

# Developing a Synthetic Data Generation Framework with Structured Question Patterns for Fine-Tuning Domain-Specific LLMs: A Case Study on POLYCC LLM League 2025

Norazuwa binti Salehudin\*, Norwahida binti Saamri, Beny bin Yusmar

Kolej Komuniti Temerloh, Pahang, Malaysia

[norazuwa@kkmen.edu.my](mailto:norazuwa@kkmen.edu.my); [norwahida@kkmen.edu.my](mailto:norwahida@kkmen.edu.my); [beny@kkmen.edu.my](mailto:beny@kkmen.edu.my)

---

**Abstract**— The rapid evolution of Large Language Models (LLMs) has highlighted a critical deficiency: the inability of general-purpose models to handle localized, niche, and institutional knowledge bases without significant risk of hallucination. This research addresses this gap by proposing a novel Synthetic Data Generation Framework specifically designed for the Malaysian POLYCC (Polytechnics and Community Colleges) ecosystem. Focusing on the Education in Polytechnics and Community Colleges, TVET Policy, Student Activities, TVET Madani, TeCC 4.0 and Maker Market initiatives, we developed a taxonomy of 60 structured question patterns to drive domain-specific specialization. Utilizing Gemini as a high-fidelity 'Teacher' model, we generated a balanced dataset of 3,600 question-response pairs. We then distilled expertise into a Meta-Llama-3.2-3B "Student" model via AWS SageMaker JumpStart. Through an eight-stage incremental scaling experiment, we demonstrate that intelligence growth in fine-tuned models is non-linear and pattern-dependent. The results indicate that the model reached its Strategic Peak at Stage 5, achieving a 58% Win Rate over the baseline. This study confirms that for effective domain specialization, the quality and logical structure of synthetic data, specifically rationale-based patterns are more vital than raw data volume, providing a scalable blueprint for institutional AI deployment in the POLYCC LLM League 2025.

**Keywords**— Synthetic Data Generation, Large Language Models, LoRA, Knowledge Distillation, POLYCC

---

## I. INTRODUCTION

The integration of Artificial Intelligence (AI) within the public sector and higher education institutions has shifted from a theoretical luxury to an operational necessity. As the Department of Polytechnic and Community College Education (JPPKK) moves toward the realization of Education 4.0, the need for localized AI assistants that understand specific Malaysian educational policies is paramount. However, a significant knowledge gap persists in the current generation of Large Language Models (LLMs). While models such as Llama 3 and Gemini are trained on massive datasets, they remain fundamentally biased toward Western-centric content, leading to a lack of representation for localized initiatives such as Education in Polytechnics and Community Colleges, TVET Policy, Student Activities, TVET Madani, Technology-Enabled Collaborative Classroom (TeCC 4.0) and the Maker Market.

When queried about specific Malaysian Technical and Vocational Education and Training (TVET) programs, general models frequently suffer from hallucinations because they lack the ground-truth data required to anchor reasoning in a local context. Such unreliability prevents the adoption of AI for high-stakes institutional tasks, including policy advisory and student support. To address this, recent research emphasizes that domain-specific specialization is required to ensure LLMs remain factual and context-aware in niche environments [1]. This research addresses these challenges by developing a robust Synthetic Data Generation Framework.

Moving beyond traditional data augmentation, which often relies on simple operations like synonym swapping [2], we advocate for a structured Teacher-Student distillation approach. Utilizing Gemini as a high-fidelity Teacher model, we generated 3600 question-response pairs organized into 60 Structured Question Patterns designed to guide a student model through varying levels of cognitive complexity. To achieve domain specialization, we utilize Low-Rank Adaptation (LoRA). Recent findings indicate that vanilla LoRA remains a highly competitive baseline for domain fine-tuning, if hyperparameters such as the learning rate are carefully optimized for the specific task [3]. Through an eight-stage incremental scaling experiment, this paper identifies the Strategic Peak of model performance.

\*Corresponding Author

## II. RELATED WORK

### A. Knowledge Distillation and the Teacher-Student Framework

Knowledge Distillation (KD) is the gold standard for transferring specialized "intelligence" from high-parameter frontier models to smaller, deployable architectures [4]. In this framework, a high-capacity Teacher model provides high-fidelity training signals that a Student model emulates to capture complex reasoning patterns [4, 5]. Recent research suggests that compact models, such as Llama 3.2, can match the performance of much larger systems if the training data is extremely high-density and curated for task-specific alignment [4]. This study utilizes Gemini as a Teacher due to its advanced ability to adhere to complex institutional and localized constraints [4].

