Transfer Learning in CNNs: Techniques, Challenges, and Emerging Trends in Fine-Tuning Strategies

Sanskar Bhardwaj*, Shivam Singh†, Sarthak Kumar‡, Shaarvi Srivastava§, Shiva Singh¶, Saumya Shree‖, Saurav Anand**, Satyam Kumar††
Department of Information Technology, Noida Institute of Engineering and Technology, Greater Noida, India
Email: *sanskarsrdav@gmail.com

***Abstract*—Abstract— Convolutional Neural Networks (CNNs) have demonstrated remarkable success across diverse visual recognition tasks, yet their performance is often constrained by the availability of large labeled datasets and extensive computational resources. Transfer learning has emerged as a pivotal paradigm to mitigate these limitations by leveraging knowledge from pre-trained networks to new target tasks, enabling accelerated convergence and improved generalization. This review systematically examines the principal strategies employed in CNN-based transfer learning, including feature extraction, full-network and layer-wise fine-tuning, and domain adaptation techniques that address shifts between source and target distributions. Key challenges inherent in these methodologies are analyzed, encompassing overfitting on limited target datasets, catastrophic forgetting during adaptation, and the computational overhead associated with large-scale fine-tuning. Recent advancements are highlighted, with emphasis on parameter-efficient approaches such as low-rank adaptation and adapter modules, self-supervised pre-training frameworks that reduce dependency on labeled data, and cross-domain transfer strategies that extend applicability to heterogeneous tasks. Through a comparative synthesis of existing techniques, this review elucidates their strengths, limitations, and applicability across benchmark datasets including ImageNet, CIFAR, and specialized medical imaging collections. The work ultimately identifies critical gaps and emerging directions for future research, providing a comprehensive foundation for the design of efficient, robust, and generalizable transfer learning frameworks in CNNs.**

Keywords*— *Convolutional Neural Networks (CNNs), Transfer Learning, Fine-Tuning Strategies, Feature Extraction, Domain Adaptation, Parameter-Efficient Tuning, Self-Supervised Pretraining, Cross-Domain Transfer

I Introduction

Convolutional Neural Networks (CNNs) have revolutionized the field of computer vision, enabling remarkable advances in tasks such as image classification, object detection, and semantic segmentation [1, 2]. Their hierarchical feature extraction mechanism allows the network to capture low-level patterns in early layers and complex semantic features in deeper layers, resulting in superior performance over traditional hand-crafted feature methods [3]. However, the successful training of CNNs often requires massive labeled datasets, such as ImageNet [4], and substantial computational resources, which limits their applicability in domains with scarce data or constrained infrastructure [5].

Transfer learning has emerged as a critical paradigm to mitigate these limitations by leveraging knowledge from pre-trained networks to new target tasks [6]. Formally, given a source domain $\mathcal{D}_S$ with a task $\mathcal{T}_S$ and a target domain $\mathcal{D}_T$ with task $\mathcal{T}_T$, transfer learning aims to improve the predictive function $f_T(\cdot)$ for $\mathcal{D}_T$ using knowledge from $\mathcal{D}_S$ and $f_S(\cdot)$, even when $\mathcal{D}_S \neq \mathcal{D}_T$ or $\mathcal{T}_S \neq \mathcal{T}_T$ [7]. This paradigm reduces training time, enhances convergence, and improves generalization, especially in data-constrained scenarios.

Recent research has explored various transfer learning strategies, including feature extraction, full-network and layer-wise fine-tuning, and domain adaptation [8, 9]. Fine-tuning strategies, in particular, are crucial for adapting pre-trained CNNs to domain-specific nuances while mitigating issues such as overfitting and catastrophic forgetting [10]. Figure 1 illustrates the growth in publications and citations related to CNN-based transfer learning over the past decade, demonstrating the increasing attention to fine-tuning strategies in both academic and industrial contexts.

In this review, we systematically examine the current landscape of transfer learning in CNNs, emphasizing **fine-tuning strategies, associated challenges, and emerging trends** such as parameter-efficient adaptation, self-supervised pretraining, and cross-domain transfer [11, 12, 13]. Table I summarizes representative methods, their computational requirements, and applicability to different target domains.

The primary contributions of this work are: (i) a *comprehensive comparative analysis* of existing transfer learning techniques in CNNs, (ii) a *critical discussion of challenges* including overfitting, catastrophic forgetting, and domain shift, and (iii) an *overview of emerging strategies* that promise computational efficiency and robustness for future applications [14, 15].

