Supervised Contrastive Fine-Tuning for Active Few-Shot Learning

Zirui Zhang, Lei Ge*, Shengyu Qiao
Data and Target Engineering College, Information Engineering University
No.62, Science Avenue, Zhengzhou 450001, Henan, China
zhangzr1029@163.com, chxy_gray@163.com,3079837836@qq.com

Abstract

Active Few-Shot Learning (AFSL) is an effective paradigm for improving the performance of large language models under limited annotation budgets. To address the inefficiency of conventional fine-tuning objectives in AFSL, this paper proposes a supervised contrastive fine-tuning framework specifically designed for natural language processing (NLP) text classification tasks. By integrating Supervised Contrastive Learning (SCL) with Hard Negative Mining (HNM), the proposed framework optimizes the embedding space through an enhanced hybrid loss function, thereby improving the utilization efficiency of labeled samples. Extensive experiments on five benchmark datasets show that, under a fixed state-of-the-art (SOTA) query strategy, our method consistently outperforms baseline models in text classification performance, and exhibits strong generalizability across different backbone architectures and acquisition functions. These findings demonstrate that optimizing how to learn—through improved learning objectives—provides a complementary direction to existing query strategies in advancing AFSL.

Keywords: Active Few-Shot Learning, Supervised Contrastive Learning, Text Classification, Representation Learning

1. Introduction

Pre-trained Language Models (PLMs) (GAO et al., 2021a; CHEN et al., 2021; KARIMI MAHABADI et al., 2022; LIN et al., 2022) have substantially advanced NLP but still rely on labeled data, making low-resource scenarios challenging. Few-Shot Learning (FSL) (BROWN et al., 2020; WANG et al., 2020; EDWARDS and CAMACHO-COLLADOS, 2024; CHANG et al., 2021) and its active variant address this by selecting informative instances for annotation under tight labeling budgets. We focus on Active Few-Shot Learning for text classification, where robust learning from few labeled samples is crucial.

AFSL iteratively selects the most informative unlabeled instances for annotation and updates the model, thus improving accuracy under budget constraints. Prior work has primarily optimized the query strategy—the "what to learn" component—(LONG et al., 2023; YUAN et al., 2020; EIN-DOR et al., 2020; AHMADNIA et al., 2025; ZHAN et al., 2022). However, the learning objective after acquisition—"how to learn"—has received far less attention, with most AFSL pipelines defaulting to cross-entropy (CE) (ZHAN et al., 2022). This choice often yields suboptimal embedding structures for few-shot updates, limiting the utility of each labeled example and forming a practical bottleneck.

CE emphasizes per-instance classification but provides limited metric structure for incremental few-shot updates (GUNEL et al., 2021). In contrast, supervised contrastive learning (SCL) (KARIMI MAHABADI et al., 2022) enforces intra-class compactness and inter-class separation, aligning the representation geometry with AFSL's incremental regime. We thus hypothesize that replacing or augmenting CE with SCL yields consistent gains. To rigorously isolate the impact of the learning objective from the confounding variables of data acquisition, we fix the query strategy to a strong state-of-the-art baseline in our main experiments. While this controlled setup demonstrates the potential of optimizing "how to learn," we acknowledge that exploring the interaction between varied query strategies and SCL-HNM is an important avenue for future generalization.

The main contributions of this paper can be summarized as follows:

- We pinpoint the neglected "how to learn" objective as a key bottleneck in AFSL.
- We introduce a SCL-based fine-tuning objective that drops into existing AFSL pipelines without altering the query strategy.
- Extensive experiments on multiple text classification benchmarks and supplementary evaluations on alternative architectures and query strategies demonstrate consistent gains over CE across acquisition rounds.
- To support reproducibility and align with the resource-sharing spirit of LREC, we make our source code and the specific pre-processed data splits publicly available to the research community.

*Corresponding author.

Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026), pages 10365–10375
11–16 May 2026. ©ELRA Language Resources Association (ELRA), 2026

2. Related Work

2.1. Active Few-Shot Learning

Active Few-Shot Learning (AFSL) centers on query strategies that select the most informative samples from large unlabeled pools for annotation. Research has evolved from early representativeness-based approaches (LOGESWARAN and LEE, 2018; KARIMI MAHABADI et al., 2022; LIN et al., 2022; GAO et al., 2021a) to methods that exploit embedding-space information and model uncertainty (AHMADNIA et al., 2025). Together, these strategies have refined the process of sample acquisition, advancing the "what to learn" question in AFSL.

However, these advances primarily address sample selection rather than model optimization, leaving the equally important "how to learn" question underexplored. Prior studies report that most AFSL frameworks still rely on cross-entropy (CE) loss for model updates after each query iteration (YUAN et al., 2020; EIN-DOR et al., 2020; ZHAN et al., 2022). CE optimizes instances independently, focusing solely on classification accuracy without enforcing structural constraints on the embedding space (GUNEL et al., 2021; KHOSLA et al., 2021). This limitation is amplified in AFSL's iterative fine-tuning, where each round introduces only a few new labeled examples. Without structural regularization, the representation space often exhibits blurred inter-class boundaries and weak intra-class compactness, which hinders efficient sample utilization even for strong query strategies. This inefficiency forms a central bottleneck in AFSL, motivating the exploration of alternative objectives that can explicitly model relationships among samples, as discussed next in contrastive learning.