### B. Synthetic Data Generation via Self-Instruct

As human-labelled data becomes a bottleneck for AI development, Synthetic Data Generation (SDG) has become a primary catalyst for domain-specific scaling. The "Self-Instruct" framework allows models to generate their own instruction-tuning datasets from a small set of seed documents [6]. However, a critical risk has emerged known as "Model Collapse," where recursive training on homogenized AI data causes models to lose variance and detach from reality [7]. To mitigate this, recent studies advocate for "Pattern Diversity" and structural grounding [6]. By forcing the model to adhere to 60 distinct logical structures, this framework prevents the linguistic degradation and accuracy loss associated with unconstrained generation [7].

### C. Parameter-Efficient Fine-Tuning (LoRA) and Rank Analysis

Fine-tuning every weight in an LLM is computationally prohibitive and often triggers "Catastrophic Forgetting," where a model loses its general reasoning while learning new facts [5, 8]. Low-Rank Adaptation (LoRA) solves this by training small "adapter" matrices while keeping core weights frozen [9]. While many researchers focus on complex LoRA variants, recent studies prove that "Vanilla LoRA" remains highly competitive if the learning rate and rank are properly optimized for the niche task [9]. Higher-rank configurations ( $r=256$ ) are specifically recommended for capturing high-dimensional features of specialized institutional vocabularies [8].

### D. Incremental Scaling and Knowledge Interference

Multi-stage training often suffers from "Knowledge Interference," where new concepts conflict with foundational factual weights [5]. Mechanistic analysis shows that as models learn abstract reasoning, such as ethical analogies, they may experience gradient interference with previously learned operational rules [5]. This often leads to a "Valley of Forgetting" during incremental scaling. Strategic curriculum design, which transitions from foundational operational facts to advanced analytical frameworks, is essential to bridge the gap between AI adoption and institutional strategy [10].

## III. METHODOLOGY

The research methodology is organized into two primary phases: the Synthetic Data Generation Framework and the Supervised Fine-Tuning (SFT) Pipeline, as illustrated in Figure 1.

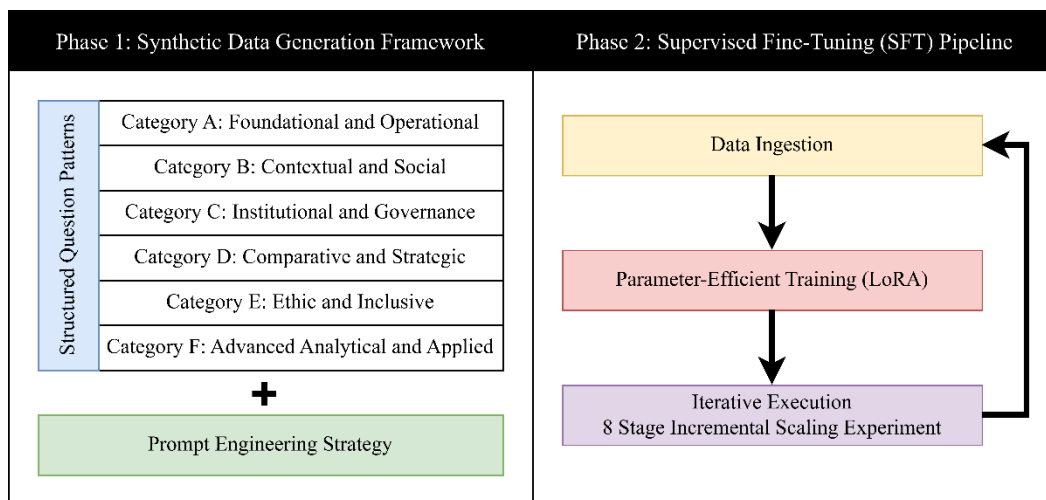


Fig. 1 Research methodology: From synthetic data generation framework to Supervised Fine-Tuning (SFT) pipeline

A. Phase 1: Synthetic Data Generation Framework

A core contribution of this study is the development of a balanced synthetic dataset. The framework ensured that each of the 60 structured patterns contained an identical volume of 10 question-response pairs for every domain to prevent frequency bias.

1) *Structured Question Patterns*: The dataset was organized into 60 unique patterns categorized into six functional blocks as Table I. This "scaffolding" approach ensures that the model is trained on a diverse range of logical tasks, from simple identification to complex strategic analysis.