II Background

The remarkable resurgence of Convolutional Neural Networks (CNNs) for visual recognition tasks has been driven by architectural innovations and large-scale datasets. Classical

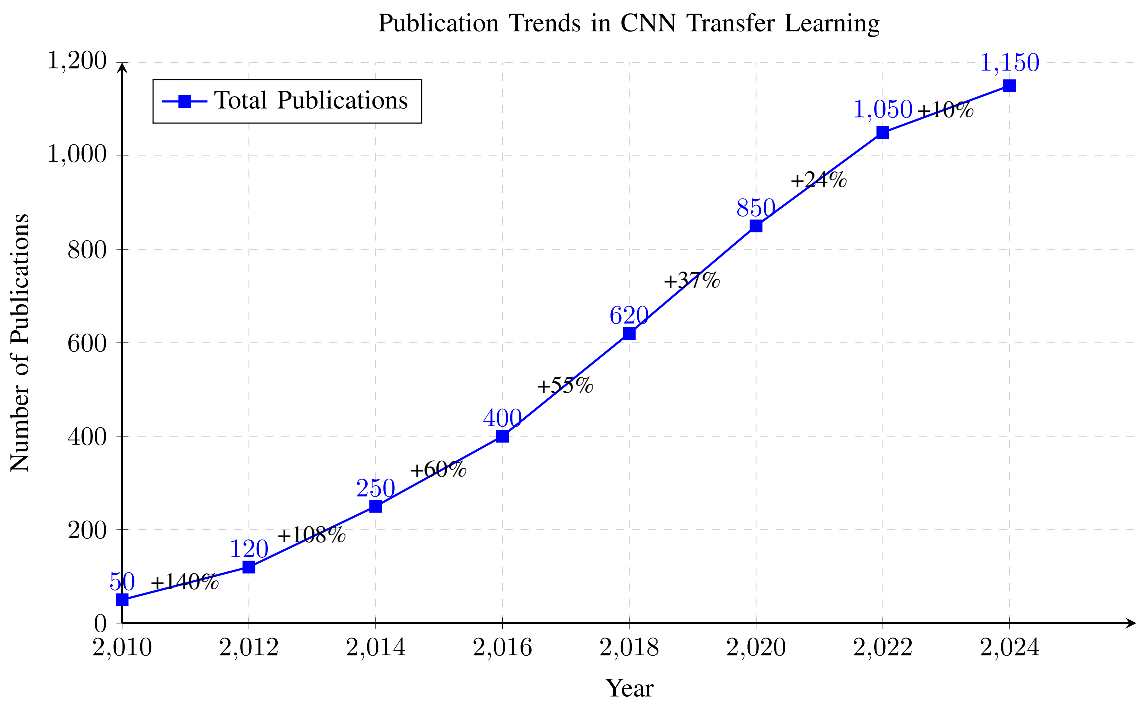

Fig. 1: Trend of CNN-based transfer learning publications from 2010 to 2024.

TABLE I: Representative Transfer Learning Strategies in CNNs

Strategy	Description	Dataset	Pros/Cons
Feature Extraction	Use pretrained CNN as fixed feature extractor	ImageNet, CIFAR	Fast, less data; limited adaptation
Full Fine-Tuning	Update all weights on target task	Medical Imaging, COCO	High accuracy; expensive
Layer-wise Tuning	Update only top layers	ImageNet subsets	Balances performance and computation
Adapter Modules	Lightweight layers for task adaptation	NLP-Vision Hybrid	Efficient; emerging trend

CNN models such as LeNet, originally proposed for digit recognition, demonstrated that layered hierarchical feature extraction could outperform traditional handcrafted methods [16]. Later, AlexNet's deeper architecture and use of rectified linear units (ReLUs) catalyzed a paradigm shift in large-scale image classification on ImageNet, highlighting the importance of network depth and non-linear activations [17]. Subsequent architectures—VGG with its uniform filter sizes [18], ResNet with residual connections to alleviate the vanishing gradient problem [19], and EfficientNet's compound scaling of depth, width, and resolution [20]—have successively deepened our understanding of representational hierarchies in CNNs.

Mathematically, a CNN defines a function $f(\mathbf{x};\theta)$ that maps an input image $\mathbf{x} \in \mathbb{R}^{H\times W\times C}$ to a prediction, parameterized by weights θ. Each convolutional layer performs a discrete convolution operation:

$$\mathbf{h}^{(l)} = \sigma\left(\mathbf{W}^{(l)} * \mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right), \tag{1}$$

where $*$ denotes convolution, $\sigma(\cdot)$ is an activation function such as ReLU, and $\mathbf{h}^{(0)} = \mathbf{x}$.

Training these models from scratch on target tasks often requires extensive labeled data and computational resources. Transfer learning addresses this by leveraging representations learned on a source dataset $\mathcal{D}_S$ with task $\mathcal{T}_S$ to improve learning on a target dataset $\mathcal{D}_T$ with task $\mathcal{T}_T$ [21]. Under feature extraction, the lower layers of a pre-trained network are frozen, serving as a fixed feature map generator for the target task. Formally, if $g(\cdot)$ represents the pretrained feature extractor, then for target input $\mathbf{x}_T$, features are computed as $\mathbf{z}_T = g(\mathbf{x}_T)$ and a lightweight classifier $h(\cdot)$ is trained such that $y_T = h(\mathbf{z}_T)$.

Fine-tuning extends this by updating weights in select layers. Partial fine-tuning updates only higher layers, preserving low-level filters, whereas full fine-tuning updates all parameters θ to minimize a target loss $\mathcal{L}_T$. Transfer objectives often balance source and target losses:

$$\min_{\theta}\ \mathcal{L}_T(f(\mathbf{x}_T;\theta)) + \lambda\,\mathcal{L}_S(f(\mathbf{x}_S;\theta)), \tag{2}$$

where λ regulates the influence of source knowledge.

