

# Transfer Learning in Natural Language Processing: Practical Techniques with Pretrained Language Models

**Debabrata Pruseth**

AI Architect & Applied AI Researcher  
Singapore

---

## Author Note

This article is a research-style companion version of author's blog post "[Transfer Learning for NLP](#)"

This research presents an applied methodology for using transfer learning in natural language processing systems built on pretrained language models. The study is positioned as an implementation-oriented framework for selecting, adapting, and governing pretrained models across practical NLP tasks such as sentiment analysis, spam detection, named entity recognition, question answering, domain-specific classification, and text generation. The work follows the author's applied AI research style of converting technical workflows into structured, reproducible decision frameworks.

The central motivation is that modern NLP development has shifted from hand-engineered linguistic rules toward reusable pretrained representations. Instead of training language models from scratch for every downstream task, practitioners can adapt shared linguistic knowledge acquired during large-scale pretraining. The framework developed in this study organizes transfer learning strategies according to model architecture, dataset size, compute budget, domain specificity, and task complexity. The paper does not report new benchmark experiments. It contributes a practical synthesis and decision methodology for applied NLP model selection and adaptation.

---

## Abstract

Transfer learning has become a foundational methodology in natural language processing because it enables practitioners to adapt pretrained language models to downstream tasks with substantially less task-specific data and engineering effort than traditional rule-based or from-scratch approaches. This paper presents a practical framework for applying transfer learning in NLP using pretrained language models such as BERT, RoBERTa, DistilBERT, BioBERT, SciBERT, GPT-style decoders, T5-style encoder-decoder models, and contemporary foundation models. The study organizes transfer learning into a decision workflow covering feature extraction, partial fine-tuning, full fine-tuning, adapter-based adaptation, prompt tuning, instruction tuning, and zero-shot or few-shot inference. The methodology emphasizes the structural roles of embedding layers, transformer blocks, and

task-specific heads, and translates these components into practical design decisions. It further provides guidance on head-only training, gradual unfreezing, discriminative learning rates, validation-based early stopping, domain-adaptive pretraining, and parameter-efficient fine-tuning methods such as LoRA, BitFit, prefix tuning, and adapter modules. The contribution is a practitioner-oriented framework that connects NLP task requirements with suitable transfer learning strategies while highlighting trade-offs in accuracy, compute cost, data availability, deployment constraints, and governance. The study is conceptual and implementation-oriented rather than experimental; therefore, its recommendations require empirical validation for each dataset, domain, and production environment.

## Keywords

Transfer Learning, Natural Language Processing, Pretrained Language Models, BERT, RoBERTa, DistilBERT, GPT, T5, Fine-Tuning, Feature Extraction, Adapter Tuning, LoRA, Prompt Tuning, Instruction Tuning, Parameter-Efficient Fine-Tuning

---

## 1. Introduction

Natural language processing has historically relied on rules, dictionaries, pattern matching, statistical features, and manually engineered linguistic representations. Such systems could support constrained applications, but they often struggled with ambiguity, context, domain variation, idioms, and evolving language. Modern NLP systems increasingly rely on pretrained language models that learn general-purpose representations from large text corpora and are then adapted to specific tasks. This shift is central to the practical value of transfer learning in NLP.

Transfer learning in NLP refers to the reuse of representations learned from one training setting, usually large-scale self-supervised pretraining, and their adaptation to downstream tasks such as sentiment classification, spam detection, named entity recognition, question answering, summarization, chatbot intent classification, or domain-specific document analysis. The practical logic is simple: language has shared structure across tasks. A sentence such as “Can I get a refund?” may indicate dissatisfaction in a sentiment system, a request in a chatbot, or a refund category in a support-ticket classifier. The same linguistic signal can be repurposed through task-specific adaptation.

The transformer architecture introduced a scalable foundation for sequence modeling based on attention mechanisms rather than recurrence or convolution, enabling parallelizable representation learning for language tasks. BERT demonstrated that deep bidirectional pretrained representations could be fine-tuned with an additional output layer across a wide range of NLP tasks. Subsequent models such as RoBERTa, DistilBERT, T5, and domain-specific variants expanded the design space for transfer learning by improving pretraining protocols, reducing model size, unifying task formats, or specializing representations for particular domains.