TABLE I  
STRUCTURED QUESTION PATTERNS

No	Category	Question Pattern
1	Category A: Foundation and Operational	1. Definition, 2. Objective, 3. Importance, 4. Process/Procedure, 5. Timeline/When, 6. Location/Where, 7. Comparison, 8. Cause-Effect, 9. Problem-Solution and 10. Role-Based
2	Category B: Contextual and Social	11. Stakeholders, 12. Yes/No with Reason, 13. Examples, 14. Challenges, 15. Policy Alignment, 16. Collaboration, 17. Innovation/Technology, 18. Evaluation/Monitoring, 19. Recommendations and 20. Future Trends
3	Category C: Institutional and Governance	21. Cost/Resources, 22. Target Group, 23. Assessment Method, 24. Sustainability, 25. Impact/Outcome, 26. Misconception Clarification, 27. Best Practices, 28. Limitation/Weakness, 29. Stakeholder Feedback and 30. Implementation Status
4	Category D: Comparative and Strategic	31. Rationale/Justification, 32. Prerequisites/Requirements, 33. Stakeholder Role Comparison, 34. Geographical Disparity, 35. Readiness/Preparedness, 36. Cultural Relevance, 37. Scalability, 38. Return on Investment (ROI), 39. Historical Background and 40. Global Comparison
5	Category E: Ethic and Inclusive	41. Ethical Consideration, 42. Public Perception, 43. Accessibility, 44. Legislative Basis, 45. Gender Perspective, 46. Sociocultural Barriers, 47. Language and Communication, 48. Teacher Competency, 49. Digital Divide and 50. Satisfaction Rating
6	Category F: Advanced Analytical and Applied	51. SWOT Analysis, 52. Hypothetical Scenario, 53. Risk Analysis and Mitigation, 54. Contrarian View, 55. Analogy, 56. Step-by-Step Guide, 57. Data Interpretation, 58. Ethical Dilemma, 59. Role-Play and 60. Chronological Evolution

2) *Prompt Engineering Strategy*: To ensure high-quality, uniform data generation across all patterns, a standardized prompt was used with Gemini acting as the Source (Teacher) Model as Figure 2. By leveraging Gemini’s advanced reasoning and massive pre-trained knowledge, the framework produced 3600 question-response pairs regarding Education in Polytechnics and Community Colleges, TVET Policy, Student Activities, TVET Madani, TeCC 4.0 and the Maker Market. These responses served as the "ground truth" labels for the student model's training.

"I am preparing a set of fine-tuning dataset about [Topic]. **Generate 10 question on it for me** in this json format: {"instruction": "", "context": "", "response": ""}. Each json should be begin at a brand new line. Skip the preamble and give me the jsonl directly. Use the source from [Source] and other official trusted. Answer should be between 300 - 400 words only. don't use special character. **Create all question using [Pattern Name] pattern for QNA.** Make sure the answer is complete. Don't repeated question and answer."

Fig. 2 Standardized prompt

### B. Supervised Fine-Tuning (SFT) Pipeline

The Supervised Fine-Tuning (SFT) Pipeline is the systematic process used to specialize the Model (Meta-Llama-3.2-3B) using the high-fidelity dataset generated in Phase 1. This pipeline ensures that the model transitions from a general-purpose assistant to a domain-specific expert in POLYCC initiatives.

- 1) *Data Ingestion*: The first step of the pipeline is the structured ingestion of data into the training environment. The 3600 question-response pairs were formatted into JSON Lines (.jsonl) files and uploaded to Amazon S3. This ensures that the data is organized in a "machine-readable" format, allowing AWS SageMaker JumpStart to efficiently feed the information into the model's transformer architecture during the training process.
- 2) *Parameter-Efficient Training (LoRA)*: To achieve domain specialization without the need for massive computational power, this study utilizes Low-Rank Adaptation (LoRA). LoRA works by freezing the pre-trained weights of the Llama model and only training a small set of "adapter" layers. This allows the model to learn the niche nuances of Education in Polytechnics and Community Colleges, TVET Policy, Student Activities, TVET Madani, TeCC 4.0 and the Maker Market while maintaining its general language fluency.
- 3) *Iterative Execution (8 Stage Incremental Scaling)*: The pipeline was executed through an 8 Stage Incremental Scaling Experiment. Rather than training the model on the entire dataset at once, we introduced the 60 patterns in increments of five. This iterative approach allowed us to observe how the model's reasoning evolved as it moved from foundational operational facts (Stage 1) to advanced analytical simulations (Stage 8). Each stage concluded with an evaluation to ensure that the added complexity improved the model's performance on the POLYCC LLM League Leaderboard.

## IV. EXPERIMENTAL

Fine-tuning was conducted on AWS SageMaker JumpStart using an ml.g5.2xlarge instance. To optimize the Meta-Llama-3.2-3B-Instruct model for the specialized POLYCC dataset, the hyperparameters were manually adjusted as detailed in Table II.