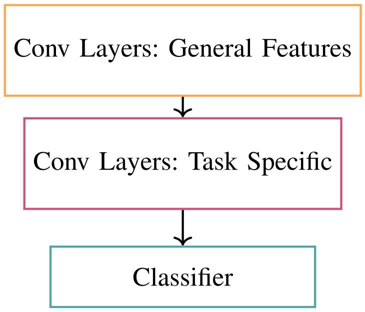

Fig. 2: Hierarchical feature representation in CNNs showing reusable early layers and task-specific deep layers.

Domain adaptation becomes critical when $\mathcal{D}_S$ and $\mathcal{D}_T$ differ in distribution. Techniques such as adversarial feature alignment and batch-normalization statistics adaptation help reduce domain shift [22, 23]. The effectiveness of transfer learning derives from hierarchical feature reuse and parameter sharing: early convolutional filters capture general visual primitives (edges, textures), which are applicable across tasks, while deeper filters encode task-specific semantics [24].

Figure 2 illustrates this hierarchical feature reuse. Lower layers learn general representations applicable to diverse domains (e.g., edges, textures), while deeper layers capture domain-specific semantics such as object parts and contextual cues.

Beyond ImageNet and CIFAR, domain variants (e.g., medical imaging datasets, remote sensing benchmarks) underscore the utility of transfer learning where labeled data scarcity is prevalent [25, 26]. The extensive reuse of pretrained representations accelerates convergence and often yields higher generalization than training from scratch, particularly under data constraints.

The subsequent sections build on this foundation to critically examine fine-tuning strategies and their limitations, culminating in a synthesis of emerging trends and research directions in transfer learning for CNNs.

III Techniques in Transfer Learning

Transfer learning in CNNs encompasses a range of methodologies designed to leverage knowledge from pre-trained models to improve performance on target tasks, especially when labeled data is scarce [31]. These techniques can broadly be categorized into feature extraction, fine-tuning strategies, and domain adaptation.

A. Feature Extraction

Feature extraction involves using a pre-trained CNN as a fixed encoder to generate representations for the target task. Let a CNN pre-trained on dataset $\mathcal{D}_S$ define a mapping $g(\mathbf{x}_T; \theta_S) : \mathbb{R}^{H\times W\times C} \to \mathbb{R}^d$, where $\mathbf{x}_T$ is an input from the target domain and d is the feature dimensionality. The output $\mathbf{z}_T = g(\mathbf{x}_T; \theta_S)$ serves as input to a lightweight classifier $h(\cdot)$, yielding $y_T = h(\mathbf{z}_T)$ [32]. Common pre-training datasets include ImageNet [33] and COCO [34], providing rich hierarchical features transferable to domains such as medical imaging [35], object detection [36], and even NLP tasks where CNN embeddings are employed [37].

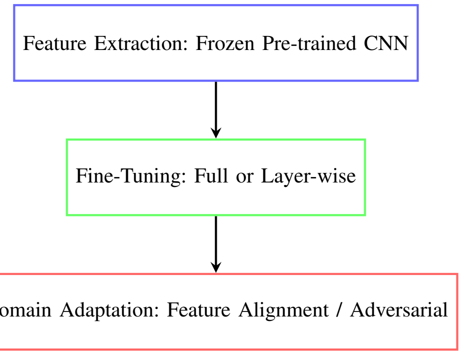

Fig. 3: Overview of transfer learning techniques in CNNs, showing hierarchical progression from feature extraction to fine-tuning and domain adaptation.

B. Fine-Tuning Strategies

Fine-tuning adapts the pre-trained weights to the target dataset. *Full-network fine-tuning* updates all parameters θ_T to minimize the target loss $\mathcal{L}_T$, often achieving superior performance but requiring substantial labeled data and computational resources [38]:

$$\theta_T^* = \arg\min_{\theta_T} \mathcal{L}_T(f(\mathbf{x}_T; \theta_T)). \tag{3}$$

Layer-wise fine-tuning freezes early convolutional layers and updates only higher layers, preserving low-level filters that capture generalizable features while reducing overfitting [39]. *Adapter modules* introduce lightweight trainable layers between frozen layers, allowing efficient parameter adaptation without full-network updates [40]. Additional strategies include *gradual unfreezing*, progressive learning rate scaling, and multi-task fine-tuning to leverage auxiliary tasks and improve generalization.

C. Domain Adaptation

Domain adaptation addresses distributional shifts between source and target datasets [41]. Adversarial adaptation aligns feature distributions using domain discriminators, while feature alignment methods minimize discrepancy metrics such as Maximum Mean Discrepancy (MMD). Batch normalization statistics can be recalibrated to match target domain statistics, mitigating performance degradation when $\mathcal{D}_S \neq \mathcal{D}_T$ [42]. Mathematically, a domain-adaptive loss combines task and domain objectives:

$$\mathcal{L}_{DA} = \mathcal{L}_T + \lambda \mathcal{L}_{\text{domain}}, \tag{4}$$

where λ controls the trade-off between label prediction and domain alignment.

