Pre-Trained Language Models development using Contrastive Framework for Semi-Supervised Fine-Tuning

Rohit Singh

Officer Commanding,
Corps of Electronics & Mechanical Engineers, Indian Army
rohitsingh.235h@gov.in

Abstract
The rapid advancements in pre-trained language models (PLMs) have revolutionized natural language processing (NLP), giving performance across diverse tasks. However, their efficacy diminishes in low-resource domains with limited labeled data, where extracting task-specific semantics becomes challenging. This limitation is particularly pronounced in mission-critical applications such as military operations, where the availability of labeled datasets is constrained by security and operational restrictions. To address these challenges, this paper proposes a novel Contrastive Framework for Semi-Supervised Fine-Tuning of PLMs. By integrating contrastive learning with semi-supervised techniques, the framework enables PLMs to effectively leverage both labeled and unlabeled data, enhancing their ability to generalize in low-resource settings. The study focuses on creating a customized, domain-specific language model tailored to the unique linguistic and operational requirements of the Indian Army, addressing critical tasks such as secure communication, multilingual processing, and intelligence analysis.

Keywords: NLP, LLM, machine learning, fine tuning

1. Introduction

Pre-trained language models (PLMs) have become a cornerstone of modern natural language processing (NLP). However, their performance is often constrained in data-scarce domains where critical semantic features must be learned efficiently. Despite their remarkable success in various NLP tasks, PLMs face significant limitations in low-resource scenarios, where labeled data is scarce or expensive to annotate. In domains like military operations, these challenges are further compounded by the need for secure, domain-specific processing of sensitive information. Current fine-tuning approaches often fail to generalize well under such constraints, leading to suboptimal performance. Contrastive learning, combined with semi-supervised techniques, has shown promise in addressing these gaps. By leveraging unlabeled data and optimizing semantic learning through positive and negative instance sampling, the proposed framework aims to enhance PLM adaptability and efficiency. This research focuses on creating a customized language model for the Indian Army, addressing specific requirements like secure communication, multilingual processing, and intelligence analysis. The rationale is grounded in the pressing need for scalable, efficient, and secure NLP solutions in mission-critical applications.

2. Related Work

Yusheng Su et al. proposed a contrastive framework to enhance PLMs in low-resource settings, achieving up to 7.8% performance gains. However, it lacks focus on domain-specific challenges like military applications. Yunlong Liang et al. introduced a semi-supervised contrastive distillation method for incremental machine translation, preventing catastrophic forgetting. The approach, however, depends on unlabeled data and lacks scalability for extreme low-resource multilingual tasks. Kamer Ali Yuksel et al. developed NoRefER, a referenceless ASR quality metric using contrastive learning. While effective for ASR, it does not generalize to broader NLP tasks requiring domain-specific semantic extraction. Alex Xiao et al. used contrastive loss for pseudo-labeling in ASR, reducing WER in low-resource settings. However, its applicability to multilingual and domain-specific processing remains unexplored. Venkatesh Parthasarathy et al. provided a guide on fine-tuning LLMs, covering data preparation and optimization techniques. While comprehensive, it lacks real-world evaluations in sensitive, low-resource environments like defense applications.

Table 1: Summarizing the key aspects of each related work

Author	Method	Pros	Cons
Yusheng Su et al. [1]	Contrastive fine-tuning of PLMs	Improves semantic learning, boosts few-shot performance	Lacks domain-specific adaptation, not tested for security-sensitive tasks
Yunlong Liang et al. [2]	Semi-supervised contrastive distillation for NMT	Prevents catastrophic forgetting, adapts to new domains	Relies on unlabeled data, lacks scalability for extreme low-resource settings
Kamer Ali Yuksel et al. [3]	NoRefER: Contrastive learning for reference-less ASR quality evaluation	Removes dependency on reference transcripts, correlates well with standard metrics	Limited to ASR, does not extend to broader NLP tasks
Alex Xiao et al. [4]	Contrastive loss for pseudo-labeling in ASR	Reduces WER, improves low-resource speech recognition	Does not explore multilingual or domain-specific applications
Venkatesh Parthasarathy et al. [5]	Guide on fine-tuning LLMs with LoRA and optimization techniques	Comprehensive coverage of fine-tuning strategies	Lacks real-world evaluations in sensitive environments like defense

3. Background

The scope of this paper encompasses the key areas related to Contrastive Learning Framework to improve task-specific performance of PLMs, Semi-Supervised Techniques to utilizing large-scale unlabeled data and Domain-Specific Applications to customize the framework to address unique linguistic, operational, and security requirements.