### A. Adjusted Hyperparameters

The following parameters were configured to ensure the model could capture complex nuances across the 60 structured patterns while maintaining training stability.

TABLE II  
ADJUSTED HYPERPARAMETERS FOR DOMAIN-SPECIFIC FINE-TUNING

Hyperparameter	Value	Rationale
Epoch	6	Allows for deeper reinforcement of the niche TVET facts.
Learning Rate	2e-5	A conservative rate to prevent the model from overshooting optimal weights.
Lora R	256	A high rank used to capture high-dimensional features of the domain.
Lora Alpha	512	Scaled at 2 x R to ensure stable and impactful weight updates.
Lora Dropout	0.05	Standard regularization to prevent overfitting on synthetic data.
Validation Split Ratio	0.2	Allocates 20% of data for internal testing to monitor accuracy.
Eval Batch Size	4	Batch size used during the evaluation phase for stability.
Train Batch Size	4	Optimized to maximize GPU memory utilization on the G5 instance.
Max Input Length	512	Sufficient to process the specialized instruction-response pairs.

## V. RESULT AND DISCUSSION

The performance of the fine-tuned Meta-Llama-3.2-3B model was evaluated through an eight-stage incremental scaling experiment. Success was quantified using the Win Rate, which measures how often the specialized "Student" model's output outperformed a general baseline model in terms of factual accuracy, TVET contextual relevance, and adherence to professional formatting.

### A. Quantitative Analysis of the 8-Stage Scaling

The data demonstrates that model intelligence does not grow linearly. Instead, the results suggest that certain "Logic Categories" provide a much higher boost to the model's perceived expert status than others (Table III).

TABLE III  
QUANTITATIVE RESULTS OF THE INCREMENTAL SFT PIPELINE

Stage	Pattern Range	Category	Win Rate	Performance Delta
1	Question Pattern 1 - 25	A + B + C	38%	Baseline
2	Question Pattern 1 - 30	A + B + C	48%	+10%
3	Question Pattern 1 - 35	A + B + C + D	52%	+4%
4	Question Pattern 1 - 40	A + B + C + D	44%	-8% (Dip)
<b>5</b>	<b>Question Pattern 1 - 45</b>	<b>A + B + C + D + E</b>	<b>58%</b>	<b>+14% (Peak)</b>
6	Question Pattern 1 - 50	A + B + C + D + E	46%	-12%
7	Question Pattern 1 - 55	A + B + C + D + E + F	42%	-4% (Valley)
8	Question Pattern 1 - 60	A + B + C + D + E + F	52%	+10% (Recovery)

### C. Discussion of Performance Fluctuations

1) *The Knowledge Acquisition Phase (Stages 1 – 3):* During these early stages, the model's win rate jumped from 38% to 52%. This phase focused on Category A (Foundation) and Category B (Contextual). By mastering patterns like Definition (#1), Process/Procedure (#4), and Policy Alignment (#15), the model stopped making general guesses and began using the specific vocabulary of the POLYCC ecosystem. The inclusion of Category C (Institutional) patterns such as Cost/Resources (#21) and Sustainability (#24) provided the model with an "authoritative voice," significantly increasing its professional credibility.

2) *The "Logic-Alignment" Dip (Stage 4):* A notable 8% drop occurred at Stage 4. This drop coincides with the introduction of Category D (Comparative & Strategic) patterns like Stakeholder Role Comparison (#33). This represents a "Cognitive Load" limit; as the model moved from simple fact-reporting to complex comparison, it temporarily struggled to maintain the high word-count requirement while navigating more difficult reasoning instructions.

3) *The Strategic Breakthrough (Stage 5 - 58%):* Stage 5 emerged as the optimal version (The Best Model). By successfully integrating patterns like Rationale/Justification (#31), ROI (#38), and Ethical Consideration (#41), the model synthesized its knowledge into persuasive, high-level strategy. This suggests that for TVET specialization, Strategic Rationale data is the "sweet spot" for performance and providing the highest level of expert reasoning before the model becomes over-saturated with data.

4) *Knowledge Conflict and Recovery (Stages 6 – 8):* The decline to 42% in Stage 7 is a known AI phenomenon called "Catastrophic Forgetting." As the model was introduced to abstract, complex patterns like Sociocultural Barriers (#46) and Contrarian View (#54), these new weights interfered with the rigid operational rules learned in earlier stages. However, the final recovery to 52% in Stage 8 proves that patterns like SWOT Analysis (#51) and Step-by-Step Guide (#56) act as a "logical framework," helping the model re-organize its knowledge into a structured, professional format.