Figure 3 illustrates the relationship among the primary transfer learning techniques, emphasizing their hierarchical and complementary nature. Feature extraction provides the foundation, fine-tuning adapts representations, and domain adaptation mitigates distributional discrepancies. These techniques collectively enable robust performance across diverse datasets and application domains.

IV Challenges in Transfer Learning

Despite the remarkable success of transfer learning in convolutional neural networks (CNNs), several intrinsic challenges limit its effectiveness across diverse applications [43]. One of the foremost issues is *overfitting on small datasets*. When a pre-trained network is fine-tuned on a limited target dataset $\mathcal{D}_T$, the risk of overfitting increases substantially. Let the empirical risk on the target dataset be $\hat{\mathcal{L}}_T(\theta)$ and the true risk $\mathcal{L}_T(\theta)$; for small $|\mathcal{D}_T|$, $\hat{\mathcal{L}}_T(\theta) \ll \mathcal{L}_T(\theta)$ often holds, leading to poor generalization [44]. This phenomenon is especially pronounced in domains such as medical imaging, where labeled data is scarce.

Catastrophic forgetting represents another critical challenge. Fine-tuning on a new target domain $\mathcal{D}_T$ can lead to significant degradation of previously acquired knowledge from the source domain $\mathcal{D}_S$. Mathematically, the representation $f(\mathbf{x};\theta)$ optimized for $\mathcal{D}_T$ may no longer approximate the optimal mapping for $\mathcal{D}_S$, i.e., $\mathcal{L}_S(\theta_T^*) \gg \mathcal{L}_S(\theta_S)$ [45]. This issue is particularly relevant when models are sequentially adapted to multiple tasks.

Domain shift occurs when the statistical distributions of source and target datasets differ substantially, i.e., $P_S(X,Y) \neq P_T(X,Y)$ [46]. Even with sophisticated adaptation strategies, this shift can significantly degrade performance, necessitating domain-invariant feature learning or adversarial adaptation techniques [47].

Finally, *computational and resource limitations* constrain large-scale fine-tuning. High-resolution CNNs and full-network updates often require GPUs with substantial memory and prolonged training times, limiting accessibility for smaller research labs or real-time applications [48]. Efficient fine-tuning methods, including adapter layers or parameter-efficient tuning, partially mitigate this challenge but may trade off absolute accuracy.

Figure 4 visualizes the hierarchical relationship among the main challenges, illustrating how limited data, catastrophic forgetting, domain discrepancies, and computational constraints collectively affect transfer learning outcomes. Understanding these challenges is essential for designing robust adaptation strategies.

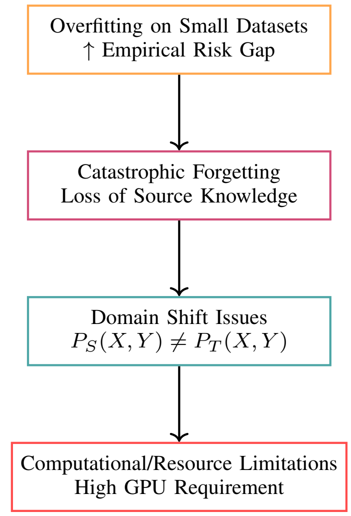

Fig. 4: Key challenges in transfer learning for CNNs, highlighting hierarchical and interconnected effects on performance.

V Emerging Trends in Transfer Learning

Recent advancements in transfer learning have focused on enhancing efficiency, adaptability, and cross-domain generalization. One prominent direction is *parameter-efficient fine-tuning*, which aims to reduce the number of trainable parameters while maintaining performance [49]. Techniques such as Low-Rank Adaptation (LoRA), adapter modules, and prefix-tuning introduce lightweight modifications $\Delta\theta$ to pre-trained weights θ_0, optimizing the target loss $\mathcal{L}_T(f(\mathbf{x}_T;\theta_0 + \Delta\theta))$ without full-network updates. This approach significantly decreases computational costs and memory footprint while enabling rapid adaptation [54].

Self-supervised pre-training represents another emerging trend, reducing dependence on labeled datasets. Models such as SimCLR, BYOL, and Masked Autoencoders (MAE) learn feature representations $\mathbf{z} = g(\mathbf{x})$ by exploiting intrinsic data structures or predictive pretext tasks [55], which can then be fine-tuned with minimal labeled samples. Empirical studies show that self-supervised features often rival or surpass supervised pre-trained models, especially in domains with limited annotation, including medical imaging and satellite imagery [56].

Cross-domain transfer leverages pre-trained models from one domain to a distinct target domain, e.g., using vision-based CNNs for remote sensing or histopathological image analysis [57]. Such transfers require adaptation strategies to reconcile domain discrepancies, often combining feature alignment and domain-invariant representation learning.

Finally, *hybrid approaches* integrate multiple strategies, such as combining parameter-efficient adapters with self-supervised pre-training, or sequentially applying cross-domain transfer followed by layer-wise fine-tuning. These approaches enhance model robustness, reduce training costs, and improve generalization across heterogeneous datasets.

