

# Building Domain-Specific Small Language Models via Guided Data Generation

Aman Kumar<sup>1</sup>, Ekant Muljibhai Amin<sup>2</sup>, Xian Yeow Lee<sup>1</sup>, Lasitha Vidyaratne<sup>1</sup>, Ahmed K. Farahat<sup>1</sup>, Dipanjan D. Ghosh<sup>1</sup>, Yuta Koreeda<sup>2</sup>, Chetan Gupta<sup>1</sup>

<sup>1</sup>Hitachi America Ltd., Santa Clara, CA, USA

<sup>2</sup>Hitachi Ltd., Tokyo, Japan

## Abstract

Large Language Models (LLMs) have shown remarkable success in supporting a wide range of knowledge-intensive tasks. In specialized domains, there is growing interest in leveraging LLMs to assist subject matter experts with domain-specific challenges. However, deploying LLMs as SaaS solutions raises data privacy concerns, while many open-source models demand significant computational resources for effective domain adaptation and deployment. A promising alternative is to develop smaller, domain-specialized LLMs, though this approach is often constrained by the lack of high-quality domain-specific training data. In this work, we address these limitations by presenting a cost-efficient and scalable training pipeline that combines guided synthetic data generation from a small seed corpus with bottom-up domain data curation. Our pipeline integrates Domain-Adaptive Pretraining (DAPT), Domain-specific Supervised Fine-tuning (DSFT), and Direct Preference Optimization (DPO) to train effective small-scale models for specialized use cases. We demonstrate this approach through *DiagnosticSLM*, a 3B-parameter domain-specific model tailored for fault diagnosis, root cause analysis, and repair recommendation in industrial settings. To evaluate model performance, we introduce four domain-specific benchmarks: multiple-choice questions (DiagnosticMCQ), question answering (DiagnosticQA), sentence completion (DiagnosticComp), and summarization (DiagnosticSum). *DiagnosticSLM* achieves up to 25% accuracy improvement over open-source models of comparable or larger size (2B-9B) on the MCQ task, while also outperforming or matching them in other tasks, demonstrating effective domain-specific reasoning and generalization capabilities.

## Introduction

Large Language Models (LLMs) have demonstrated remarkable potential in supporting a wide range of knowledge-intensive tasks across domains such as law (Lai et al. 2024), finance (Li et al. 2023), and medicine (Thirunavukarasu et al. 2023). This success has prompted increasing interest in applying LLMs to specialized industrial domains, including fault diagnosis and repair. These domains are often hindered by fragmented documentation, shortage of skilled technicians, and the high cost and time investment required to train new personnel (Gao, Cecati, and Ding 2015; Nonaka

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

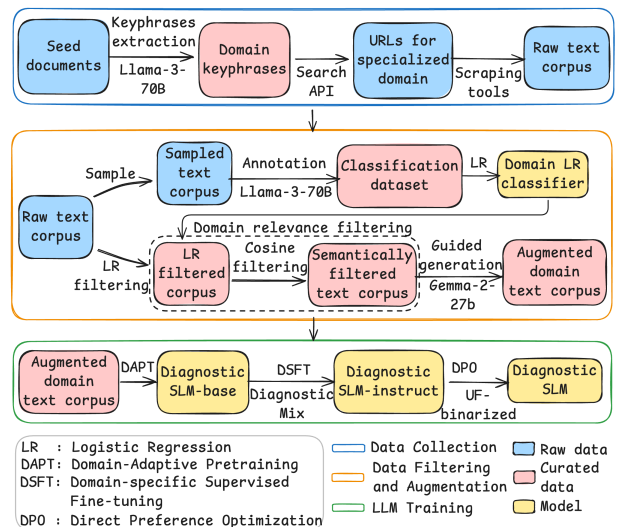


Figure 1: Overview of the *DiagnosticSLM* pipeline, illustrating domain-specific data collection, guided synthetic data generation, and three-stage model training.

and Takeuchi 1996). While technicians performing low-skill operations may rely on manuals and standard procedures, high-skill tasks (such as fault identification, root cause analysis, and repair planning) require deep experiential knowledge that is rarely formalized and is often tacit (Argote and Ingram 2000).

Despite the promise of LLMs, there are several barriers to their adoption in industrial settings. Proprietary models like GPT-4o (Hurst et al. 2024), Gemini (Team et al. 2023), and Claude (Bai et al. 2022) pose privacy risks due to cloud-only inference and lack of transparency, making them unsuitable for regulated environments. Open-source LLMs, though more accessible, often demand substantial computational resources for fine-tuning and deployment. Additionally, general-purpose models typically lack the domain-specific vocabulary, reasoning capability, and contextual grounding needed to support technicians in high-stakes diagnostic workflows (Khodabandelou et al. 2024).