2.2. Contrastive Learning

To address the limitations of CE-based optimization, recent research has increasingly explored contrastive learning objectives that explicitly model pairwise relations in the embedding space. Because AFSL involves incremental fine-tuning under limited supervision, objectives that impose such structure are particularly beneficial. Supervised Contrastive Learning (SCL) (KHOSLA et al., 2021) shapes the embedding geometry by pulling positive samples together and pushing negatives apart, thereby producing more discriminative representations in data-scarce settings (BAIK et al., 2021; YIN et al., 2021; LI et al., 2020). When applied to pre-trained language models, SCL has shown clear improvements over CE in standard few-shot fine-tuning (GUNEL et al., 2021; KARIMI MAHABADI et al., 2022).

Building upon the SCL formulation, researchers have further enhanced discriminative power through hard negative mining (HNM) (LONG et al., 2023). Hard negatives—samples highly similar to the anchor in feature space—are essential for refining decision boundaries. For example, SCHaNe by LONG et al. (2023) emphasizes these cases by assigning higher weights to difficult negatives. Such extensions have proven effective in standard few-shot fine-tuning and remain an active research direction in contrastive learning.

Despite this progress, adapting SCL and HNM to AFSL's iterative setting remains underexplored. The AFSL paradigm introduces unique challenges: (i) each iteration provides only a few new labeled samples, demanding objectives that balance stability and adaptability; (ii) PLMs bring strong priors but also anisotropy in embedding space (ZHAN et al., 2022); and (iii) linguistic semantics introduce high intra-class variance, making hard negative definition more nuanced (GAO et al., 2021b; HE et al., 2020). Addressing these challenges, this work integrates Supervised Contrastive Learning with Hard Negative Mining (SCL-HNM) into AFSL, aiming to bridge the gap between representation structure and active sample selection, and offering a unified perspective on "how to learn" under limited supervision.

3. Methods

This section introduces a supervised contrastive fine-tuning framework for Active Few-Shot Learning (AFSL), designed to address the bottleneck of inefficient sample utilization. In contrast to prior studies that primarily focus on query strategies, our method optimizes the learning objective itself, aiming to enhance sample efficiency through structured representation learning.

3.1. Task Definition

Our framework follows the standard active-learning protocol. Starting with a large unlabeled pool U, the system gradually constructs a labeled support set S through multiple iterations. Each iteration consists of three steps: (1) the query strategy $\mathcal{A}$ selects a batch of M unlabeled samples for annotation, (2) the annotated samples are incorporated into the support set, and (3) the model is fine-tuned on the updated support set using the proposed learning objective.

To rigorously evaluate the impact of the learning objective, we adopt an iterative evaluation protocol summarized in Algorithm 1. A model f_θ is trained for T iterations, with performance assessed whenever the support set reaches a predefined size $N_{\max}$. At each checkpoint, the model is reinitialized and fine-tuned from scratch on S to ensure fair comparison across iterations.

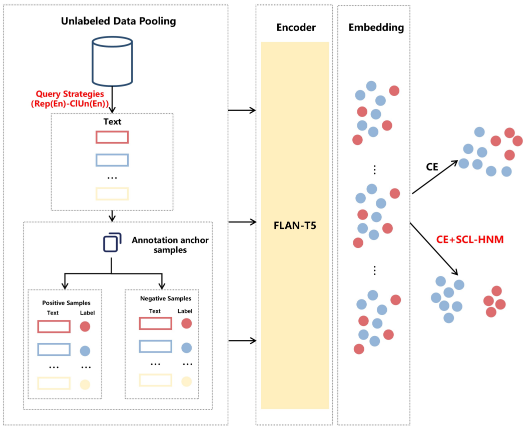

Figure 1: Overview of the proposed AFSL pipeline. Each iteration includes three stages: (1) sample querying using a fixed strategy; (2) support-set expansion with newly annotated samples; and (3) model update via supervised contrastive fine-tuning. Evaluation is conducted on a held-out test set.

Algorithm 1 Iterative Evaluation Protocol for AFSL

Require: Unlabeled pool U; initial support set S_0; model f_θ; fixed query strategy $\mathcal{A}$; iterations T; batch size M; target support size $N_{\max}$; held-out test set $\mathcal{D}_{\text{test}}$.
Ensure: Performance on $\mathcal{D}_{\text{test}}$.

$S \leftarrow S_0$
for $t = 1$ **to** T **do**
  $X_{\text{sel}} \leftarrow \mathcal{A}(U, M)$
  $Y_{\text{sel}} \leftarrow \textsc{Annotate}(X_{\text{sel}})$
  $S \leftarrow S \cup \{(x, y) \mid x \in X_{\text{sel}}, y \in Y_{\text{sel}}\}$
  $U \leftarrow U \setminus X_{\text{sel}}$
  $f_\theta \leftarrow \textsc{FineTune}(f_\theta; S)$
  if $|S| \geq N_{\max}$ **then break**
  end if
end for
$f_\theta \leftarrow \textsc{Reinitialize}()$
$f_\theta \leftarrow \textsc{TrainFromScratch}(f_\theta; S)$
return $\textsc{Evaluate}(f_\theta, \mathcal{D}_{\text{test}})$

To isolate the influence of our learning objective, all experiments employ a fixed query strategy. Specifically, we adopt the state-of-the-art Rep(En)-ClUn(En) (AHMADNIA et al., 2025), which balances representativeness in the embedding space with model uncertainty. Thus, this study focuses on improving the *model update* step rather than proposing a new querying algorithm.