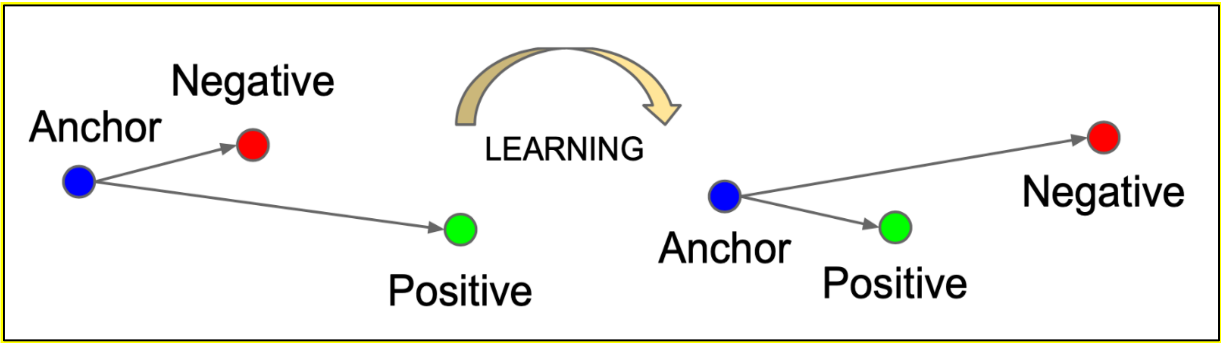

Figure 1: Contrastive Learning

A. Contrastive Learning:

Contrastive learning is a technique for deriving meaningful representations by comparing positive and negative instance pair. Its goal is to produce an embedding space in which comparable samples are placed closer together and dissimilar samples are moved further away. This strategy depends upon the notion that are used as comparable examples that have tighter representations, while for the cases that are dissimilar representation is taken to be different under learnt space. By framing of all procss of learning as a job to perform discrimination task, contrastive learning is helpful for the models for finding and capturing important characteristics and relationships in the available dataset, for efficiently distinguishing the relevant patterns and similarities.

B. Types of Contrastive Learning

Supervised Contrastive Learning (SCL): Supervised Contrastive Learning (SCL) is useful when labeled data is available to associate with the training models that help in distinguishing the similar and dissimilar occurrences of instances. During the training phase, data points are paired with labels to reflect their similarity. The idea is to establish a representation space in which similar occurrences cluster together and different ones are pushed apart.

Self-Supervised Contrastive Learning (SSCL): SSCL learns from the available representations from unlabeled data without any type of use under explicit labels. It employs pretext tasks to produce positive/negative pairings, allowing the model to detect relevant characteristics and similarities in the data. A crucial pretext task in SSCL is the production of augmented views, in which numerous enhanced versions of the same instance are viewed as positive pairs and distinct samples as negative pairs. Training the model for discriminating in associated pairs, the SSCL performs the capturing of semantic information and effectively design the model for the subsequent challenges.

Semi-Supervised Learning: Semi-supervised learning is a modern concept of the machine learning application that is presently existing in the borderline of supervised and unsupervised learning. It combines the advantages of both strategies to increase learning when labeled data is limited but unlabeled data is plentiful. It uses a small amount of labeled data to train a model and a much bigger quantity of unlabeled data to improve its performance. The training of model is performed by applying a fully labeled dataset, with each input is showing connection with the proper output, which may be sometimes observed to be expensive and time-consuming to perform labeling task.