A promising alternative is to develop small, domain-specialized LLMs that can be deployed on-premise and cus-

tomized for specific industrial use cases. Prior work on domain-adapted models such as LawGPT (Nguyen 2023), BloombergGPT (Wu et al. 2023), and ChipNeMo (Liu et al. 2023) has shown that targeted pretraining on specialized corpora significantly improves performance in expert domains. In particular, Domain-Adaptive Pretraining (DAPT) (Gururangan et al. 2020) followed by Supervised Fine-tuning (SFT) (Akatsuka et al. 2024) has emerged as a cost-effective strategy for domain adaptation.

However, a key bottleneck in this process is access to high-quality domain-specific training data. Industrial data is often siloed, proprietary, and inconsistently documented, limiting the ability to build or evaluate customized models. Moreover, curating clean, relevant subsets from massive corpora like Common Crawl is labor-intensive and prone to low recall (Remus and Biemann 2016). Even with access to data, training from scratch or full fine-tuning of large models is often cost-prohibitive due to the extensive GPU hours and memory requirements involved (Xia et al. 2024).

To address these challenges, we propose a scalable pipeline for building domain-specialized small language models (SLMs). Our approach starts with bottom-up domain data curation and synthetic data generation from a curated seed corpus, followed by a three-stage training process to produce effective, lightweight models suitable for real-world deployment.

Our contributions are as follows: (i) a pragmatic, low-resource pipeline that combines Domain-Adaptive Pretraining (DAPT), Domain-specific Supervised Fine-tuning (DSFT), and Direct Preference Optimization (DPO), tailored for small on-premise LLMs deployed in regulated industrial domains such as automotive diagnostics; (ii) a novel guided synthetic data augmentation strategy that leverages LLMs to enrich a bottom-up curated seed corpus for domain-specific data construction; (iii) a targeted approach for generating supervised fine-tuning data aligned with key domain tasks; and (iv) the introduction of four new domain-specific benchmarks, multiple-choice questions (DiagnosticMCQ), question answering (DiagnosticQA), sentence completion (DiagnosticComp), and summarization (DiagnosticSum), for evaluating LLMs in automotive diagnostics. An overview of our method is shown in Figure 1.

## Methodology

We adopted a three-step training pipeline to develop our proposed domain-specific SLM, *DiagnosticSLM*. First, we perform additional pretraining of an open-source SLM, Llama-3.2-3B (Grattafiori et al. 2024; Meta 2024), on a text corpus of Automotive domain. Next, we perform supervised fine-tuning using automotive domain related tasks adapted from general supervised fine-tuning Alpaca dataset (Taori et al. 2023). Lastly, we perform direct preference optimization using the UltraFeedback dataset (Cui et al. 2023). For all experiments, we utilized 2× NVIDIA RTX 4090 GPUs.

### Domain Adaptive Pretraining (DAPT)

**Automotive Domain Data Curation** We selected the automotive domain as the primary focus for diagnostics, with

an emphasis on diagnostic procedures, repair operations, and core automotive concepts. We adopted a bottom-up data collection strategy, extracting keyphrases from internal documents and technical manuals using an internally hosted Llama-3-70B-Instruct model (Grattafiori et al. 2024) in a 4-bit configuration (Mangrulkar et al. 2022). Outputs were verified by an expert for correctness and then used to guide targeted web searches via the Google Custom Search API. This process yielded approximately 403 million tokens from 706,971 webpages. Representative keyphrases are provided in the Appendix (Kumar et al. 2025).

**Domain Relevance Filtering and Classification** To further refine the dataset, we performed domain relevance filtering using a combination of LLM annotation and classical classification techniques. A random sample of 20,000 webpages was selected for annotation of the classification dataset. We used Llama-3-70B-Instruct model to label each instance as relevant or irrelevant to the automotive domain. The LLM annotated 11,621 samples as relevant and 8,379 as irrelevant, with the full annotation process taking approximately 22 GPU-hours. This highlights the considerable time and computational cost of large-scale LLM-based labeling, making it too expensive to annotate the entire dataset by experts. A manual review of a subset was conducted to validate label quality. While no corrections were applied, the low observed error rate and minimal review time (2 person-hours) supported continued use of the LLM-labeled data.

This labeled subset was split 80–20 into training and test sets. We trained a logistic regression classifier with L2 regularization ( $C = 10$ ) using the SAGA solver. The resulting model achieved an accuracy of 88.2% on the test set. Applying this classifier to the full dataset yielded 356,312 instances (50.39%) classified as *automotive-related*, and 350,624 (49.61%) as *non-automotive-related*.

Using eight topic-specific documents representing major automotive systems, we computed cosine similarity between each topic and all automotive-related samples, retaining those with a similarity score greater than 0.25. We also recovered possible false negatives by applying the same similarity check to the non-relevant set and retaining the top 20%. After merging and deduplication, the curated dataset contained 387,572 samples (257M tokens). Details of the topic creation process and similarity computation are in Appendix (Kumar et al. 2025).