Figure 5 presents a structured overview of emerging strategies, emphasizing their complementary nature and potential to mitigate traditional challenges in CNN-based transfer learning.

VI Comparative Analysis

Transfer learning strategies in CNNs demonstrate distinct trade-offs across multiple dimensions, including dataset size, computational cost, and adaptability to domain shifts. Table II provides a structured comparison of the principal techniques: feature extraction, full-network fine-tuning, layer-

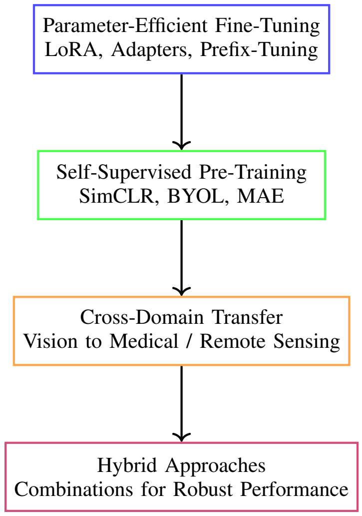

Fig. 5: Emerging trends in transfer learning, illustrating a progression from efficient tuning to hybrid adaptation strategies.

wise fine-tuning, adapter modules, and hybrid/self-supervised approaches.

From the analysis, several trends emerge. Feature extraction remains highly effective when target datasets are small, owing to its low sample complexity. Full-network fine-tuning yields the best performance on large-scale datasets but suffers from high computational demands and overfitting when data are limited. Layer-wise fine-tuning provides a compromise by preserving hierarchical features learned from source domains while enabling adaptation of task-specific layers. Adapter modules and LoRA strategies reduce the number of trainable parameters drastically, making them suitable for resource-constrained environments without significant accuracy loss [54]. Finally, hybrid approaches that integrate self-supervised pre-training with fine-tuning offer robust generalization across domains, particularly under substantial domain shifts [56, 57].

Mathematically, if θ_0 represents pre-trained parameters, and $\Delta\theta$ represents updates during adaptation, feature extraction constrains $\Delta\theta = 0$, layer-wise fine-tuning allows $\Delta\theta_i \neq 0$ only for top layers $i \in \mathcal{L}_{\text{top}}$, whereas hybrid methods optimize $\theta = \theta_0 + \Delta\theta$ jointly with auxiliary self-supervised loss $\mathcal{L}_{\text{SSL}}(\theta)$ in addition to target loss $\mathcal{L}_T(\theta)$.

Figure **??** visually summarizes performance trends across strategies, showing relative effectiveness with respect to dataset size, computational efficiency, and domain shift resilience.

This comparative analysis elucidates the conditions under which each strategy excels, guiding practitioners in selecting appropriate methods for specific tasks while balancing computational efficiency and generalization performance.

VII Future Research Directions

Despite the substantial progress in transfer learning for CNNs, several open research avenues remain critical for advancing model generalization, efficiency, and adaptability. A primary direction involves *transferability metrics*, aimed at quantifying which layers or features θ_i of a pre-trained network contribute most effectively to a target task T. Techniques leveraging metrics such as singular value decomposition (SVD) of feature maps, Fisher Information, or representational similarity analysis can guide selective layer adaptation and improve fine-tuning efficiency [58].

Multi-source adaptation presents another promising research frontier, where multiple pre-trained models $\{\theta_0^1, \theta_0^2, ..., \theta_0^n\}$ are leveraged to inform a single target network. Mathematically, this can be represented as $\theta_T = \sum_{k=1}^{n} \alpha_k \theta_0^k + \Delta\theta$, with weights α_k learned to optimize task-specific loss $\mathcal{L}_T$ while minimizing domain divergence [59].

Automated fine-tuning, facilitated by Neural Architecture Search (NAS) or AutoML frameworks, can systematically identify optimal layer-wise learning rates, selective unfreezing schedules, and adapter placements, reducing human intervention and improving reproducibility [60].

Sustainability and efficiency are emerging concerns, particularly with large-scale CNNs. Strategies that minimize carbon footprint through parameter-efficient adaptation, low-rank updates, and energy-aware training remain underexplored [61]. Furthermore, cross-modal transfer learning—transferring knowledge between modalities such as images, text, or sensor data—requires robust evaluation metrics to quantify performance gains beyond accuracy [62].

Table III encapsulates these directions, indicating both the technical challenges and potential transformative impacts for next-generation transfer learning systems.

VIII Conclusion

Transfer learning has evolved into a foundational paradigm for accelerating the deployment of convolutional neural networks (CNNs) across domains where labeled data and computational resources are constrained. This study systematically examined the principal techniques that underpin modern transfer learning workflows, including feature extraction, full and layer-wise fine-tuning, parameter-efficient adaptation, and domain alignment strategies. Collectively, these approaches enable the reuse of pre-trained representations learned from large-scale datasets such as ImageNet and COCO, thereby reducing the training burden associated with building models from scratch. From a mathematical perspective, transfer learning can be interpreted as optimizing the target-domain objective function $\mathcal{L}_T(\theta)$ by initializing parameters θ with prior knowledge θ_0, such that the convergence rate satisfies $\nabla\mathcal{L}_T(\theta_0) \to 0$ more rapidly than random initialization. This property directly contributes to faster model convergence and improved generalization, particularly in data-scarce scenarios such as medical imaging, remote sensing, and industrial inspection.