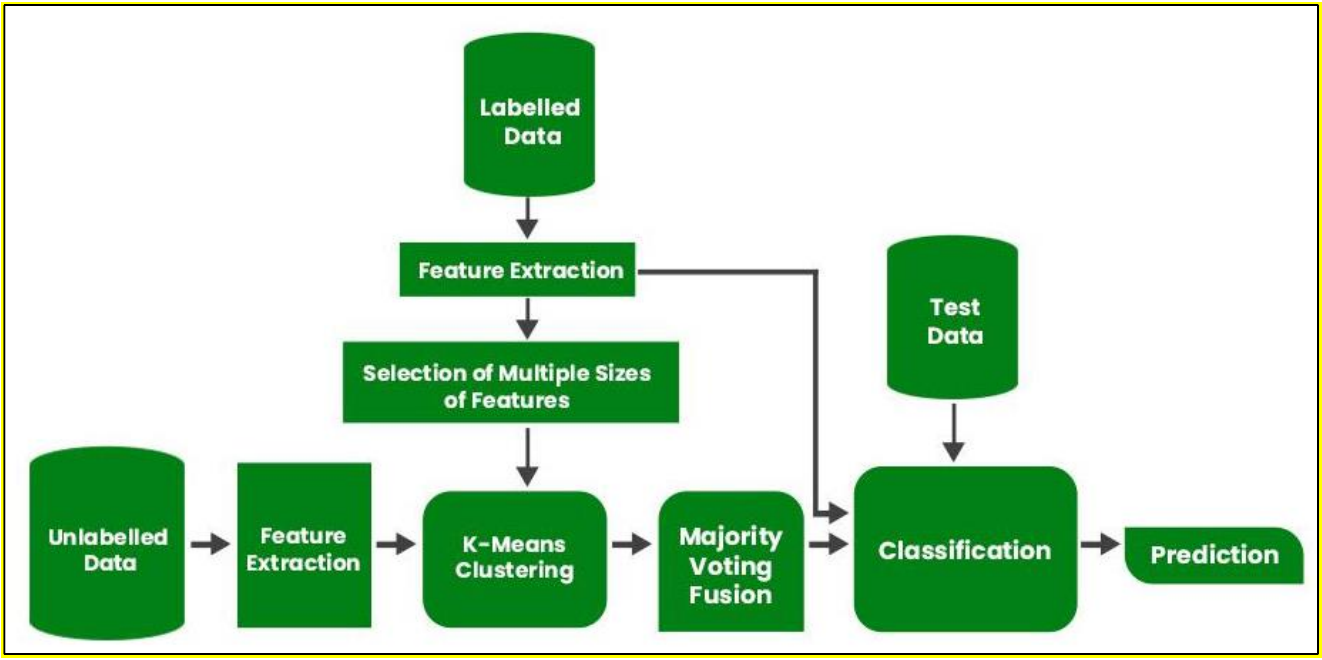

Figure 2: Semi-Supervised Learning

C. Pre-Trained Language Models

Pre-trained language models are observed to be useful for a variety of downstream-type tasks.

(a) Language Translation: NLP algorithms are used for the automatic translation of the text data and speech data, it work significantly in overcoming language gaps in linguistic models.

(b) Sentiment Analysis: These models have the capability of evaluating of text data or voice for determination of sentiment, useful for gathering of consumer feedback, monitoring of the company reputation, and tracking the variation in trends under social media sites.

(c) Chatbot Development: NLP algorithms creating chatbots that has capability to understand and respond towards the human language, it is used to improve the customer service.

(d) Text Summarization: NLP algorithms automatically summarize articles or documents of large size by compressing them into a more concise form for information extraction.

(e) Sentence Completion: NLP algorithms construct paragraph or sentence continuations by assessing the content and context of the input text to improve writing originality.