**Guided Synthetic Data Augmentation** To enhance the dataset, we employed a teacher model to expand the automotive content which is prompted to generate additional domain-relevant text and remove non-relevant portions. This augmentation step aimed to create a more comprehensive and detailed corpus, thereby improving the model’s ability to understand and generate automotive-specific information. A relatively smaller model, Gemma-2-27B (Team et al. 2024), was selected to facilitate efficient parallel inference across two GPUs. For data points exceeding the single-GPU context window, we partitioned the text into smaller chunks, processed them independently, and merged the outputs.

The augmentation process focused on enriching all relevant samples by prompting the teacher model to add factu-

ally accurate details and more comprehensive explanations, guided by its internal automotive knowledge. The underlying assumption is that the teacher model already possesses a certain amount of the necessary domain knowledge; rather than generating information from scratch, the prompts encourage the model to elaborate, clarify, and expand upon existing content. In this way, the teacher model plays a constructive role in enhancing the dataset by injecting technical depth and contextual breadth in a guided manner. This prompt-based strategy not only deepened the technical content and improved specificity across the dataset but also facilitated the filtering of non-relevant samples and removal of low-value, non-technical content inside the sample such as addresses, geographic references, and irrelevant forum conversations. The entire expansion and modification pipeline consumed approximately 5,400 GPU-hours. Following augmentation, we applied fuzzy de-duplication using MinHash (Shrivastava and Li 2014) to remove redundant sentences including content from our automotive benchmark. The final automotive corpus comprised 206 million tokens.

**Model and framework** We leveraged a 3B-parameter Llama-3.2 model for DAPT using our curated automotive corpus. The model was initialized from its publicly released pretrained weights. We employed the LlamaFactory Framework (Zheng et al. 2024b) for training and utilized fully sharded data parallelism (FSDP) (Zhao et al. 2023) via the Accelerate library (Gugger et al. 2022) to distribute model parameters across GPUs.

**DAPT Training** We adopted a full-parameter fine-tuning strategy on the automotive domain-specific data. DAPT was conducted using the causal language modeling (CLM) objective, consistent with the pretraining objective of Llama-3 models. CLM can be defined as the negative log-likelihood of the next token given all previous tokens:

$$\mathcal{L}_{\text{CLM}} = - \sum_{t=1}^T \log P_{\theta}(x_t | x_{<t}) \quad (1)$$

where  $x_t$  is the token at position  $t$ ,  $x_{<t}$  are all preceding tokens, and  $P_{\theta}$  is the probability distribution parameterized by the model weights  $\theta$ .

Full-parameter fine-tuning was chosen over parameter-efficient methods (e.g., LoRA) to maximize domain alignment and capture deeper semantic shifts specific to the automotive domain, and was computationally feasible given the model’s relatively small size (3B parameters). The model was trained using the AdamW optimizer with a learning rate of  $1 \times 10^{-4}$ , and a cosine learning rate scheduler with 10% warmup steps. We set the per-device batch size to 1 and used gradient accumulation over 8 steps, resulting in an effective global batch size of 16 across two GPUs. The input sequence length was set to 2,048 tokens, yielding a total of 16,384 tokens processed per forward pass. Training was conducted for 1 epoch, consisting of 5,789 steps in FP16 precision. The entire training process took approximately 118 GPU-hours. We call this model *DiagnosticSLM-base*.

## Domain-specific Supervised Fine-tuning (DSFT)

**DSFT Data Generation** To adapt the DAPT model to task-specific objectives, we constructed a domain-specific instruction-tuning dataset for the automotive domain. We curated 10 distinct task types and provided three example prompts per task to the GPT-4o model. These prompts guided the generation of instruction–response pairs across eight automotive topics introduced earlier in the data collection phase (e.g., Engine Repair, Brakes, etc.).

Using this setup, we generated approx. 20,000 examples spanning 10 automotive-domain instruction tasks inspired by Alpaca dataset tailored for the automotive domain. Details of each task are provided in the Appendix (Kumar et al. 2025). To balance domain-specific and general instruction-following capabilities, we combined our domain task dataset with the 52,000 samples from the Alpaca. This produced a mixed training set for ablation studies evaluating the effects of different dataset combinations. The final DSFT corpus comprises approximately 72,000 instruction–response pairs, and we refer to it as *DiagnosticMix*.

**DSFT Training** We fine-tuned the DAPT model on the DSFT dataset using an auto-regressive next-token prediction objective.

$$\mathcal{L}_{\text{SFT}} = - \sum_{t=1}^T \log P_{\theta}(y_t | y_{<t}, x) \quad (2)$$

where  $x$  is the instruction prompt and  $y = (y_1, \dots, y_T)$  is the ground-truth response. The model is trained to generate  $y$  conditioned on  $x$ .