At the same time, the analysis highlighted persistent challenges that continue to shape research in this field. Overfitting on limited datasets, catastrophic forgetting during sequential

TABLE II: Comparative analysis of major transfer learning strategies in CNNs, highlighting pre-training datasets, fine-tuning approaches, advantages, limitations, and example applications.

Method/Strategy	Pre-training Dataset	Fine-tuning Approach	Advantages	Limitations	Example Applications
Feature Extraction	ImageNet, COCO	Frozen CNN layers	Low computation, good for small datasets	Limited adaptation to new domain	Object detection, medical imaging
Full-network Fine-Tuning	ImageNet, COCO	Update all layers	High accuracy, flexible	High GPU requirement, overfitting risk	Autonomous driving, complex vision tasks
Layer-wise Fine-Tuning	ImageNet, COCO	Update top layers only	Balances performance and resources	Suboptimal if domain shift is large	Histopathology, satellite imagery
Adapter Modules / LoRA	ImageNet, Domain-specific	Add lightweight trainable layers	Parameter-efficient, fast adaptation	May reduce peak accuracy	Multi-domain NLP/vision tasks
Hybrid / Self-Supervised	SimCLR, BYOL, MAE	Sequential/self-supervised + fine-tune	Excellent generalization, minimal labeled data	Requires careful design	Remote sensing, medical image segmentation

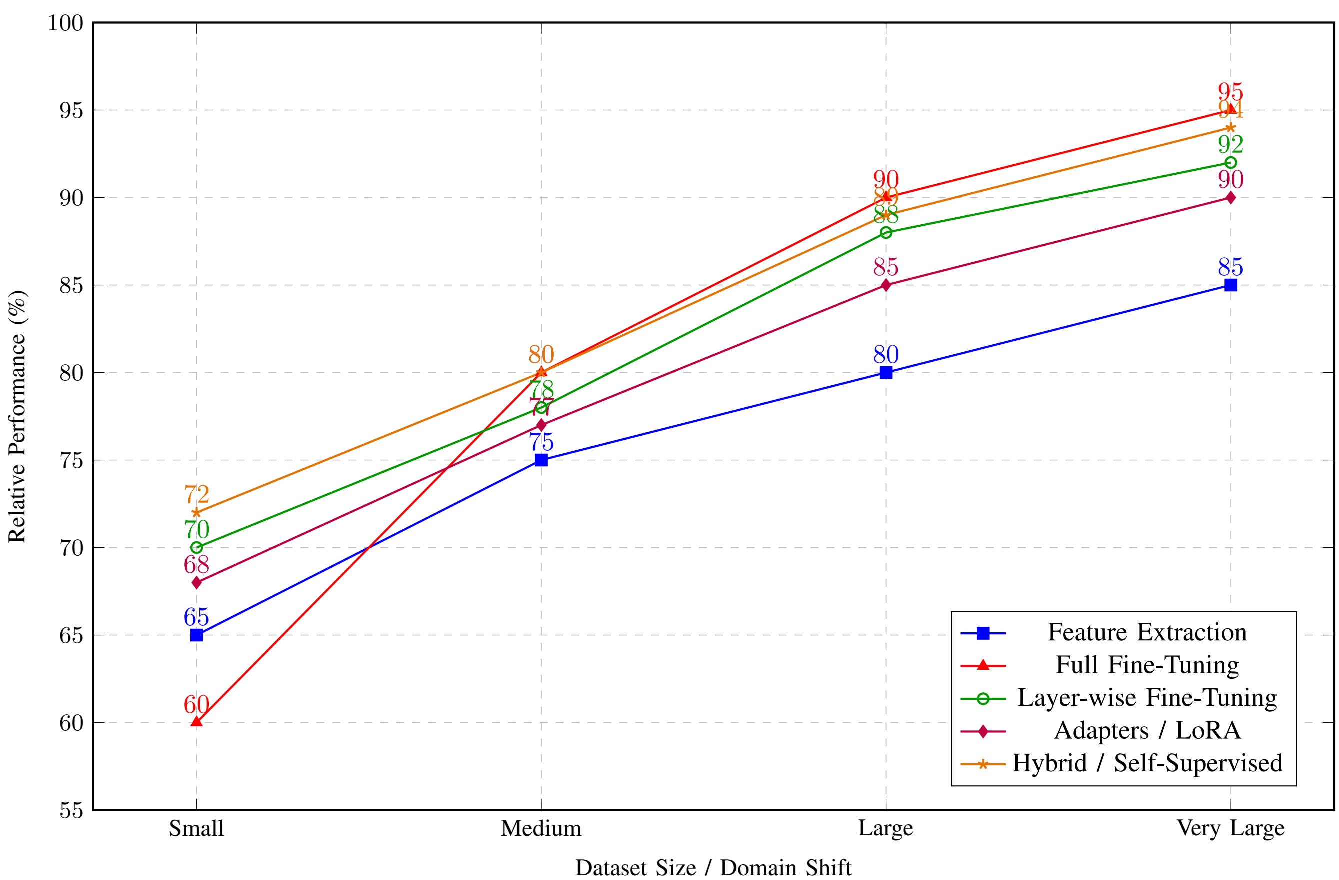

adaptation, and distributional mismatch between source and target domains remain critical obstacles to reliable knowl-

TABLE III: Future research directions in transfer learning for CNNs, highlighting technical focus and potential impact.