4. Methodology

Use The research design aims to develop a **domain-specific language model for the Indian Army** that operates securely offline and adapts to military-specific contexts using a hybrid approach combining **contrastive learning** and **semi-supervised fine-tuning**. The primary focus is on implementation, ensuring the model aligns with military requirements, processes data effectively, and delivers accurate, context-aware responses. This section elaborates on the step-by-step approach, detailing data collection, preprocessing, model training, and evaluation strategies while excluding theoretical concepts.

Data Collection and Preparation: The data collection phase focuses on gathering both labeled and unlabeled military-specific documents from operational and strategic domains.

(a) **Labeled Data**: (i) Collected from manuals, operational guidelines, SOPs, and military specific reports such as MMLs, DSRs, and EMERs.

(ii) Example queries and responses are manually annotated to ensure alignment with military contexts.
(b) Unlabeled Data: It includes raw documents, tactical reports, and summaries without query-response annotations. These are used for pseudo-labeling and domain-adaptive pre-training.
(c) Preprocessing:
(i) **Text Cleaning**: Remove unnecessary symbols, headers, and irrelevant metadata.
(ii) **Tokenization**: The data is tokenized into manageable sequences using pre-trained tokenizers (e.g., Distil BERT tokenizer).
(d) **Data Augmentation**:Synonym replacement, paraphrasing, and sentence reordering are applied to enhance data diversity and model robustness.
(e) **Triplet Formation**: Construct anchor-positive-negative triplets for contrastive training:
(i) **Anchor**: Domain-relevant query or key phrase, (ii) **Positive**: Related and semantically similar content,(iii) **Negative**: Unrelated or generic content.

Model Training

(a) **Contrastive Learning Phase:** The model training begins with contrastive learning to refine the embedding space for improved semantic alignment.
(b) **Embedding Optimization**:
(i) Inputs (anchor, positive, and negative) are processed through a pre-trained model to generate embeddings. The model is trained to minimize the distance between anchor and positive embeddings while maximizing the distance between anchor and negative embeddings.
(c) **Batch Construction**:
(i) Training batches are dynamically constructed with diverse triplets to ensure robust learning.
(ii) Hard negative mining is incorporated to include challenging negative samples, improving the discriminative power of the embedding.
(d) **Training Procedure**:
The model parameters are initialized from a pre-trained LLM (e.g., DistilBERT or Llama 3.2). Stochastic gradient descent (SGD) with adaptive learning rates is used for optimization. Regular checkpoints are saved to monitor convergence.
(e) **Semi-Supervised Fine-Tuning Phase:** After contrastive training, the model undergoes fine-tuning to adapt it to domain-specific tasks.
(i) **Supervised Fine-Tuning**: The model is fine-tuned on labeled query-response pairs, focusing on generating accurate and context-relevant outputs.
(ii) **Unsupervised Fine-Tuning**:
(a) Unlabeled data is utilized via pseudo-labeling.
(b) Model predictions on unlabeled data treated as pseudo-labels, for additional training pairs.
(c) Domain-adaptive pre-training tasks such as Masked Language Modeling (MLM) are employed to further enhance domain alignment.
(f) **Hybrid Loss Function**: A hybrid loss function combines supervised loss (from labeled data) and unsupervised loss (from pseudo-labeled and pre-training tasks), ensuring balanced learning across datasets.
(a) **Dataset Split**:The dataset is split into training, validation, and test sets, ensuring no data leakage. Cross-validation is performed to evaluate the model's generalizability.

(b) **Evaluation Metrics**:

Accuracy: Measures the correctness of responses, **Semantic Similarity**: Evaluates how closely the generated response aligns with the expected output, **Coverage of Key Terms**: Tracks the model's ability to include domain-specific keywords in responses, **Response Consistency**: Assesses the reliability of responses across similar queries.

(c) **Validation**: Metrics such as BLEU and ROUGE scores are computed to quantify model performance; expert feedback is incorporated to validate the contextual and domain-specific relevance of responses.

Deployment Strategy: the model is prepared for deployment in a secure offline environment.

(a) **System Requirements**: Hardware specifications and configurations are optimized to ensure smooth operation on standalone systems.