Training was performed using the LlamaFactory framework with nearly identical hyperparameter settings as DAPT, except for a reduced learning rate of  $1 \times 10^{-5}$ . The per-device batch size was set to 2, with gradient accumulation over 8 steps, and an effective global batch size of 32.

The model was trained for 1 epoch, totaling 2,130 training steps, using precision FP16. The complete training process took approximately 47 GPU-hours. We call this model *DiagnosticSLM-instruct*.

## Direct Preference Optimization (DPO)

While DAPT and DSFT enable the *DiagnosticSLM-instruct* to internalize domain knowledge, they do not guarantee that the model’s outputs align with human preferences. To address this, we performed additional fine-tuning through preference alignment using the DPO (Rafailov et al. 2023). DPO is an alternative to Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al. 2022) that directly optimizes a language model to prefer responses that align with human-annotated preferences, without requiring a separate reward model or reinforcement learning loop. Given a set of preference pairs  $(x, y_{\text{pos}}, y_{\text{neg}})$ , where  $x$  is the input and  $y_{\text{pos}}$  and  $y_{\text{neg}}$  are the preferred and less preferred responses respectively, DPO fine-tunes the model by minimizing the following loss:

$$\mathcal{L}(\theta) = - \log \sigma(\beta(\log \pi_{\theta}(y_{\text{pos}} | x) - \log \pi_{\theta}(y_{\text{neg}} | x))) \quad (3)$$

Model	Params	Training Type	Accuracy (Correct/Total)	A	B	C	D
<i>Ground Truth (label dist.)</i>	–	–	–	240	246	205	185
<b>DiagnosticSLM (ours)</b>	<b>3B</b>	<b>DAPT+DSFT+DPO</b>	<b>0.4532 (397/876)</b>	<b>288</b>	<b>311</b>	<b>201</b>	<b>76</b>
Phi-3.5-mini-instruct	3.8B	Original Instruct	0.4384 (384/876)	105	203	372	196
Phi-4-mini-instruct	3.8B	Original Instruct	0.4098 (359/876)	91	200	364	221
Gemma-2-2B-it	2B	Original Instruct	0.3733 (327/876)	265	83	317	211
Llama-3.2-3B-Instruct	3B	Original Instruct	0.3653 (320/876)	166	130	426	154
Qwen2.5-3B-Instruct	3B	Original Instruct	0.3630 (318/876)	115	154	352	223
Gemma-2-2B	2B	Base	0.3368 (295/876)	173	406	166	131
Qwen2.5-3B	3B	Base	0.3288 (288/876)	112	122	414	161
Llama-3.2-3B	3B	Base	0.2705 (237/876)	155	63	621	37
DiagnosticSLM-base (ours)	3B	DAPT	0.2352 (206/876)	4	3	868	1
<b>Llama-3.1-8B-Instruct</b>	<b>8B</b>	<b>Original Instruct</b>	<b>0.4692 (411/876)</b>	<b>136</b>	<b>220</b>	<b>305</b>	<b>215</b>
Gemma-2-9b-It	9B	Original Instruct	0.4521 (396/876)	147	152	338	162
Ministral-8B-Instruct-2410	8B	Original Instruct	0.4281 (375/876)	153	118	385	216
Llama-3.1-Tulu-3-8B	8B	Original Instruct	0.4041 (354/876)	140	241	129	366
c4ai-command-r7b-12-2024	7B	Original Instruct	0.3984 (349/876)	216	241	331	87
Llama-3.1-8B	8B	Base	0.3095 (247/876)	171	110	402	46

Table 1: Comparison of model variants on the *DiagnosticMCQ* evaluation using 5-shot prompting. Columns A, B, C, and D indicate how many times each model selected each answer option across the full set of questions.

where  $\pi_\theta$  is the model’s output distribution,  $\sigma(\cdot)$  is the sigmoid function, and  $\beta$  is a temperature parameter controlling the sharpness of preference. This encourages the model to increase the log-probability gap between preferred and non-preferred responses. Compared to RLHF, DPO offers a simpler and more stable training paradigm, while maintaining strong empirical performance. We adopt DPO in our alignment stage due to its computational efficiency, reduced implementation complexity, and its ability to effectively guide the model toward preferred behaviors.

**Dataset** We used the UltraFeedback Binarized dataset (Cui et al. 2023) as a general-purpose source of pairwise preferences for DPO and apply it to further fine-tune *DiagnosticSLM-instruct*. A brief dataset summary appears in Appendix (Kumar et al. 2025). Because UltraFeedback is not domain-specific, we plan to construct an automotive preference dataset and repeat DPO with that domain data.