This study examines the following research question:

## **How can pretrained language models be selected and adapted systematically for practical NLP tasks under varying constraints of data, compute, domain specificity, and task complexity?**

The central contribution is a structured decision framework for transfer learning in NLP. The framework converts practical adaptation techniques into a method that supports model selection, fine-tuning design, parameter-efficient adaptation, validation, deployment readiness, and responsible use.

---

## **2. Research Objective**

The objective of this study is to develop a practical methodology for using pretrained language models in applied NLP systems.

The specific objectives are:

1. To define transfer learning in NLP as a reusable adaptation strategy for pretrained language representations.
  2. To explain the architectural components that influence adaptation decisions.
  3. To classify commonly used pretrained NLP model families by task suitability.
  4. To organize transfer learning methods into feature extraction, fine-tuning, adapter-based tuning, prompt tuning, instruction tuning, and zero-shot or few-shot inference.
  5. To provide a decision framework based on dataset size, compute budget, task complexity, and domain specificity.
  6. To identify best practices for stable and responsible fine-tuning.
  7. To define limitations and governance considerations for applied NLP transfer learning.
- 

## **3. Background and Conceptual Foundation**

### **3.1 From Rule-Based NLP to Pretrained Language Models**

Earlier NLP systems often required human-designed rules. For example, a spam detector might search for terms such as “free,” “winner,” or “urgent.” Such rules can be interpretable, but they are brittle because they do not reliably capture context, sarcasm, domain shift, or adversarial variation.

Pretrained language models changed this workflow by learning contextual representations from large-scale text before task-specific adaptation. BERT introduced bidirectional transformer-based pretraining for language understanding tasks. ULMFiT previously demonstrated the practical value of inductive transfer learning for text classification and introduced important fine-tuning techniques such as discriminative learning rates and gradual unfreezing. T5 later framed NLP tasks in a unified text-to-text format, further generalizing the transfer learning paradigm.

### 3.2 Why Transfer Learning Works in NLP

Transfer learning works well in NLP because linguistic structure is reusable across tasks. Syntax, semantics, entity relationships, discourse patterns, sentiment cues, and contextual dependencies learned during pretraining can support multiple downstream objectives. A pretrained encoder can represent a customer complaint, legal clause, medical note, or research abstract in ways that can later be specialized for classification, extraction, ranking, retrieval, or generation.

The practical effect is that practitioners often do not need to build a language representation system from the beginning. Instead, they can select a pretrained backbone and adapt it using a method appropriate to the task, data, and deployment environment.

### 3.3 Core Components of a Pretrained Language Model

A modern pretrained language model can be interpreted through three practical components.

Component	Function	Adaptation Decision
Embedding layer	Converts tokens into vector representations	Often frozen or updated slowly
Transformer blocks	Learn contextual relationships through attention mechanisms	Frozen, partially fine-tuned, fully fine-tuned, or adapted with parameter-efficient modules
Task-specific head	Maps learned representations to task outputs	Usually replaced for supervised downstream tasks

The embedding layer functions as the model's learned input representation system. Transformer blocks act as the contextual reasoning engine. The task-specific head converts representations into outputs such as class labels, entity tags, answer spans, or generated text. In many practical workflows, the head is trained first, the upper transformer layers are adapted next, and lower layers are frozen or updated with smaller learning rates.

---

## 4. Pretrained Model Families for NLP Transfer Learning

Different pretrained model families are suited to different categories of NLP tasks. The selection should be based on whether the task primarily requires understanding, generation, domain specialization, sequence-to-sequence transformation, or multimodal reasoning.