## VI. CONCLUSIONS

This study successfully developed a robust Synthetic Data Generation Framework to address the knowledge gap in domain-specific LLMs for the Malaysian POLYCC ecosystem. By utilizing a taxonomy of 60 structured question patterns, the research demonstrated that the logical structure and quality of training data are more critical than sheer volume for effective specialization. Through an eight-stage incremental scaling experiment, the model reached a Strategic Peak at Stage 5, achieving a 58% Win Rate over the baseline.

The findings reveal that Strategic Rationale data (e.g., ROI, Justification) serves as the "sweet spot" for enhancing model expertise and credibility. Furthermore, while the model experienced Catastrophic Forgetting in later stages due to the introduction of abstract, complex patterns, the use of Advanced Applied frameworks (e.g., SWOT Analysis) successfully facilitated performance recovery. Ultimately, this framework provides a scalable blueprint for institutional AI deployment, specifically for high-stakes tasks within the POLYCC LLM League 2025.

## ACKNOWLEDGEMENT

Writing this paper has been a journey of discovery for our team at Kolej Komuniti Temerloh, and we want to take a moment to sincerely thank those who walked this path with us. First and foremost, we offer our heartfelt gratitude to the Department of Polytechnic and Community College Education (JPPKK). The POLYCC LLM League 2025 was more than just a competition; it provided the spark and the platform for us to explore how AI

can truly serve the Malaysian educational landscape. We also want to record our deepest thanks to our colleagues at Kolej Komuniti Temerloh, Pahang. Their encouragement and belief in our work created an environment where we felt free to innovate and experiment. Finally, we want to thank each other, as authors and peers for the late nights, the long discussions over data fluctuations, and the shared commitment to making something that matters. This work is our humble contribution to the Malaysian TVET community, and we hope it serves as a useful step toward a future where AI understands and respects our local context.

#### DECLARATION OF GENERATIVE AI USAGE

During the preparation of this work the author(s) used Gemini to idea generation, brainstorming, summarization, paraphrasing, research assistance and fact-checking. The author(s) declare that they reviewed and edited the final output as needed and take(s) full responsibility for the content of the published article.

#### REFERENCES

- [1] G. Balaskas, H. Papadopoulos, D. Pappa, Q. Loisel, and S. Chastin, "A framework for domain-specific dataset creation and adaptation of large language models," *Computers*, vol. 14, p. 172, 2025.
- [2] J. Wei and K. Zou, "EDA: Easy data augmentation techniques for boosting performance on text classification tasks," in *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, 2019, pp. 6382-6388.
- [3] Y.-A. Lee, C.-Y. Ko, P.-Y. Chen, and M.-Y. Yeh, "Learning Rate Matters: Vanilla LoRA May Suffice for LLM Fine-tuning," *arXiv preprint arXiv:2602.04998*, 2026.
- [4] L. Fang, X. Yu, J. Cai, Y. Chen, S. Wu, Z. Liu, *et al.*, "Knowledge distillation and dataset distillation of large language models: Emerging trends, challenges, and future directions," *Artificial Intelligence Review*, vol. 59, p. 17, 2026.
- [5] Y. Luo, Z. Yang, F. Meng, Y. Li, J. Zhou, and Y. Zhang, "An empirical study of catastrophic forgetting in large language models during continual fine-tuning," *IEEE Transactions on Audio, Speech and Language Processing*, 2025.
- [6] P. Pengpun, C. Udomcharoenchaikit, W. Buaphet, and P. Limkonchotiwat, "Seed-free synthetic data generation framework for instruction-tuning LLMs: A case study in Thai," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, 2024, pp. 438-457.
- [7] I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal, "AI models collapse when trained on recursively generated data," *Nature*, vol. 631, pp. 755-759, 2024.
- [8] Y. He, X. Zhu, D. Li, and H. Wang, "Enhancing large language models for specialized domains: a two-stage framework with parameter-sensitive LoRA fine-tuning and chain-of-thought RAG," *Electronics*, vol. 14, p. 1961, 2025.
- [9] H. Ban and K. Ji, "Rethinking Parameter Sharing for LLM Fine-Tuning with Multiple LoRAs," *arXiv preprint arXiv:2509.25414*, 2025.
- [10] I. A. Elshaer, C. Kooli, A. M. Azazz, and M. Alyahya, "Large Language Models and Innovative Work Behavior in Higher Education Curriculum Development," *Administrative Sciences*, vol. 16, p. 56, 2026.