**DPO Fine-tuning** We fine-tune the Llama-3.2-3B-based *DiagnosticSLM-instruct* model using DPO with LoRA, selected due to computational constraints from the additional memory overhead of preference-pair inputs and gradient storage. Training is conducted for 1 epoch with a per-device batch size of 1, a learning rate of  $1 \times 10^{-5}$ , and a cosine learning rate scheduler with a 10% warm-up ratio. We use bfloat16 precision training and apply a sigmoid-based preference loss with  $\beta = 0.1$ . For LoRA, we set the rank to 16 and the  $\alpha$  to 4. The input sequence length is capped at 2,048 tokens, and data preprocessing is parallelized across 8 workers. Evaluation is performed every 500 steps, with 10% of the dataset reserved for validation. The entire DPO training process took approximately 144 GPU-hours. We refer to this final model as *DiagnosticSLM*.

## Evaluation

The Automotive Service Excellence (ASE) certification exams are widely regarded as the standard for assessing technical proficiency in automotive diagnostics, repair, and maintenance across vehicle systems (Thompson 2014). To evaluate domain-specific knowledge coverage and enable quantitative performance assessment, we constructed four benchmarks: multiple-choice questions (*DiagnosticMCQ*), question answering (*DiagnosticQA*), sentence completion (*DiagnosticComp*), and summarization (*DiagnosticSum*), inspired by the ASE exam format and subject areas. *DiagnosticMCQ* was built from commercially licensed ASE-aligned practice materials, while the other benchmarks were derived from it through reformatting and GPT-4o-generated variations, followed by subject matter expert validation to ensure technical accuracy. We applied strict de-duplication against all training corpora to prevent overlap and contamination. For all experiments, we compare our model with Llama-3.2 (Meta 2024), Qwen2.5 (Yang et al. 2024), Phi-3.5/Phi-4 (Abdin et al. 2024), and Gemma-2 (Team et al. 2024).

### Multiple Choice Question Task

*DiagnosticMCQ* contains 876 multiple-choice questions across various automotive subtopics, each with four answer choices and one correct ground truth. Each evaluation prompt includes five example QA pairs followed by a test question to guide the model in understanding the structure and answer format. Model predictions are compared against the ground truth to compute accuracy. As shown in Table 1, our final model, *DiagnosticSLM*, achieves an accuracy of 45.32% (397/876), surpassing the Llama-3.2-3B-Instruct baseline (36.53%) by a substantial margin. Despite its smaller size, *DiagnosticSLM* also outperforms several larger

models, including Ministral (8B, 42.81%), Llama3.1-Tulu (8B, 40.41%), and C4ai-command-r (7B, 39.84%) and performs on par with Gemma-2 (9B, 45.21%). It also exceeds the performance of similarly sized or slightly larger models such as Phi-4-mini-instruct (3.8B, 40.98%), Qwen2.5-3B-Instruct (3B, 36.30%), and Gemma-2 (2B, 37.33%).

We observed that nearly all base versions of models exhibit skewed option selection when evaluated on *DiagnosticMCQ*, often favoring a specific choice regardless of the question content. This trend shifts noticeably in their instruction-tuned variants. We observed a similar trend in our own model: after DAPT, the model predominantly selected option C. However, following DSFT and DPO, the answer distribution becomes more balanced. For comparison, the distribution of correct options in the ground truth is also included in the Table 1.

### Question-Answering Task

LLMs often exhibit selection biases, i.e., token bias towards label tokens (A/B/C/D) (Wang et al. 2023; Pezeshkpour and Hruschka 2023), positional bias in answer ordering (Zheng et al. 2024a), and context priming from in-context examples (Zhao et al. 2021). To introduce variation in evaluation and better assess the instruction-following capabilities of the models, we converted the *DiagnosticMCQ* dataset into a question-answering (QA) format using GPT-4o, referred to as *DiagnosticQA* benchmark. In this setting, the answer choices are embedded directly within the question in natural language, which reduces surface-level biases by requiring models to interpret the full input and generate an explicit answer rather than selecting a label. An example question is shown in Appendix Figure 4 (Kumar et al. 2025). We evaluate models using accuracy, comparing their generated answers against the ground truth.

Model	Params	Accuracy
<b>DiagnosticSLM (ours)</b>	<b>3B</b>	<b>0.3831</b>
Phi-4-mini-instruct	3.8B	0.3628
Qwen2.5-3B-Instruct	3B	0.3299
Gemma-2-2B-it	2B	0.3169
Phi-3.5-mini-instruct	3.8B	0.2342
Llama-3.2-3B-Instruct	3B	0.2208

Table 2: Accuracy comparison of different models evaluated on the *DiagnosticQA* dataset using 5-shot prompting.

To assess performance in this free-form QA setting, we evaluated multiple models with similar parameter sizes. As shown in Table 2, all models experienced a drop in accuracy when shifting from MCQ to QA format. For example, *DiagnosticSLM* decreased from 45.32% to 38.31%, *Phi-4* from 40.98% to 36.28%, *Qwen2.5* from 36.30% to 32.99%, *Gemma-2* from 37.33% to 31.69%, *Phi-3.5* from 43.84% to 23.42%, and *Llama-3.2* from 36.53% to 22.08%. Despite the overall performance degradation observed in the QA setting, *DiagnosticSLM* consistently outperformed all other models by a substantial margin.