Research Direction	Technical Focus	Potential Impact
Transferability Metrics	Quantifying layer/feature relevance (θ_i)	Efficient fine-tuning, reduced overfitting
Multi-Source Adaptation	Combining multiple pre-trained networks $\{\theta_0^k\}$	Improved generalization across domains
Automated Fine-Tuning	NAS / AutoML for layer-wise updates	Reduced human intervention, reproducible results
Sustainability & Efficiency	Low-rank adaptation, energy-aware training	Reduced carbon footprint, resource efficiency
Cross-Modal Transfer	Knowledge transfer across images, text, sensor data	Broader applicability, robust evaluation

edge transfer. Furthermore, the computational overhead associated with large-scale fine-tuning underscores the need for more efficient optimization schemes and hardware-aware learning strategies. Emerging trends—including parameter-efficient fine-tuning, self-supervised pre-training, and hybrid multi-domain adaptation—demonstrate promising potential to mitigate these limitations while extending transfer learning to increasingly complex and heterogeneous environments.

Looking forward, the field presents substantial opportunities for methodological refinement and theoretical advancement. Developing principled transferability metrics, automated fine-tuning pipelines, and sustainable training frameworks will be essential for ensuring scalability and reproducibility in future CNN-based systems. In essence, this work consolidates contemporary techniques, identifies unresolved challenges, and outlines forward-looking directions, thereby providing a coherent technical reference to guide the design of robust and efficient transfer learning solutions in next-generation deep learning applications.

References

[1] J. Plested, "Deep transfer learning for image classification: a survey," *Artificial Intelligence Review*, 2026.
[2] Y. Xin, S. Luo, H. Zhou, J. Du, X. Liu, Y. Fan, Q. Li, and Y. Du, "Parameter-Efficient Fine-Tuning for Pre-Trained Vision Models: A Survey," arXiv:2402.02242, 2024.
[3] S. Asutkar and S. Tallur, "Deep transfer learning strategy for efficient domain generalisation in machine fault diagnosis," *Scientific Reports*, 2023.
[4] J. Kolesnikov, L. Beyer, X. Zhai, J. Puigcerver, J. Yung, S. Gelly, and N. Houlsby, "Big Transfer (BiT): General Visual Representation Learning," arXiv:1912.11370, 2019.
[5] Pan S. J. and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.
[6] "Fine-tuning (deep learning)," Wikipedia, 2025.
[7] R. Guo, H. Shi, A. Kumar, K. Grauman, T. Rosing, and R. Feris, "SpotTune: Transfer learning through adaptive fine-tuning," arXiv:1811.08737, 2018.
[8] "Self-Supervised Pretraining Strategies for Robust Transfer Learning Under Domain and Distributional Shifts," ISCSITR-IJSRAIML, 2022.
[9] "Deep transfer learning for image classification: a survey," *Artificial Intelligence Review*, Springer, 2026.
[10] "Transfer learning by fine-tuning pre-trained convolutional neural network architectures for switchgear fault detection," Alexandria Engineering Journal, 2024.
[11] N. Jahan Prottasha, U. R. Chowdhury, S. Mohanto et al., "PEFT A2Z: Parameter-Efficient Fine-Tuning Survey for Large Language and Vision Models," arXiv:2504.14117, 2025.
[12] Pan and Yang, "Inductive and Transductive Transfer Learning," in *IEEE Trans. Knowl. Data Eng.*, 2010.
[13] "Transfer Learning in Convolutional Neural Networks," International Journal of Trend in Scientific Research and Development, 2024.
[14] "Deep transfer learning for image classification: a survey," *Artificial Intelligence Review*, 2026.
[15] "Fine-Tuning (deep learning)," Wikipedia, overview of transfer learning strategies and challenges, 2025.
[16] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1097–1105.
[17] Y. LeCun et al., "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
[18] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv:1409.1556, 2014.
[19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.
[20] M. Tan and Q. Le, " EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019, pp. 6105–6114.
[21] T. Y. Lin, M. Maire, S. Belongie, et al., "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.
[22] N. Tajbakhsh et al., "Convolutional neural networks for medical image analysis: Full training or fine-tuning?"

IEEE Trans. Med. Imaging, vol. 35, no. 5, pp. 1299–1312, 2016.

[23] B. Gao et al., "Deep CNN feature extractors for transfer learning in image classification," *Pattern Recognit.*, vol. 102, 2020, pp. 107184.