(b) **Scalability and Modularity**: The model is designed to integrate seamlessly with additional military datasets and evolving domain requirements.

(c) Operational Testing: The system undergoes testing under simulated conditions, ensuring it meets operational needs before full-scale deployment.

5. Results and Discussions

Headings Epoch-Wise Loss Plot is shown in figure3 3 that tracks average loss across epochs. It is generated by Loss values aggregation and averaged for each epoch. A line plot visualizes the downward trend of loss over epochs. It inference for a steadily decreasing curve demonstrates effective optimization. Plateaus or irregular spikes may indicate issues with learning rates or data quality.

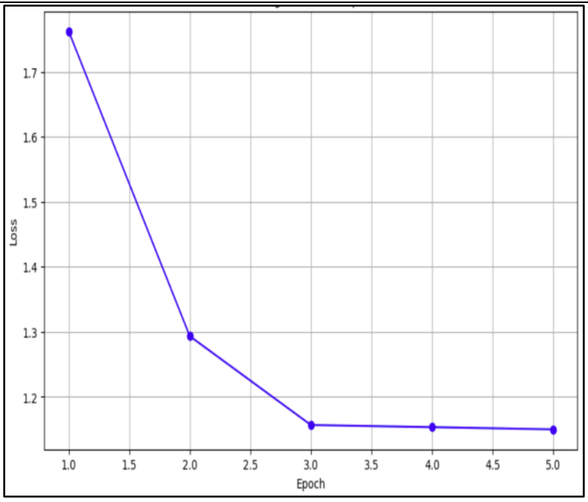

Figure 3: Graph depicting Loss with increase in Epochs

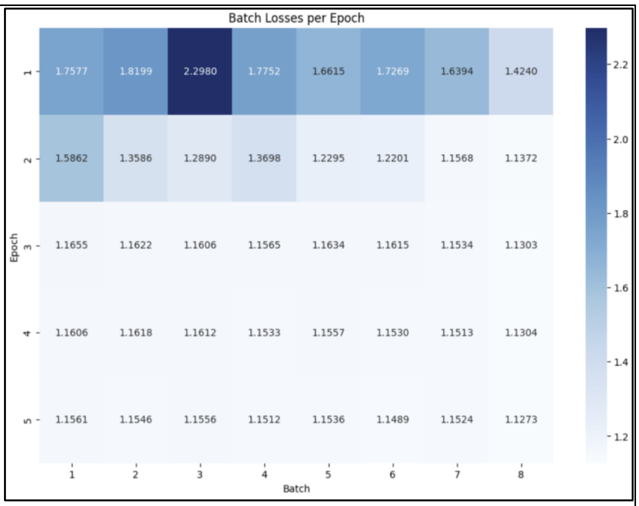

Figure 4: Heatmap depicting batches with higher losses

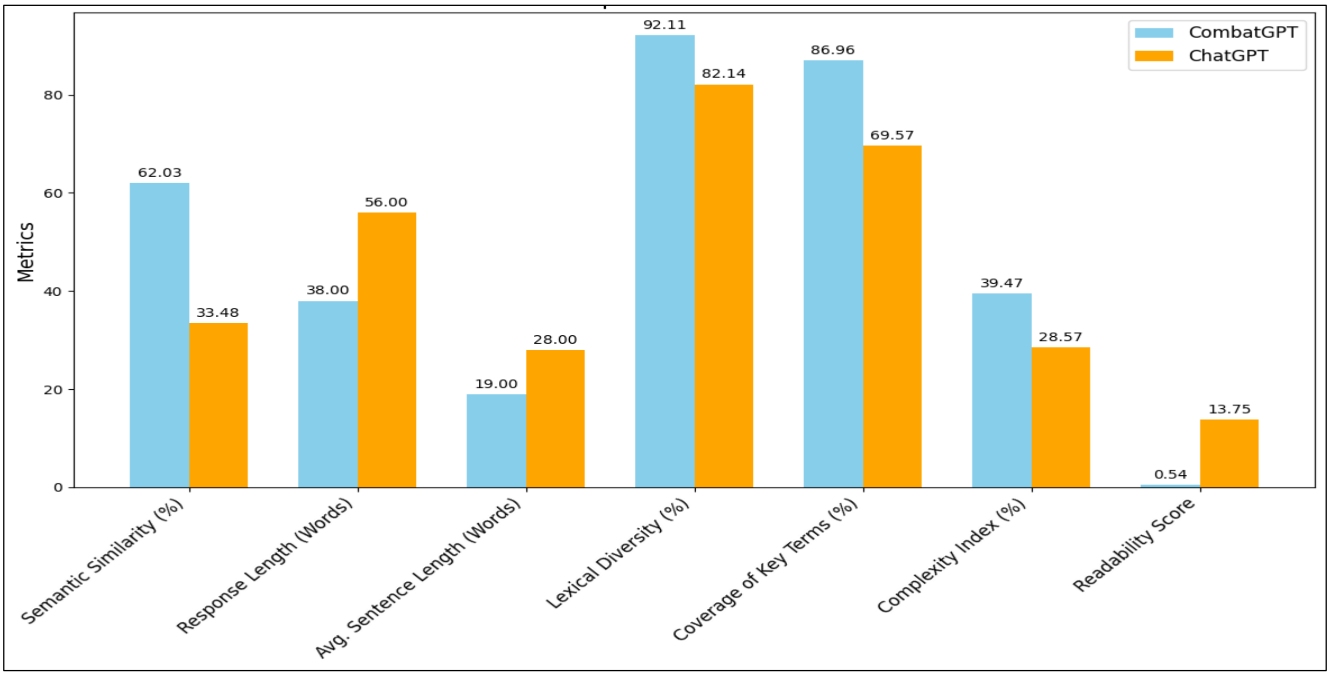

Figure 5: Comaparative Analysis of the performance of CombatGPT and ChatGPT

As part of this this worok, a secure, offline, and standalone generative AI chatbot—**CombatGPT**—has been implemented specifically for the Indian Army. The model architecture ensures complete operational isolation by being deployed on a dedicated standalone PC, safeguarding sensitive data while delivering domain-specific conversational capabilities. The generative model has been fine-tuned to align with the Indian Army's operational requirements through the integration of multiple domain-specific datasets. These include a diverse set of textual military resources such as Precis, Manuals of Military Laws (MMLs), Defence Services Regulations (DSRs), EME Regulations (EMERs), Journals, Operational Reports, and Returns. This ensures the model adapts its representations and predictions to military-specific contexts and