3.2. Model Selection

For our experiments, we adopt `google/flan-t5-base` as the backbone model f_θ. This model has demonstrated strong few-shot performance across various NLP tasks, providing a competitive baseline for our study. To obtain high-quality sentence-level representations, we follow standard practice and utilize its encoder component. For each input sequence x_i, we extract the hidden states of all tokens from the final encoder layer and apply mean pooling to produce a fixed-dimensional sentence embedding $h_i \in R^d$. This embedding h_i serves as a comprehensive semantic representation of x_i and acts as the basic unit for all subsequent learning objectives.

To ensure rigorous and comparable evaluation across different support set sizes ($K = 10, 20, \ldots, 100$), we adopt a strict *weight-reset policy*. Before fine-tuning at each evaluation checkpoint K, the model parameters are reset to the original pre-trained weights θ_0. The model is then fine-tuned exclusively on the current support set S_t. This procedure ensures that any performance improvement arises solely from the quantity and quality of the samples in S_t, rather than from cumulative effects of continuous training. Consequently, it provides a fair and isolated assessment of our learning method's efficacy. After obtaining stable representations from f_θ, we now focus on how to optimize them effectively through supervised contrastive fine-tuning.

3.3. Supervised Contrastive Fine-tuning with HNM

Before presenting our hybrid training objective, we clarify a key architectural decision. Many contrastive-learning studies add a non-linear projection head to the encoder output. We depart from this design and remove the projection head. Instead, we optimize the core semantic representations produced by the encoder f_θ.

Parameter efficiency. In few-shot settings, extra trainable parameters increase the risk of overfitting. Avoiding them improves stability.

Direct optimization. Our goal is to refine the very representation used by the downstream classifier. This keeps the model simple and the optimization path direct under limited data.

We acknowledge a trade-off between parameter efficiency and representation disentanglement. In AFSL, where labeled data are extremely scarce, preventing overfitting has higher priority. We therefore accept the loss of a projection head's potential disentanglement benefits in exchange for greater stability.

To improve sample utilization and obtain a well-structured embedding space, we adopt a hybrid objective. It combines a conventional classification loss with a representation regularizer. We keep cross-entropy (CE) as the primary objective to drive predictive accuracy. We then add a supervised contrastive regularizer to enforce intra-class compactness and inter-class separability. This division of roles aligns the classifier with task accuracy while shaping the geometry of the embedding space.

The supervised contrastive loss (KHOSLA et al., 2021) is defined for each anchor sample i as:

$$\mathcal{L}_{\text{SCL}} = -\frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(h_i \cdot h_p/\tau)}{\sum_{a \in A(i)} \exp(h_i \cdot h_a/\tau)} \tag{1}$$

where $P(i)$ is the set of positive samples sharing the same label as i, $A(i)$ includes all other samples in the batch, and τ is the temperature parameter controlling distribution sharpness. This loss pulls embeddings of the same class together while pushing those of different classes apart.

To enhance inter-class discrimination, we integrate a *hard negative mining* (HNM) mechanism into the supervised contrastive loss. For each anchor i we sort all negative samples in the batch by their similarity to z_i (cosine similarity between the L2-normalized embeddings) and select the top-k to form the hard negative set $H(i)$. These hard negatives are then included alongside the positives in the denominator of the loss (Equation 2), encouraging the model to pay particular attention to the most confusable cases.

The resulting *supervised contrastive with hard negatives* (SCL-HNM) loss for an anchor i is defined as

where $\mathrm{sim}(u, v) = u^\top v$ denotes cosine similarity between L2-normalized embeddings, $P(i)$ is the set of positives for anchor i, and $H(i)$ is the set of the k hardest negatives selected as described above. The batch-level loss $\mathcal{L}_{\text{SCL-HNM}}$ is obtained by averaging $\mathcal{L}^{(i)}_{\text{SCL-HNM}}$ over all anchors in the mini-batch.

Finally, the overall training objective is a weighted combination of cross-entropy and the proposed SCL-HNM loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \alpha\, \mathcal{L}_{\text{SCL-HNM}}, \tag{3}$$

where $\alpha > 0$ is a tunable hyperparameter that balances classification accuracy and representation structuring. This hybrid objective allows the model to maintain predictive stability from $\mathcal{L}_{\text{CE}}$ while gaining geometric consistency from $\mathcal{L}_{\text{SCL-HNM}}$.

4. Experiments Setup

4.1. Datasets

To comprehensively evaluate the effectiveness and generalization ability of our proposed method, five benchmark datasets were selected to cover diverse text classification scenarios. Detailed statistics for each dataset are summarized in Table 1, and representative task examples are shown in Table 2.

Dataset	Multi-Label	Train	Val	Test	#Class	$U\%$
MPQA Type		4,248	1,060	1,327	4	85.1
MPQA Polarity	×	4,505	1,123	1,404	2	8.9
MPQA Intensity	×	4,505	1,123	1,404	5	34.6
AG News	×	118,800	1,200	7,600	4	0.0
Amazon Reviews	×	200,000	1,200	5,000	5	0.0

Table 1: Statistics of the benchmark datasets, including label non-uniformity $U\%$ as defined in Equation 4.

The column $U\%$ in Table 1 measures the degree of label non-uniformity. Following MÜLLER et al. (2022), we compute it as

$$U = \sum_{l \in L} \left| f(l) - \tfrac{1}{|L|} \right|, \tag{4}$$

where L denotes the set of labels and $f(l)$ is the relative frequency of label l in the dataset. A value of $U = 0$ indicates a perfectly uniform distribution, while larger values reflect greater imbalance.

