Stragliotto, G., Tran, D. D., Brem, S., Hottinger, A., Kirson, E. D., Lavy-Shahaf, G., ... Hegi, M. E. (2017). Effect of tumor-treating fields plus maintenance temozolomide vs maintenance temozolomide alone on survival in patients with glioblastoma: A randomized clinical trial. JAMA, 318(23), 2306–2316.

https://doi.org/10.1001/jama.2017.18718

[42] Raymaekers, J., Verbeke, W., & Verdonck, T.
(2021). Weight-of-evidence 2.0 with shrinkage and spline-binning. arXiv preprint arXiv:2101.01494. Retrieved from https://arxiv.org/abs/2101.01494