Model Family	Architecture Type	Practical Strength	Representative Use Cases
BERT	Encoder	Contextual language understanding	Classification, NER, extractive QA
RoBERTa	Optimized BERT-style encoder	Strong encoder performance through improved pretraining design	Classification, ranking, QA, NER
DistilBERT	Compressed encoder	Lower latency and smaller footprint	Edge NLP, mobile inference, constrained deployment
BioBERT	Domain-specific encoder	Biomedical language understanding	Biomedical NER, relation extraction, biomedical QA
SciBERT	Domain-specific encoder	Scientific text representation	Literature mining, research classification
GPT-style models	Decoder	Text generation and instruction following	Dialogue, completion, summarization, creative generation
T5	Encoder-decoder	Text-to-text task unification	Translation, summarization, QA, multi-task learning
Foundation models	Large general-purpose models	Zero-shot, few-shot, and instruction-based use	Rapid prototyping, reasoning, multi-domain assistance

RoBERTa showed that careful pretraining design and hyperparameter choices could substantially affect downstream performance. DistilBERT demonstrated that knowledge distillation can produce a smaller and faster BERT-style model while retaining much of its language understanding capability. BioBERT illustrates the value of domain-adaptive pretraining when the target domain has specialized vocabulary and distributional patterns, such as biomedical literature.

## 5. Proposed Framework and Methodology

The proposed methodology is a decision-driven transfer learning framework for NLP. It consists of six stages: task definition, model family selection, adaptation strategy selection, fine-tuning design, validation and error analysis, and deployment governance.

### 5.1 Stage 1: Define the NLP Task and Output Format

The first stage is to define the task in operational terms. The practitioner should specify the input, expected output, label structure, evaluation metrics, and deployment context.

Task Type	Input	Output	Common Metrics
Sentiment analysis	Text review or message	Sentiment class	Accuracy, F1-score, AUC

<b>Task Type</b>	<b>Input</b>	<b>Output</b>	<b>Common Metrics</b>
Spam detection	Email or message	Spam / not spam	Precision, recall, F1-score
Named entity recognition	Sentence or document	Token-level entity labels	Entity-level F1-score
Question answering	Question and context	Answer span or generated answer	Exact match, F1-score
Text summarization	Long document	Short summary	ROUGE, human evaluation
Intent classification	User query	Intent label	Accuracy, macro F1-score
Domain document classification	Specialized text	Domain label	F1-score, calibration metrics

The task definition determines whether an encoder, decoder, encoder-decoder, or large instruction-tuned model is appropriate.

## **5.2 Stage 2: Select the Model Family**

Model selection should match the problem structure.

For classification, NER, retrieval ranking, and extractive QA, encoder models such as BERT, RoBERTa, DistilBERT, BioBERT, or SciBERT are typically suitable because they are optimized for language understanding. For generation, dialogue, completion, and instruction-following tasks, decoder models are more appropriate. For tasks that can be naturally expressed as input-output text transformation, such as translation, summarization, or multi-task QA, encoder-decoder models such as T5 provide a unified interface.

Domain specificity is critical. A general-purpose model may perform well on broad internet-style text but underperform in medical, legal, scientific, financial, or enterprise-specific settings. In such cases, domain-adaptive pretraining or domain-specific pretrained models should be considered.

## **5.3 Stage 3: Select the Transfer Learning Strategy**

The core adaptation decision is choosing how much of the pretrained model should be modified.

Strategy	Description	Suitable Conditions	Main Trade-Off
Feature extraction	Freeze pretrained model and use embeddings with a separate classifier	Small data, low compute, simple tasks	Limited adaptation
Head-only training	Replace and train only the task head	Small labeled dataset, baseline creation	May underfit complex tasks
Partial fine-tuning	Unfreeze top transformer layers	Medium data, moderate complexity	Requires tuning stability
Full fine-tuning	Update all model parameters	Large labeled data, high compute	Expensive and overfitting-prone
Adapter-based tuning	Insert or train small modules while freezing backbone	Low compute, many tasks, modular deployment	Requires adapter management
Prompt tuning	Learn or design prompts while keeping base model fixed	Limited data, large model access	Sensitive to prompt design
Instruction tuning	Fine-tune on instruction-formatted tasks	General instruction-following behavior	Requires curated instruction data
Zero-shot / few-shot inference	Use task description or examples without model training	API access, exploratory tasks	Less controllable and harder to validate

LoRA is a prominent parameter-efficient fine-tuning method that freezes pretrained weights and injects trainable low-rank matrices into transformer layers, substantially reducing trainable parameters and memory requirements relative to full fine-tuning. Prefix tuning learns continuous task-specific prompt vectors while keeping the base model frozen, representing another parameter-efficient adaptation strategy.