Dataset	Task	Input Sentence	Output
MPQA-T	Opinion classification (type)	"We don’t hate the sinner," he says, "but we hate the sin."	sentiment
MPQA-P	Opinion classification (intention)	My public affairs keepers couldn’t care less.	negative
MPQA-I	Opinion classification (Sentiment)	European countries are critical of Bush’s "axis of evil" remark.	high
AG News	News classification (topic)	New technology applies electrical fuses to help identify faults.	Sci/Tec
Amazon Reviews	Sentiment analysis (amazon)	The instructions were confusing, but my kid loves this toy!	5 (stars)

Table 2: Representative task examples for each benchmark dataset.

The MPQA Opinion Corpus (WIEBE et al., 2005) provides fine-grained opinion annotations. Following the dataset revision and task formulation by AHMADNIA et al. (2024), we construct three tasks: *Type* (multi-label), *Polarity* (binary), and *Intensity* (five-class).

AG News (AGN, 2005) is a topic-classification corpus; as no validation split is provided, we randomly sample $1{,}200$ examples from the training set for validation.

Amazon Reviews (Ama, 2020) is used for 1–5 star sentiment prediction. We employ its English subset and down-sample $1{,}200$ validation instances to avoid overfitting in few-shot fine-tuning.

Dataset selection was guided by three criteria: (1) **Task diversity**—covering binary, multi-class, and multi-label classification; (2) **Semantic granularity**—including both coarse-grained topics (AG News) and fine-grained sentiment distinctions (Amazon Reviews, MPQA Intensity); (3) **Distributional challenge**—datasets with varying $U\%$ values enable evaluation of robustness under label imbalance.

4.2. Baselines

To evaluate the effectiveness of the proposed approach, we define two core baselines for comparison. These baselines are designed to isolate

$$\mathcal{L}^{(i)}_{\text{SCL-HNM}} = -\frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\text{sim}(z_i, z_p)/\tau)}{\sum_{p' \in P(i)} \exp(\text{sim}(z_i, z_{p'})/\tau) + \sum_{h \in H(i)} \exp(\text{sim}(z_i, z_h)/\tau)}. \quad (2)$$

the specific contributions of individual components within the AFSL framework.

Random + CE. This basic baseline adopts a random sampling strategy for instance selection and employs the standard cross-entropy (CE) loss for model fine-tuning. It represents a naive configuration without informed query selection or an advanced learning objective. This baseline establishes the performance floor and serves as a reference to quantify the improvements achieved by more sophisticated methods.

SOTA Query + CE. This is the primary and most competitive baseline. It combines the state-of-the-art (SOTA) query strategy Rep(En)-ClUn(En) (AHMADNIA et al., 2025) with the conventional CE loss. The strategy is highly effective, as it jointly considers embedding-space representativeness and model uncertainty to select informative samples. Comparison with this baseline enables us to isolate the contribution of the proposed learning objective. By keeping the SOTA query strategy fixed, any observed performance improvement can be attributed solely to the superiority of our proposed SCL-HNM loss over the standard CE objective.

Our Full Method (SOTA Query + SCL-HNM). Our complete method also employs the same SOTA query strategy but replaces the conventional CE loss with the hybrid objective that integrates CE and SCL-HNM. This setting allows a direct evaluation of the effectiveness of the proposed fine-tuning objective under an identical query configuration.

The following subsection presents the fine-tuning experiments and evaluation protocol used to compare these configurations.

4.3. Fine-Tuning Experiment

All experiments follow the standard Active Few-Shot Learning (AFSL) iterative evaluation protocol. The goal is to systematically assess the effectiveness of different learning objectives under various support set sizes $|S|$. A strict *weight-reset policy* is enforced to guarantee fairness and independence across iterations.

Specifically, after each query iteration t, the model parameters f_θ are re-initialized to the same pre-trained state before the next fine-tuning phase begins. This prevents knowledge accumulation across iterations and ensures that performance reflects the learning efficiency of each objective rather than residual effects from previous rounds.

Each iteration proceeds as follows:

1. The query strategy $\mathcal{A}$ selects a batch of M unlabeled samples from pool U.

2. These samples are annotated to obtain the corresponding labels.

3. The updated support set S_t is used to fine-tune f_θ under the given learning objective.

4. Evaluation is conducted on a fixed test set $\mathcal{D}_{\text{test}}$ to record performance metrics.

This controlled setting isolates the contribution of the learning objective by removing confounding factors such as query order or accumulated gradient history. The experimental workflow is summarized in Algorithm 1.

4.4. Implementation Details

Evaluation Metrics. Model performance is primarily measured using the Micro-F1 score, which reflects both precision and recall across all categories. Each AFSL experiment is repeated with five different random seeds to ensure robustness. We report the mean and standard deviation of all results.

Experimental Environment. All experiments are conducted on a single NVIDIA RTX 4090 GPU (24 GB memory). Our implementation is built upon the PyTorch framework, and we use the Hugging Face Transformers library to load and fine-tune the `google/flan-t5-base` model.

Hyperparameter Configuration. The main hyperparameters for model fine-tuning are summarized in Table 3. Three key parameters of the proposed SCL-HNM loss—the loss weight α, the temperature τ, and the number of hard negatives k—are determined through a systematic sensitivity analysis on the validation split of the MPQA Polarity dataset. This procedure aims to evaluate the robustness of our approach and identify a single, stable configuration with strong generalization ability for all subsequent experiments. We employ the AdamW optimizer and apply early stopping based on validation performance, with a patience of 20 epochs, to avoid overfitting. To ensure full reproducibility, random seeds are fixed for all PyTorch, NumPy, and data-splitting operations.