### Sentence Completion Task

To further assess domain-specific knowledge with reduced bias, we introduce the *DiagnosticComp* benchmark. Each question in *DiagnosticMCQ* is reformulated into four natural language sentences, each as an independent completion prompt, by appending one answer option to the shared question stem. The completions are generated using GPT-4o. Appendix Figure 4 (Kumar et al. 2025) shows a sample prefix and its four completions.

Model	Params	Accuracy
<b>DiagnosticSLM (ours)</b>	<b>3B</b>	<b>0.5556</b>
Phi-4-mini-instruct	3.8B	0.5128
Phi-3.5-mini-instruct	3.8B	0.5076
Llama-3.2-3B-Instruct	3B	0.4786
Gemma-2-2B-it	2B	0.4701
Qwen2.5-3B-Instruct	3B	0.4188

Table 3: Accuracy comparison of different models evaluated on the *DiagnosticComp* dataset using log-likelihood scoring.

For each sentence prompt, we extract the model’s output logits at every generation step, apply a softmax over the vocabulary dimension to obtain token-level probabilities  $P(x_i | x_{<i})$ , and compute the *aggregate log-likelihood*:

$$\log P(x_{1:n} | \text{prefix}) = \sum_{i=1}^n \log P(x_i | x_{<i}). \quad (4)$$

Since all four prompts share the same prefix, differences in total log-likelihoods reflect only the option texts. The option with the highest probability score is selected.

To focus on surface-level knowledge, we filter out questions requiring multi-step or implicit reasoning based on: (1) patterns such as reasoning chains or negations, and (2) option length. Longer options tend to accumulate lower log-likelihoods, so we restrict options to a maximum of four words and remove MCQs with longer options. After filtering, the final *DiagnosticComp* set contains 117 completion questions. As shown in Table 3, *DiagnosticSLM* outperforms other models, achieving 55.56% accuracy.

### Summarization Task

To evaluate models’ ability to summarize domain-specific technical content, we introduce the *DiagnosticSum* task. Each input consists of a 3–6 line ground-truth explanation from the *DiagnosticMCQ* dataset. Using GPT-4o, we generate concise two-line summaries of these explanations as reference outputs. Models are prompted to produce two-line summaries of the input, aiming to preserve key technical information in compact form.

We evaluate summarization using ROUGE (1, 2, L), BLEU, BERTScore (F1), and cosine similarity, utilizing the rouge-score, NLTK, bert-score, and sentence-transformers (all-MiniLM-L6-v2 (Wang et al. 2020)) libraries. To account for generative variability, we report results over five independent trials at a temperature of 0.5. Aggregated results are

reported in Appendix Table 5 (Kumar et al. 2025). Phi-4-mini-instruct achieves the highest performance, while *DiagnosticSLM* remains competitive with models of similar size. *DiagnosticSLM*'s lower performance likely stems from a lack of domain-specific summarization data.

### Dataset Ablation Study

Model Name	Training Type	Accuracy
Llama-3.2-3B-Instruct	Original Instruct	36.53
Llama-3.2-3B	Base + SFT	30.71
Llama-3.2-3B	Base + DAPT + SFT	37.79
Llama-3.2-3B	Base + DSFT	38.24
DiagnosticSLM-instruct (ours)	Base + DAPT + DSFT	44.41
<b>DiagnosticSLM (ours)</b>	<b>Base + DAPT + DSFT + DPO</b>	<b>45.32</b>

Table 4: Ablation study on different datasets evaluated on *DiagnosticMCQ* using 5-shot prompting. *Note:* SFT refers to fine-tuning on the Alpaca dataset; DSFT refers to fine-tuning on the *DiagnosticMix* dataset.

To evaluate the individual and combined effects of DAPT, DSFT, and DPO, we conducted a series of ablation studies using the Llama-3.2-3B model. Our results, in Table 4, demonstrate that neither DAPT nor DSFT alone is sufficient to achieve optimal performance. Specifically, applying DAPT followed by SFT using Alpaca dataset yields a modest improvement in *DiagnosticMCQ* accuracy (37.79) compared to the original instruction-tuned model (36.53), indicating that DAPT introduces some domain-awareness. Similarly, fine-tuning the base model directly with *DiagnosticMix* data without prior DAPT results in a comparable *DiagnosticMCQ* accuracy of 38.24, suggesting that task-specific supervision alone is also moderately effective. However, combining DAPT with *DiagnosticMix* DSFT leads to a substantial performance boost, achieving a *DiagnosticMCQ* accuracy of 44.41, demonstrating that domain pretraining and domain-specific supervision work synergistically. Further incorporating DPO pushes the *DiagnosticMCQ* accuracy even higher to 45.32, underscoring the added value of preference alignment in improving model behavior.