## 5.4 Stage 4: Fine-Tuning Design

Fine-tuning should be treated as a controlled adaptation process rather than a single training command.

### 5.4.1 Establish a Head-Only Baseline

The recommended starting point is to freeze the pretrained backbone and train only the task-specific head. This provides a low-cost baseline and helps determine whether the existing representation is already sufficient for the target task.

#### **5.4.2 Apply Discriminative Learning Rates**

Different layers should learn at different rates. The task head can use the highest learning rate because it starts from random or newly initialized weights. Upper transformer layers can use a smaller learning rate because they require task adaptation. Lower layers should be frozen or updated conservatively because they often encode general linguistic features.

#### **5.4.3 Use Gradual Unfreezing**

Gradual unfreezing reduces instability. The workflow begins with the classification or task head, then progressively unfreezes upper transformer blocks, and only later considers deeper layers if validation results justify additional adaptation. ULMFiT established gradual unfreezing and discriminative fine-tuning as important techniques for transfer learning in text classification.

#### **5.4.4 Monitor Validation Metrics**

Validation should use metrics aligned with task risk. Accuracy alone may be insufficient for imbalanced classification. Macro F1-score, weighted F1-score, AUC, precision, recall, calibration error, and confusion matrices may be more informative depending on the use case.

#### **5.4.5 Prevent Catastrophic Forgetting**

Catastrophic forgetting occurs when aggressive fine-tuning overwrites useful pretrained knowledge. It can be mitigated through lower learning rates, layer freezing, gradual unfreezing, early stopping, regularization, and parameter-efficient adaptation.

#### **5.4.6 Use Domain-Adaptive Pretraining When Needed**

If the downstream corpus differs substantially from the original pretraining distribution, domain-adaptive pretraining can be useful. For example, biomedical, legal, scientific, and enterprise-support datasets often contain specialized vocabulary and discourse structures. BioBERT provides evidence that continued pretraining on biomedical corpora can improve biomedical text mining tasks.

### **5.5 Stage 5: Decision Table for Practical NLP Transfer Learning**

The following decision table translates the framework into operational guidance.

<b>Data Availability</b>	<b>Compute Budget</b>	<b>Task Complexity</b>	<b>Recommended Approach</b>
Small labeled dataset	Low	Simple classification	Feature extraction plus logistic regression, SVM, or shallow classifier
Small labeled dataset	Medium	Domain-specific classification	Head-only training followed by partial fine-tuning
Small labeled dataset	Very low	Any moderate task	LoRA, BitFit, prefix tuning, or adapter-based tuning
Medium labeled dataset	Medium	NER, QA, ranking, sequence tasks	Partial fine-tuning of upper transformer layers
Large labeled dataset	High	Complex or high-stakes task	Full fine-tuning with strong validation and monitoring
No labeled dataset	API access only	Exploratory task	Zero-shot or few-shot prompting
Large unlabeled domain corpus	Low to medium	Domain shift	Domain-adaptive pretraining followed by supervised fine-tuning
Mixed text and image input	High	Multimodal task	Multimodal model or multimodal foundation model

This table is not a substitute for empirical evaluation. It provides a starting point for selecting a strategy that can then be tested on validation data.

## 6. Technical Analysis and Practical Findings

This study does not introduce new benchmark results. The practical findings are derived from framework analysis, model architecture characteristics, and established transfer learning principles.

### 6.1 Transfer Learning Reduces the Need for Hand-Crafted Rules

A pretrained model already encodes reusable linguistic patterns. Instead of manually creating extensive rule sets for tasks such as spam detection or support-ticket classification, practitioners can adapt pretrained contextual representations to the target label space.

## 6.2 Model Architecture Should Match the Task

Encoder models are generally appropriate for understanding tasks. Decoder models are appropriate for generation. Encoder-decoder models are suitable when the task can be expressed as text-to-text transformation. Domain-specific encoders are preferable when vocabulary and discourse patterns differ significantly from general text.