Parameter	Value
Model	`google/flan-t5-base`
Batch size	10
Learning rate	1×10^{-4}
Dropout rate	0.1
Optimizer	AdamW
Early stopping patience	20 epochs
Loss weight α	0.1
Temperature τ	0.05
Hard negatives k	5

Table 3: Key hyperparameters used for model fine-tuning.

5. Results and Analysis

5.1. Main Results

We evaluate the proposed **Supervised Contrastive Learning with Hard Negative Mining (SCL-HNM)** objective across multiple data scales to assess its effectiveness and generalization capability. Experiments are conducted on five benchmark datasets and compared against strong baselines. Table 4 reports Micro-F1 (%) as the mean $\pm$ standard deviation over five random seeds.

Performance Upper Bound with Full Data. Under the full-data setting, replacing the conventional CE loss with **SCL-HNM** consistently improves performance across all datasets (e.g., MPQA-Polarity: 95.8 vs. 94.2; MPQA-Intensity: 51.3 vs. 50.0). This demonstrates that explicit structural constraints in the embedding space benefit standard supervised fine-tuning beyond the active few-shot regime.

Independent Contribution of the Objective. To isolate the effect of the learning objective, we compare CE and **SCL-HNM** under random sampling. **SCL-HNM** generally performs better, especially as K increases (e.g., MPQA-Polarity at $K = 100$: 90.3 vs. 88.4). A slight decline at very small K (e.g., MPQA-Type) suggests that minimal informative supervision is necessary to stabilize representation learning under extreme scarcity.

Synergy with a State-of-the-Art Query Strategy. Comparing **Ours (SOTA Query + SCL-HNM)** with **SOTA Query + CE** yields higher or comparable results on four of five datasets. The only exception is AG News at $K \in \{10, 20\}$, where topical diversity favors CE. Overall, optimizing *how to learn* (objective) complements optimizing *what to learn* (query), producing additional AFSL gains.

Cold-Start Advantages. The benefits of **SCL-HNM** are particularly pronounced at small K (e.g., MPQA-Polarity at $K = 20$: 84.1 vs. 81.5; Amazon Reviews at $K = 20$: 53.8 vs. 52.0), indicating higher accuracy with fewer annotations and lower initial labeling cost.

Generalization Across Task Types. **SCL-HNM** performs robustly across binary (MPQA-Polarity), multi-class (MPQA-Intensity, AG News, Amazon Reviews), and multi-label (MPQA-Type) tasks, validating its role as a general regularizer for structured representation learning.

Generalization Across Architectures and Query Strategies. In active learning, the choice of query strategy strongly influences the distribution and geometry of the support set (SETTLES, 2009). To ensure that the observed gains are not artifacts of a specific model backbone or a favorable interaction with a single acquisition function, we extended our evaluation to an alternative setup utilizing the `roberta-base` encoder and a standard Uncertainty-based (Entropy) sampling strategy. As detailed in Appendix A, **SCL-HNM** consistently outperforms standard CE fine-tuning under this distinct configuration. This confirms that the structural benefits of SCL-HNM maintain their robustness and generalizability across varying PLM architectures and distinct active learning query distributions.

Summary. These findings confirm that optimizing the objective is an *orthogonal yet complementary* axis to optimizing the query strategy in AFSL. With **SCL-HNM**, the framework achieves higher overall performance and greater sample efficiency without altering the query mechanism.

5.2. Ablation Study

To quantify the contribution of each component, we conduct ablation experiments on MPQA-Intensity and MPQA-Polarity with the **SOTA Query** fixed and $K = 100$, varying only the learning objective. Table 5 presents the results.

SCL as the Foundation (RQ1). Replacing CE with SCL consistently improves performance (e.g., +0.7 on MPQA-Intensity, +1.2 on MPQA-Polarity), demonstrating that enforcing intra-class compactness and inter-class separability enhances sample efficiency in data-scarce settings.

HNM as a Catalyst (RQ2). Adding HNM on top of SCL further enhances performance (e.g., MPQA-Intensity: 41.6 vs. 40.4; MPQA-Polarity: 92.5 vs. 92.0), showing that focusing on the most confusable negatives sharpens decision boundaries and provides additional gains.

Table 4: Micro-F1 (%) on five datasets under different query strategies and objectives. Results are mean $\pm$ std over five seeds. A dash (–) indicates a configuration is not applicable (the active query strategy is not used in full-data training).