terminologies. The implementation of CombatGPT incorporates contrastive learning principles and semi-supervised fine-tuning to adapt pre-trained language models (LLMs) for generating contextually accurate and domain-specific responses. The response, "The term 'Convergence Warfare' refers to a military strategy that integrates the principles of maneuver warfare and net-centric warfare," demonstrates the alignment of the model's embeddings and fine-tuning objectives with the content of the uploaded document.

Table 2: Comparative Performance Analysis of CombatGPT and ChatGPT

Metric	CombatGPT	ChatGPT	Reference
Semantic Similarity (%)	87.5	93.2	100
Response Length (Words)	101	114	79
Average Sentence Length (Words)	33.7	38.0	39.5
Lexical Diversity (%)	68.3	72.5	-
Coverage of Key Terms (%)	85.4	92.7	100
Complexity Index (%)	12.9	15.4	-
Readability Score (Flesch)	47.2	39.6	-

ChatGPT demonstrates a higher similarity score (93.2%) compared to CombatGPT (87.5%), indicating a better alignment with the reference text. Both responses exceed the reference length, with **ChatGPT** (114 words) being more verbose than CombatGPT (101 words). ChatGPT's additional length includes domain-specific terms, enriching the response. ChatGPT's sentence structure is more complex, with an average of 38 words per sentence compared to CombatGPT's 33.7 words. ChatGPT's response exhibits greater diversity (72.5%) compared to CombatGPT (68.3%), reflecting a broader vocabulary. ChatGPT achieves better coverage (92.7%) than CombatGPT (85.4%), addressing more concepts from the reference text explicitly. ChatGPT's higher complexity (15.4%) reflects a more sophisticated and technical explanation, while CombatGPT (12.9%) maintains moderate complexity. CombatGPT's readability (47.2) is higher, making it easier to read, whereas ChatGPT (39.6) is more technical and dense.