## 6.3 Fine-Tuning Depth Should Match Data and Risk

Small datasets do not usually justify updating every parameter. Head-only training, partial fine-tuning, and parameter-efficient methods are often more stable starting points. Full fine-tuning should be reserved for cases with sufficient data, compute, and validation capacity.

## 6.4 Parameter-Efficient Fine-Tuning Supports Practical Deployment

Adapter-based methods, LoRA, BitFit, and prefix tuning allow practitioners to adapt large models without modifying all base weights. This supports modular deployments where one frozen backbone can support multiple task-specific adapters. LoRA specifically reduces trainable parameter requirements by injecting low-rank trainable matrices while freezing pretrained weights.

## 6.5 Validation Must Be Task-Specific

Validation metrics should reflect real operational risk. A sentiment model may require macro F1-score if classes are imbalanced. A medical entity extraction system may require high recall. A spam detector may require both false-positive and false-negative analysis. A chatbot intent classifier may require confusion analysis across similar intents.

---

# 7. Discussion

The proposed framework positions transfer learning as a practical design discipline rather than a single technique. The central decision is not merely which pretrained model to use, but how to adapt it under real constraints.

A low-resource startup building a support-ticket classifier may prefer feature extraction, DistilBERT, or LoRA-style adaptation. A biomedical research team may use BioBERT or domain-adaptive pretraining. A general-purpose assistant may rely on instruction-tuned foundation models and few-shot prompting. A regulated enterprise system may require smaller, auditable, task-specific encoders with clear evaluation records.

The main trade-off is between adaptation power and operational cost. Full fine-tuning can provide high task specificity but requires more data, compute, and governance. Feature extraction is cheap and stable but may not capture task-specific nuance. Parameter-efficient fine-tuning provides a middle path by adapting selected components while preserving the pretrained backbone.

Another important trade-off concerns control. Prompt-based zero-shot inference can be fast and flexible, but it may be difficult to validate and reproduce compared with supervised fine-tuning. Conversely, a fine-tuned classifier may be narrower but easier to evaluate, version, monitor, and audit.

The framework also highlights that transfer learning is not limited to supervised fine-tuning. It includes feature extraction, domain-adaptive pretraining, instruction tuning, adapter-based learning, prompt design, and few-shot inference. These approaches form a continuum from fully frozen model use to complete parameter adaptation.

---

## 8. Research Contribution

This study contributes a structured applied AI framework for transfer learning in NLP. The contribution is methodological and implementation-oriented rather than experimental.

The main contributions are:

1. **A practical NLP transfer learning decision framework**  
The study organizes model selection and adaptation around task type, dataset size, compute budget, domain specificity, and deployment constraints.
  2. **A component-level interpretation of pretrained model adaptation**  
The framework connects embedding layers, transformer blocks, and task heads to practical decisions about freezing, fine-tuning, and replacement.
  3. **A strategy taxonomy for applied transfer learning**  
The paper classifies feature extraction, head-only training, partial fine-tuning, full fine-tuning, adapter tuning, prompt tuning, instruction tuning, and zero-shot or few-shot inference.
  4. **A fine-tuning workflow for practitioners**  
The methodology recommends head-only baselines, discriminative learning rates, gradual unfreezing, validation-based early stopping, and domain-adaptive pretraining where appropriate.
  5. **A responsible-use perspective**  
The work emphasizes validation, transparency, domain fit, error analysis, and avoidance of overclaiming when adapting pretrained NLP models.
- 

## 9. Limitations

This study has several limitations.

First, the paper does not report new experiments, benchmark scores, statistical tests, or production deployment results. The framework therefore requires empirical validation for each dataset and use case.

Second, transfer learning performance depends heavily on data quality, label consistency, domain fit, preprocessing, class imbalance, and evaluation design. These factors cannot be resolved by model selection alone.

Third, the framework is intended for practical NLP design and does not provide a full mathematical treatment of transfer learning, optimization, representation learning, or generalization theory.

Fourth, pretrained language models may encode biases from their pretraining data. Fine-tuning can reduce, amplify, or redirect such biases depending on the downstream dataset.

Fifth, parameter-efficient methods such as LoRA and adapters reduce training cost but do not eliminate the need for validation, monitoring, or domain-specific evaluation.