Dataset	Method (Query + Objective)	$K = 10$	$K = 20$	$K = 50$	$K = 100$	FULL
MPQA-Type	Random + CE	56.2±4.6	60.0±2.9	65.3±2.6	66.7±2.7	80.7
	Random + SCL-HNM	54.8±4.7	59.3±1.9	62.5±2.6	**69.7±1.7**	**81.9**
	SOTA Query + CE	53.6±6.2	59.2±3.6	64.6±2.0	69.3±0.8	–
	Ours (SOTA Query + SCL-HNM)	**57.0±4.3**	**60.0±3.2**	**65.1±3.1**	69.4±0.5	–
MPQA-Polarity	Random + CE	76.5±2.4	80.6±2.4	85.3±1.4	88.4±0.9	94.2
	Random + SCL-HNM	**78.3±2.7**	82.0±1.7	84.1±2.0	90.3±0.8	**95.8**
	SOTA Query + CE	77.5±5.1	81.5±1.8	87.5±1.4	90.8±0.8	–
	Ours (SOTA Query + SCL-HNM)	77.8±2.1	**84.1±1.5**	**88.5±2.1**	**92.5±0.7**	–
MPQA-Intensity	Random + CE	33.0±3.8	34.0±3.6	35.5±1.6	35.5±1.5	50.0
	Random + SCL-HNM	33.5±2.0	33.6±3.7	35.7±2.9	40.2±0.9	**51.3**
	SOTA Query + CE	33.9±2.2	35.1±2.9	38.0±1.6	39.7±1.8	–
	Ours (SOTA Query + SCL-HNM)	**34.5±2.3**	**36.0±2.0**	**38.0±1.7**	**41.6±1.6**	–
AG News	Random + CE	71.8±6.3	87.3±1.2	88.4±1.2	89.3±0.7	94.4
	Random + SCL-HNM	73.5±5.1	87.0±1.5	88.9±0.9	89.8±0.6	**95.2**
	SOTA Query + CE	86.9±0.7	87.5±0.7	88.1±2.0	89.1±0.8	–
	Ours (SOTA Query + SCL-HNM)	**86.2±1.3**	**86.4±1.4**	**88.9±0.6**	**90.0±0.9**	–
Amazon Reviews	Random + CE	47.2±4.8	52.7±4.3	55.0±2.9	59.3±0.9	65.7
	Random + SCL-HNM	48.1±3.9	53.5±3.1	56.2±2.5	60.1±0.8	**68.8**
	SOTA Query + CE	51.5±0.2	52.0±3.5	58.3±1.6	59.9±1.1	–
	Ours (SOTA Query + SCL-HNM)	**52.4±0.8**	**53.8±2.3**	**59.1±1.3**	**60.5±0.9**	–

Table 5: Ablation results (Micro-F1, %) with $K = 100$ and a fixed SOTA query strategy. Mean $\pm$ std over five runs.

Dataset	Method (Objective)	Micro-F1
MPQA-Intensity	SOTA Query + CE	39.7±1.8
	SOTA Query + SCL	40.4±1.1
	SOTA Query + SCL-HNM	**41.6±1.6**
MPQA-Polarity	SOTA Query + CE	90.8±0.8
	SOTA Query + SCL	92.0±0.6
	SOTA Query + SCL-HNM	**92.5±0.7**

Ablation Summary. SCL establishes a structured embedding space superior to CE, while HNM amplifies its discriminative capacity by refining class margins. Together, they offer complementary benefits that account for the observed overall improvements.

5.3. Analysis of Embedding Space

While previous sections demonstrated the quantitative effectiveness of **SCL-HNM**, this section examines the underlying mechanism driving these improvements. We conduct both qualitative and quantitative analyses to explore how **SCL-HNM** enhances the geometric organization of the embedding space.

5.3.1. Visualization of Spatial Structure

To examine the influence of the learning objective, we employ t-SNE to project sentence embeddings from the Amazon Reviews test set ($K = 100$) into two dimensions. Figure 2 compares the representations generated by the baseline CE model and our proposed **SCL-HNM** method.

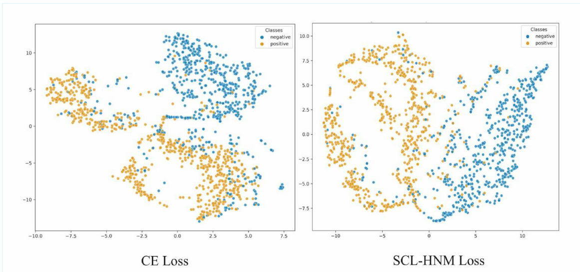

Figure 2: t-SNE visualization of sentence embeddings from the Amazon Reviews test set ($K = 100$): (a) SOTA baseline (CE loss); (b) Proposed **SCL-HNM** method.

SOTA baseline (CE loss). The CE-based embeddings exhibit an *entangled structure* with significant overlap between classes and diffuse intra-class distributions, resulting in poorly defined class boundaries and unstable downstream classification.

Proposed SCL-HNM. By contrast, the embeddings learned by **SCL-HNM** display high intra-class compactness and large inter-class margins. Samples form dense, coherent clusters, while distinct classes are well-separated. This structured, nearly linearly separable representation space offers a clear advantage for classification tasks.

Overall, these qualitative findings align with quan-

Table 6: Quantitative evaluation of embedding space quality on the Amazon Reviews test set ($K = 100$). Mean $\pm$ std over five runs. Lower Alignment indicates tighter intra-class clustering.

Method	**Alignment** $\downarrow$	**Uniformity**
SOTA Query + CE	0.6009	-1.5695
SOTA Query + SCL-HNM	**0.4287**	**-1.2812**

titative improvements, confirming that **SCL-HNM** strengthens discriminative representation learning through explicit geometric regularization.

5.3.2. Quantitative Analysis of Embedding Geometry

To complement the qualitative visualization, we quantitatively assess the structure of the embedding space using two standard metrics for contrastive representation learning (WANG and ISOLA, 2020):

- **Alignment** ($\mathcal{L}_{\text{align}}$): Measures the expected squared distance between embeddings of the same class (lower is better, indicating tighter intra-class clustering).
- **Uniformity** ($\mathcal{L}_{\text{uniform}}$): Measures how well the embeddings are distributed on the unit hypersphere (calculated as the logarithm of average pairwise exponential distances).