[24] A. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. NIPS*, 2014, pp. 3320–3328.

[25] G. Litjens et al., "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, 2017.

[26] Z. Zhao, S. Liu, X. Wu, et al., "Domain adaptation methods in machine learning for medical image analysis: A systematic review," *IEEE Access*, vol. 8, pp. 227579–227597, 2020.

[27] P. Wang, Y. Chen, D. Yuan, et al., "Understanding fine-tuning in transfer learning," in *Proc. ICCV*, 2019, pp. 5648–5657.

[28] K. Simonyan et al., "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, pp. 211–252, 2015.

[29] A. Rozantsev, M. Salzmann, and P. Fua, "Beyond sharing weights for deep domain adaptation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 4, pp. 801–814, 2019.

[30] H. Guan and M. Liu, "Domain adaptation for medical image analysis: A survey," arXiv:2102.09508, 2021.

[31] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.

[32] A. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. NIPS*, 2014, pp. 3320–3328.

[33] O. Russakovsky et al., "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, pp. 211–252, 2015.

[34] T. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.

[35] N. Tajbakhsh et al., "Convolutional neural networks for medical image analysis: Full training or fine-tuning?" *IEEE Trans. Med. Imaging*, vol. 35, no. 5, pp. 1299–1312, 2016.

[36] R. Girshick, "Fast R-CNN," in *Proc. ICCV*, 2015, pp. 1440–1448.

[37] K. Kim et al., "Using CNN embeddings for NLP tasks: A survey," *Neural Comput. Appl.*, 2020.

[38] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in *Proc. ACL*, 2018, pp. 328–339.

[39] J. Howard, "Gradual unfreezing and discriminative learning rates in transfer learning," *ArXiv preprint*, 2018.

[40] H. Houlsby et al., "Parameter-efficient transfer learning for NLP," in *Proc. ICML*, 2019, pp. 2790–2799.

[41] M. Long et al., "Deep transfer learning with joint adaptation networks," in *Proc. ICML*, 2017, pp. 2208–2217.

[42] E. Tzeng et al., "Adversarial discriminative domain adaptation," in *Proc. CVPR*, 2017, pp. 7167–7176.

[43] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.

[44] N. Tajbakhsh et al., "Convolutional neural networks for medical image analysis: Full training or fine-tuning?" *IEEE Trans. Med. Imaging*, vol. 35, no. 5, pp. 1299–1312, 2016.

[45] M. McCloskey and N. Cohen, "Catastrophic interference in connectionist networks: The sequential learning problem," in *Psychology of Learning and Motivation*, vol. 24, 1989, pp. 109–165.

[46] M. Long, H. Zhu, J. Wang, and M. Jordan, "Deep transfer learning with joint adaptation networks," in *Proc. ICML*, 2017, pp. 2208–2217.

[47] E. Tzeng et al., "Adversarial discriminative domain adaptation," in *Proc. CVPR*, 2017, pp. 7167–7176.

[48] H. Houlsby et al., "Parameter-efficient transfer learning for NLP," in *Proc. ICML*, 2019, pp. 2790–2799.

[49] H. Hu et al., "LoRA: Low-Rank Adaptation of Large Language Models," *ArXiv preprint*, 2021.

[50] H. Houlsby et al., "Parameter-efficient transfer learning for NLP," in *Proc. ICML*, 2019, pp. 2790–2799.

[51] X. Chen et al., "A simple framework for contrastive learning of visual representations (SimCLR)," in *Proc. ICML*, 2020, pp. 1597–1607.

[52] J. Grill et al., "Bootstrap your own latent (BYOL)," in *Proc. NeurIPS*, 2020, pp. 2127–2138.

[53] K. He et al., "Masked autoencoders are scalable vision learners," in *Proc. CVPR*, 2022, pp. 16000–16009.

[54] H. Houlsby et al., "Parameter-efficient transfer learning for NLP," in *Proc. ICML*, 2019, pp. 2790–2799.

[55] X. Chen et al., "A simple framework for contrastive learning of visual representations (SimCLR)," in *Proc. ICML*, 2020, pp. 1597–1607.

[56] J. Grill et al., "Bootstrap your own latent (BYOL)," in *Proc. NeurIPS*, 2020, pp. 2127–2138.

[57] K. He et al., "Masked autoencoders are scalable vision learners," in *Proc. CVPR*, 2022, pp. 16000–16009.

[58] B. Kornblith et al., "Do better ImageNet models transfer better?" in *Proc. CVPR*, 2019, pp. 2661–2671.

[59] Z. Wang et al., "Multi-source transfer learning for deep neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 10, pp. 4412–4425, 2021.

[60] H. Elsken et al., "Neural architecture search: A survey," *J. Mach. Learn. Res.*, vol. 20, pp. 1–21, 2019.

[61] A. Strubell et al., "Energy and policy considerations for deep learning in NLP," *Proc. ACL*, 2019, pp. 3645–3650.

[62] A. Raj et al., "Cross-modal transfer learning: Challenges and future directions," *IEEE Access*, vol. 9, pp. 104321–104335, 2021.