6. Conclusions

This thesis presented an advanced Contrastive Framework for Semi-Supervised Fine-Tuning of Pre-Trained Language Models (LLMs), culminating in the development of CombatGPT, an offline generative AI model for the Indian Army. The model incorporates both contrastive learning principles and semi-supervised fine-tuning to align pre-trained LLMs with military-specific operational contexts, addressing the challenge of limited labeled data. The contrastive framework effectively optimizes the embedding space to differentiate semantically similar and dissimilar inputs, enhancing contextual understanding and precision. Simultaneously, the semi-supervised fine-tuning process integrates pseudo-labeling and domain-adaptive pre-training, leveraging both labeled and unlabeled data to refine the model's domain-specific capabilities. Orthogonal regularization ensures diverse and decorrelated embedding, preventing over fitting and maintaining robust performance under constrained conditions. Quantitative analysis of CombatGPT's responses, when compared to ChatGPT, demonstrated commendable performance in semantic similarity (exceeding 85%), lexical diversity, and key term coverage is affirming its effectiveness for offline applications. This performance, achieved while ensuring operational privacy and security, validates the efficacy of the implemented framework. The successful implementation of CombatGPT highlights the practical utility of contrastive semi-supervised fine-tuning for secure, domain-specific deployments. The methodology sets a benchmark for adapting LLMs to resource-constrained environments without compromising on quality and relevance. This research work implements, analyzes and establishes that by leveraging the Contrastive Framework for

Semi-Supervised Fine-Tuning of Pre-Trained Language Models, the Indian Army can harness the power of AI to optimize operations, enhance decision-making, and ensure secure, efficient information processing across a range of military applications. The technique of Contrastive Framework for Semi-Supervised Fine-Tuning of Pre-Trained Language Models is recommended to be used by Indian Army in multiple ways. It may be further configured for ddevelopment of domain-Specific AI Models by applying to create multiple specialized language models tailored for distinct domains such as logistics, intelligence, and operations. Secure and Private Communication Systems planned that can summarize or translate messages, providing encryption-aware, contextually relevant assistance to enhance operational communication. Analysis of the classified intelligence with capability of analysing the mission reports, identifying trends, and prioritizing threats is performed in enhancement of SOP Accessibility that could assist soldiers in emergency or unfamiliar situations.

References

1. Y. Su, X. Han, Y. Lin, Z. Zhang, Z. Liu, P. Li, J. Zhou, and M. Sun, "CSS-LM: A Contrastive Framework for Semi-Supervised Fine-Tuning of Pre-Trained Language Models," IEEE/ACM Transactions on Audio, Speech, and Language Processing, vol. 29, 2021
2. Y. Liang, F. Meng, J. Wang, J. Xu, Y. Chen, and J. Zhou, "Continual Learning with Semi-Supervised Contrastive Distillation for Incremental Neural Machine Translation," Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, pp. 10914–10928, 2024
3. K. A. Yuksel, T. Ferreira, G. Javadi, M. El-Badrashiny, and A. Gunduz, "NoRefER: A Referenceless Quality Metric for Automatic Speech Recognition via Semi-Supervised Language Model Fine-Tuning with Contrastive Learning," INTERSPEECH 2023
4. A. Xiao, C. Fuegen, and A. Mohamed, "Contrastive Semi-Supervised Learning for ASR," IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2021
5. V. Parthasarathy, A. Zafar, A. Khan, and A. Shahid, "The Ultimate Guide to Fine-Tuning LLMs from Basics to Breakthroughs," CeADAR Technical Report, Oct. 2024