### Pathway to Deployment

DiagnosticSLM is being prepared for on-premise deployment, and is designed to run inside customer or BU-managed infrastructure where data privacy, network isolation, and performance requirements must be met. Figure 2 illustrates the end-to-end pipeline used to develop and operate the system. Data is collected from internal knowledge bases and selected external sources using the Data Collection API, then filtered and normalized through the Data Filtering and Augmentation APIs to create task-specific datasets. The SLM Training API and Model Training Platform manage model training, validation, and version-controlled registra-

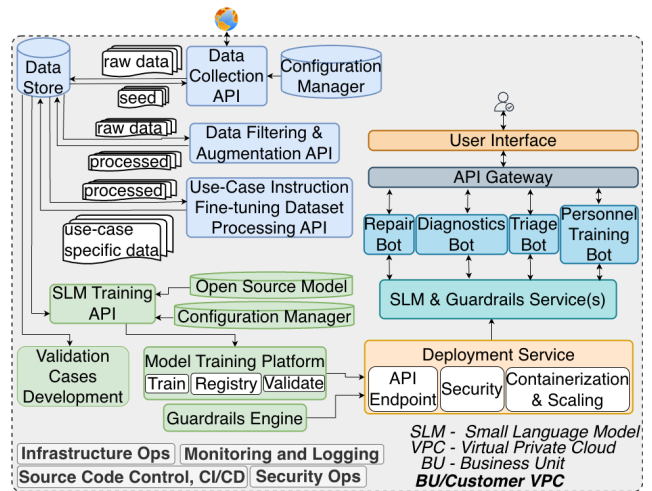


Figure 2: *DiagnosticSLM* deployment architecture

tion, with a Guardrails Engine enforcing domain constraints, including abstention when model confidence is low.

For deployment, models are packaged as containerized inference services and hosted in customer or Business Unit (BU) virtual private cloud (VPC) environments using Kubernetes. The Deployment Service provisions API endpoints, authentication, authorization, and resource allocation. Inference primarily targets GPU nodes, with quantized CPU configurations available for lower-throughput environments. The deployed SLM and Guardrails services expose standard APIs consumed by downstream applications such as Repair, Diagnostics, Triage, and Personnel Training assistants through a unified API Gateway and User Interface. Monitoring of latency, abstention behavior, drift indicators, and error logs supports routine operation, while CI/CD workflows enable controlled updates and rollback. Post-deployment evaluation will examine how end-users interact with the system during routine maintenance and troubleshooting workflows, which may help to improve task efficiency.

### Conclusion

We introduced *DiagnosticSLM*, a small language model tailored for diagnostics and repair to assist frontline workers, and it is lightweight for edge deployment in industrial environments. Our three-stage training pipeline consisting of DAPT, DSFT, and DPO, enabled effective domain adaptation using curated and augmented data. Across tasks such as multiple-choice questions, question-answering, sentence completion and summarization, *DiagnosticSLM* outperformed or matched larger open-source models, underscoring the impact of specialized training on domain-specific corpora. Our ablation study highlights the complementary contributions of each training stage. While this work focused on validating the end-to-end pipeline, future work will incorporate retrieval-augmented inference and parameter-efficient fine-tuning techniques to further enhance performance and adaptability.