Sixth, zero-shot and few-shot prompting can be useful for exploratory workflows, but results may vary across prompts, model versions, providers, and inference settings.

Finally, model behavior can change over time when using external API-based foundation models. This creates reproducibility and governance challenges for regulated or high-stakes applications.

---

## 10. Governance and Responsible Use

Transfer learning in NLP requires governance because pretrained models can produce confident but incorrect, biased, or unsafe outputs. Responsible implementation should include the following controls.

<b>Governance Area</b>	<b>Risk</b>	<b>Responsible Practice</b>
Bias	Model inherits or amplifies social and domain biases	Bias testing, representative data, human review
Transparency	Users may not understand model limitations	Model cards, documentation, decision logs
Privacy	Fine-tuning data may contain sensitive information	Data minimization, anonymization, access controls
Validation	Model may perform well on validation but fail in production	Out-of-distribution testing, monitoring, drift detection
Overclaiming	Transfer learning may be treated as automatically reliable	Report limitations and confidence boundaries
Reproducibility	API model behavior may change	Version tracking, fixed evaluation sets, audit records

Governance Area	Risk	Responsible Practice
Human oversight	Errors may affect users or decisions	Human-in-the-loop review for high-impact tasks

In high-stakes domains such as healthcare, law, finance, hiring, education, or public services, transfer learning should not be deployed solely on the basis of convenience or pretrained model reputation. It should be validated against domain-specific requirements and reviewed by relevant experts.

---

## 11. Future Work

Future work can extend this framework in several directions.

First, the methodology should be validated empirically across representative NLP tasks, including sentiment analysis, NER, extractive QA, summarization, and domain-specific classification.

Second, comparative studies should evaluate feature extraction, partial fine-tuning, full fine-tuning, LoRA, BitFit, prefix tuning, and adapters under controlled data and compute budgets.

Third, domain-adaptive pretraining should be studied across specialized corpora such as legal contracts, clinical notes, scientific papers, support tickets, and financial reports.

Fourth, the framework can be extended with reproducible implementation templates using open-source libraries such as Hugging Face Transformers, PEFT, AdapterHub, and evaluation toolkits.

Fifth, governance workflows should be formalized into model cards, risk assessments, monitoring dashboards, and deployment checklists.

Sixth, future research should examine how transfer learning interacts with retrieval-augmented generation, tool use, multimodal models, and enterprise knowledge systems.

---

## 12. Conclusion

This paper presented a practical framework for transfer learning in natural language processing using pretrained language models. The study examined how pretrained representations can be adapted to downstream tasks through feature extraction, fine-tuning, adapter-based methods, prompt tuning, instruction tuning, and zero-shot or few-shot inference.

The central conclusion is that transfer learning provides a powerful and practical foundation for NLP system development, but its value depends on disciplined model selection, careful adaptation, validation, and governance. Encoder models support understanding tasks, decoder models support generation, encoder-decoder models support text-to-text transformation, and domain-specific models improve suitability for specialized language distributions.

The contribution of this work is a decision-oriented methodology that helps practitioners connect task requirements with appropriate transfer learning strategies. The framework is not a substitute for empirical evaluation. Instead, it provides a structured starting point for building reliable, efficient, and responsible NLP systems with pretrained language models.

---

## References

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. arXiv.

Howard, J., & Ruder, S. (2018). *Universal Language Model Fine-tuning for Text Classification*. Proceedings of ACL 2018.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2021). *LoRA: Low-Rank Adaptation of Large Language Models*. arXiv.

Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., & Kang, J. (2020). *BioBERT: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining*. arXiv.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). *RoBERTa: A Robustly Optimized BERT Pretraining Approach*. arXiv.

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P. J. (2020). *Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer*. Journal of Machine Learning Research. arXiv.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). *Attention Is All You Need*. Advances in Neural Information Processing Systems.

Pruseth, D. (2026). *Transfer Learning for NLP*. Debabrata Pruseth AI blog.

## Suggested Citation

Pruseth, D. (2026). *Transfer Learning in Natural Language Processing: Practical Techniques with Pretrained Language Models*. Debabrata Pruseth AI Blog.