Table 6 presents the comparison between the CE baseline and **SCL-HNM**. The proposed objective significantly improves Alignment (reducing the score from 0.6009 to 0.4287), indicating highly cohesive intra-class structures. Notably, while unsupervised contrastive learning strictly minimizes Uniformity to scatter instances evenly, supervised contrastive learning (SCL) intentionally sacrifices some absolute uniformity to form distinct, widely separated class clusters on the hypersphere. This geometric dynamic explicitly explains the observed shift in the Uniformity metric (from -1.5695 to -1.2812).

These quantitative results strictly align with the t-SNE visualization, confirming that **SCL-HNM** reshapes the embedding geometry toward tighter class representations and stable feature organization, which directly translates to the superior classification accuracy observed in prior experiments.

6. Conclusion

This work investigates a longstanding yet underexplored bottleneck in **Active Few-Shot Learning (AFSL)**: the inefficient utilization of labeled samples induced by a suboptimal learning objective. Prior research has predominantly optimized *what to learn* (i.e., query strategies), while largely overlooking *how to learn* (i.e., the training objective). The widespread reliance on cross-entropy (**CE**)—which lacks explicit structural constraints on the embedding space—can limit further progress in AFSL.

To address this limitation, we propose **Supervised Contrastive Learning with Hard Negative Mining (SCL-HNM)** as a fine-tuning objective within the AFSL paradigm. Without altering any state-of-the-art (SOTA) query strategy, our approach integrates a structural regularizer (the **SCL-HNM** loss) with the conventional CE objective to improve sample utilization efficiency. Comprehensive experiments on five benchmarks show consistent and meaningful gains over strong baselines. Notably, in the challenging cold-start phase, our method yields up to a 2.6% absolute Micro-F1 improvement on **MPQA-Polarity** when $K = 20$, supporting both effectiveness and generalization.

Overall, our results provide systematic evidence that, in **AFSL**, *how to learn* is as critical as *what to learn*. Optimizing the training objective offers an **orthogonal and complementary** avenue to query design, enabling higher performance and better use of scarce annotations. We advocate treating the learning objective not as a fixed default but as an optimizable component for overcoming current performance plateaus.

Limitations

Although the proposed method demonstrates promising performance, several limitations remain and motivate future research.

(1) Generalization across diverse query strategies and architectures. As highlighted by recent active learning literature, the choice of query strategy strongly dictates the geometric distribution and structural density of the support set (SETTLES, 2009). In our primary experiments, we fixed the query strategy to Rep(En)-ClUn(En) to rigorously isolate the contribution of the learning objective. However, this raises a valid concern regarding whether the observed performance gains are inherently driven by SCL-HNM or merely a byproduct of a favorable interaction with this specific acquisition function. For instance, uncertainty-based sampling tends to select boundary instances closely clustered near decision margins (SETTLES, 2009), which might make hard negative mining (HNM) overly aggressive. Conversely, diversity-based methods select highly scattered instances (ZHAN et al., 2022), potentially diluting the intra-class clustering effect enforced by SCL. Recognizing these geometric dynamics is crucial for advancing AFSL. Future work must systematically evaluate SCL-HNM across alternative query strategies (e.g., pure uncertainty, diversity, and hybrid methods) to confirm its robustness. Furthermore, our empirical validation is currently limited to the encoder of `google/flan-t5-base`. Investigating how this hybrid objective interacts with the distinct embedding topologies of varied architectures, such as masked language models (e.g., RoBERTa) and decoder-only large language models (LLMs), remains a critical next step. **Preliminary experiments addressing this generalization on `roberta-base` with Uncertainty Sampling are provided in Appendix A.**

(2) Scalability and computational overhead. While **SCL-HNM** is simple to integrate, it introduces inherent computational overhead compared to standard CE. The contrastive term requires computing pairwise similarities within a mini-batch, resulting in an $\mathcal{O}(B^2)$ time complexity (where B is the batch size) compared to the $\mathcal{O}(B)$ complexity of CE. Furthermore, the hard negative mining mechanism involves sorting these similarities, which adds further non-linear scaling costs. For larger models or datasets, these overheads can become a scalability bottleneck. Promising future directions include utilizing parameter-efficient fine-tuning (e.g., LoRA) and gradient caching mechanisms to mitigate the memory and computational burden.

(3) Theoretical grounding. As pointed out in the review process, while our work provides strong empirical evidence for SCL-HNM, the theoretical explanation for why this hybrid objective consistently improves AFSL across diverse tasks and architectures remains limited. Future work should delve into the joint optimization dynamics—e.g., how the gradient interactions between CE and SCL-HNM affect generalization bounds and margin formation—aiming to establish a more principled theoretical foundation for contrastive objectives in active learning.

(4) Failure modes in extreme low-data regimes. A critical failure mode of SCL-HNM occurs in the extreme early stages of AFSL (e.g., $K < 5$). As the efficacy of supervised contrastive learning inherently relies on the availability of positive pairs (KHOSLA et al., 2021), a severe scarcity of labeled data may result in mini-batches where certain classes lack positive pairs entirely. In such scenarios, the contrastive objective cannot be properly computed, potentially destabilizing the training dynamics or causing the model to underperform compared to standard CE. Future iterations of this framework could incorporate momentum encoders or memory banks (HE et al., 2020) to cache historical representations, thereby alleviating the scarcity of positive pairs during the initial active querying rounds.