## References

- Abdin, M.; Aneja, J.; Behl, H.; Bubeck, S.; Eldan, R.; Gunasekar, S.; Harrison, M.; Hewett, R. J.; Javaheripi, M.; Kauffmann, P.; et al. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.
- Akatsuka, S.; Kumar, A.; Lee, X. Y.; Vidyaratne, L.; Ghosh, D. D.; and Farahat, A. K. 2024. Rule-Guided Language Model Alignment for Text Generation Management in Industrial Use Cases. In *Neurips Safe Generative AI Workshop 2024*.
- Argote, L.; and Ingram, P. 2000. Knowledge transfer: A basis for competitive advantage in firms. *Organizational behavior and human decision processes*, 82(1): 150–169.
- Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Cui, G.; Yuan, L.; Ding, N.; Yao, G.; Zhu, W.; Ni, Y.; Xie, G.; Liu, Z.; and Sun, M. 2023. UltraFeedback: Boosting Language Models with High-quality Feedback. *arXiv:2310.01377*.
- Gao, Z.; Cecati, C.; and Ding, S. X. 2015. A survey of fault diagnosis and fault-tolerant techniques—Part I: Fault diagnosis with model-based and signal-based approaches. *IEEE transactions on industrial electronics*, 62(6): 3757–3767.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Gugger, S.; Debut, L.; Wolf, T.; Schmid, P.; Mueller, Z.; Mangrulkar, S.; Sun, M.; and Bossan, B. 2022. Accelerate: Training and inference at scale made simple, efficient and adaptable. <https://github.com/huggingface/accelerate>.
- Gururangan, S.; Marasović, A.; Swayamdipta, S.; Lo, K.; Beltagy, I.; Downey, D.; and Smith, N. A. 2020. Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Khodabandelou, G.; Khodabandelou, R.; Ghasemi, M.; Aghaei, M.; et al. 2024. Challenges and Barriers of Using Large Language Models (LLM) Such as ChatGPT for Diagnostic Medicine: A Scoping Review. *Journal of Clinical Medicine*, 13(2): 365.
- Kumar, A.; Amin, E. M.; Lee, X. Y.; Vidyaratne, L.; Farahat, A. K.; Ghosh, D. D.; Koreeda, Y.; and Gupta, C. 2025. Building Domain-Specific Small Language Models via Guided Data Generation. *arXiv preprint arXiv:2511.21748*.
- Lai, J.; Gan, W.; Wu, J.; Qi, Z.; and Yu, P. S. 2024. Large language models in law: A survey. *AI Open*.
- Li, Y.; Wang, S.; Ding, H.; and Chen, H. 2023. Large language models in finance: A survey. In *Proceedings of the fourth ACM international conference on AI in finance*, 374–382.
- Liu, M.; Ene, T.-D.; Kirby, R.; Cheng, C.; Pinckney, N.; Liang, R.; Alben, J.; Anand, H.; Banerjee, S.; Bayraktaroglu, I.; et al. 2023. Chipnemo: Domain-adapted llms for chip design. *arXiv preprint arXiv:2311.00176*.
- Mangrulkar, S.; Gugger, S.; Debut, L.; Belkada, Y.; Paul, S.; and Bossan, B. 2022. PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods. <https://github.com/huggingface/peft>.
- Meta, A. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. *Meta AI Blog*. Retrieved December, 20: 2024.
- Nguyen, H.-T. 2023. A brief report on lawgpt 1.0: A virtual legal assistant based on gpt-3. *arXiv preprint arXiv:2302.05729*.
- Nonaka, I.; and Takeuchi, H. 1996. The knowledge-creating company: How Japanese companies create the dynamics of innovation. *Long range planning*, 29(4): 592.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pezeshkpour, P.; and Hruschka, E. 2023. Large language models sensitivity to the order of options in multiple-choice questions. *arXiv preprint arXiv:2308.11483*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741.
- Remus, S.; and Biemann, C. 2016. Domain-specific corpus expansion with focused webcrawling. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 3607–3611.
- Shrivastava, A.; and Li, P. 2014. In defense of minhash over simhash. In *Artificial intelligence and statistics*, 886–894. PMLR.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford Alpaca: An Instruction-following LLaMA model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; Millican, K.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Team, G.; Riviere, M.; Pathak, S.; Sessa, P. G.; Hardin, C.; Bhupatiraju, S.; Hussenot, L.; Mesnard, T.; Shahriari, B.; Ramé, A.; et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nature medicine*, 29(8): 1930–1940.
- Thompson, J. 2014. *The perceived benefits of Automotive Service Excellence (ASE) certifications for graduates of four-year Automotive Technology programs*. University of Arkansas.

Wang, P.; Li, L.; Chen, L.; Zhu, D.; Lin, B.; Cao, Y.; Liu, Q.; Liu, T.; and Sui, Z. 2023. Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926*.

Wang, W.; Wei, F.; Dong, L.; Bao, H.; Yang, N.; and Zhou, M. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neural information processing systems*, 33: 5776–5788.

Wu, S.; Irsoy, O.; Lu, S.; Dabrovolski, V.; Dredze, M.; Gehrmann, S.; Kambadur, P.; Rosenberg, D.; and Mann, G. 2023. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564*.

Xia, Y.; Kim, J.; Chen, Y.; Ye, H.; Kundu, S.; Hao, C. C.; and Talati, N. 2024. Understanding the performance and estimating the cost of llm fine-tuning. In *2024 IEEE International Symposium on Workload Characterization (IISWC)*, 210–223. IEEE.

Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.

Zhao, Y.; Gu, A.; Varma, R.; Luo, L.; Huang, C.-C.; Xu, M.; Wright, L.; Shojanazeri, H.; Ott, M.; Shleifer, S.; et al. 2023. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*.

Zhao, Z.; Wallace, E.; Feng, S.; Klein, D.; and Singh, S. 2021. Calibrate Before Use: Improving Few-shot Performance of Language Models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 12697–12706. PMLR.

Zheng, C.; Zhou, H.; Meng, F.; Zhou, J.; and Huang, M. 2024a. Large Language Models Are Not Robust Multiple Choice Selectors. In *The Twelfth International Conference on Learning Representations*.

Zheng, Y.; Zhang, R.; Zhang, J.; Ye, Y.; Luo, Z.; Feng, Z.; and Ma, Y. 2024b. LlamaFactory: Unified Efficient Fine-Tuning of 100+ Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*. Bangkok, Thailand: Association for Computational Linguistics.