Data and Code Availability

To facilitate reproducibility and support future research, the source code implementing the proposed SCL-HNM objective, along with the specific alternative data splits used in our experiments, are publicly available at: `https://github.com/zzz-Sgr/AFSL-SCL-HNM`.

7. Bibliographical References

S. AHMADNIA, A. Y. JORDEHI, M. H. K. HEYRAN, et al. 2025. Active few-shot learning for text classification. ArXiv.

S. AHMADNIA, A. YOUSEFI JORDEHI, M. HOSSEINI KHASHEH HEYRAN, et al. 2024. Opinion mining using pre-trained large language models: Identifying the type, polarity, intensity, expression, and source of private states. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 12481–12495, Torino, Italia. ELRA and ICCL.

S. BAIK, J. CHOI, H. KIM, et al. 2021. Meta-learning with task-adaptive loss function for few-shot learning. ArXiv.

T. B. BROWN, B. MANN, N. RYDER, et al. 2020. Language models are few-shot learners. ArXiv.

E. CHANG, X. SHEN, H. S. YEH, et al. 2021. On training instance selection for few-shot neural text generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 8–13, Online. Association for Computational Linguistics.

Y. CHEN, Y. ZHANG, C. ZHANG, et al. 2021. Revisiting self-training for few-shot learning of language model. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9125–9135, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

A. EDWARDS and J. CAMACHO-COLLADOS. 2024. Language models for text classification: Is in-context learning enough? ArXiv.

L. EIN-DOR, A. HALFON, A. GERA, et al. 2020. Active learning for bert: An empirical study. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7949–7962, Online. Association for Computational Linguistics.

T. GAO, A. FISCH, and D. CHEN. 2021a. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.

T. GAO, X. YAO, and D. CHEN. 2021b. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

B. GUNEL, J. DU, A. CONNEAU, et al. 2021. Supervised contrastive learning for pre-trained language model fine-tuning. ArXiv.

K. HE, H. FAN, Y. WU, et al. 2020. Momentum contrast for unsupervised visual representation learning. ArXiv.

R. KARIMI MAHABADI, L. ZETTLEMOYER, J. HENDERSON, et al. 2022. Prompt-free and efficient few-shot learning with language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3638–3652, Dublin, Ireland. Association for Computational Linguistics.

P. KHOSLA, P. TETERWAK, C. WANG, et al. 2021. Supervised contrastive learning. ArXiv.

A. LI, W. HUANG, X. LAN, et al. 2020. Boosting few-shot learning with adaptive margin loss. ArXiv.

X. V. LIN, T. MIHAYLOV, M. ARTETXE, et al. 2022. Few-shot learning with multilingual language models. ArXiv.

Lajanugen LOGESWARAN and Honglak LEE. 2018. An efficient framework for learning sentence representations. In *arXiv preprint arXiv:1803.02893*. Accessed: 2025-09-17.

Z. LONG, G. KILLICK, R. MCCREADIE, et al. 2023. When hard negative sampling meets supervised contrastive learning. ArXiv.

T. MÜLLER, G. PÉREZ-TORRÓ, A. BASILE, et al. 2022. Active few-shot learning with fasl. ArXiv.

B. SETTLES. 2009. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin-Madison.

T. WANG and P. ISOLA. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 9929–9939. PMLR.

Y. WANG, Q. YAO, J. KWOK, et al. 2020. Generalizing from a few examples: A survey on few-shot learning. ArXiv.

T. YIN, N. LIU, and H. SUN. 2021. Self-paced active learning for deep cnns via effective loss function. *Neurocomputing*, 424:1–8.

M. YUAN, H. T. LIN, and J. BOYD-GRABER. 2020. Cold-start active learning through self-supervised language modeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7935–7948, Online. Association for Computational Linguistics.

X. ZHAN, Q. WANG, K. HUANG, et al. 2022. A comparative survey of deep active learning. ArXiv.

8. Language Resource References

2005. *AG's corpus of news articles*.

2020. *The Multilingual Amazon Reviews Corpus*.

WIEBE, J. and WILSON, T. and CARDIE, C. and others. 2005. *Annotating Expressions of Opinions and Emotions in Language (MPQA Corpus)*.

A. Generalization Across Architectures and Query Strategies

To address potential concerns regarding the generalizability of SCL-HNM across different model architectures and active learning acquisition functions, we conducted preliminary experiments using an alternative setup.

We replaced the `flan-t5-base` encoder with `roberta-base` and swapped the Rep(En)-ClUn(En) query strategy for a standard Uncertainty-based sampling (Entropy). Table 7 presents the Micro-F1 scores on the MPQA-Polarity dataset under this new configuration.

Table 7: Micro-F1 (%) on MPQA-Polarity using `roberta-base` and Entropy sampling. Results are mean ± std over 5 random seeds.

Method	$K = 20$	$K = 50$	$K = 100$
Uncertainty + CE	60.0 ± 5.6	74.9 ± 4.1	78.7 ± 6.1
Uncertainty + SCL-HNM	**69.4** ± **3.3**	**76.6** ± **3.7**	**79.1** ± **4.4**

The results indicate that SCL-HNM continues to outperform standard CE fine-tuning, demonstrating that the structural benefits of the proposed hybrid loss are not strictly bound to a single model backbone or a specific query strategy. Furthermore, as indicated by the reduced standard deviation, SCL-HNM significantly improves the training stability across different active learning initialization seeds.