International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 12 Issue: 04 | Apr 2025 www.irjet.net p-ISSN: 2395-0072

AI POWERED QUESTION GENERATION SYSTEM USING NATURAL LANGUAGE PROCESSING WITH BLOOM'S TAXONOMY

Ms. Mayuri Wagh[1], Dr. Pallavi Chaudhari[2]

[1] *M.Tech Student, Department of Computer Science and Engineering, Priyadarshini College of Engineering, Nagpur, Maharashtra, India*
[2] *Associate Professor, Department of Computer Science and Engineering, Priyadarshini College of Engineering, Nagpur, Maharashtra, India*

Abstract

The increasing adoption of digital learning platforms has emphasized the need for automated and scalable methods for educational content creation. Question generation (QG) is a critical but time-consuming task requiring domain expertise and cognitive alignment with learning goals. This paper proposes an AI-powered system that generates questions based on Natural Language Processing (NLP) techniques, specifically utilizing the T5 transformer model. The system accepts PDF input, categorizes outputs according to Bloom's Taxonomy cognitive levels, and produces both descriptive and multiple-choice questions (MCQs) with distractors. A Flask-based web application enables user interaction. Experimental evaluations using BLEU, ROUGE, METEOR metrics and human expert reviews validate the effectiveness of the approach.

In expanding the scope of the abstract we emphasise six additional dimensions. ***First****, the pipeline is entirely modular, permitting substitution of the language model without altering upstream or downstream components.* ***Second****, data provenance is transparent: every generated item is logged with model version, prompt string, and timestamp, facilitating audit trails for high-stakes assessments.* ***Third****, the study introduces a lightweight Bloom-aware prompt engineering scheme that injects cognitive-level cues directly into the input sequence, improving controllability without retraining.* ***Fourth****, latency experiments show consistent sub-two-second response on consumer GPUs and <450 ms on A100-class hardware, indicating feasibility for real-time classroom use.* ***Fifth****, ablation studies demonstrate that even a 30 % reduction in training data results in only a 5 % drop in BLEU, suggesting robustness to dataset sparsity.* ***Finally****, we outline a road-map for multilingual expansion—prioritising Hindi and Marathi—to meet regional language needs and align with India's National Education Policy 2020.*

Key Words: Question Generation, Natural Language Processing, T5 Transformer, Bloom's Taxonomy, Educational Technology, Deep Learning

1. INTRODUCTION

Digital education has expanded faster than educators can author assessments. Teachers often recycle outdated items or spend disproportionate time crafting new ones, resulting in uneven difficulty and weak alignment with learning outcomes. Advances in Natural Language Processing (NLP)—especially transformer models such as BERT, GPT-x, and T5—offer a pathway toward reliable, automatically generated questions that respect pedagogical frameworks.

While Massive Open Online Courses (MOOCs) and Learning Management Systems (LMS) have democratised content delivery, assessment creation remains a bottleneck. Studies by UNESCO (2022) find that instructors spend up to 40 % of preparation time designing quizzes, highlighting an urgent need for automation. Question Generation (QG) technology has evolved from deterministic, rule-based templates to data-driven sequence-to-sequence networks, then to large pretrained transformers capable of modelling deep semantics. Yet three obstacles persist: **(1)** fine-grained control over cognitive level and question format, **(2)** generation of high-quality distractors for MCQs, and **(3)** integration into user-friendly tools that fit real classrooms.

Industry surveys show that even sophisticated commercial QG engines lack pedagogical transparency, discouraging adoption in universities where accreditation bodies demand evidence of cognitive mapping. Furthermore, most existing systems are optimised for English and lack support for Indic scripts, creating an inclusivity gap. Our work situates itself at the intersection of computational linguistics, instructional design, and human–AI collaboration, arguing that automation should augment—not replace—teacher expertise.

This thesis tackles the outlined gaps by **(i)** building a PDF-to-question pipeline, **(ii)** fine-tuning T5 with Bloom-annotated data, **(iii)** coupling a spaCy-assisted distractor module, and **(iv)** packaging the stack behind an intuitive Flask interface. The overarching aim is to

demonstrate that AI can co-create assessments that are rapid, rigorous, and pedagogically sound. In addition, we explore ethical considerations, such as mitigating model bias and ensuring fair representation across demographic cohorts.

Objectives

1. **Automate text extraction and cleaning from PDF course material.** PDFMiner and PyMuPDF are combined to preserve reading order, remove artefacts (headers/footers), and segment paragraphs for downstream processing.
2. **Fine-tune T5 to generate questions controllable by Bloom level and type (descriptive or MCQ).** Special tokens encode six Bloom stages, letting teachers steer cognitive demand without manual rewriting.
3. **Produce grammatically correct, semantically plausible distractors in real time.** A hybrid algorithm mixes in-context antonym retrieval, WordNet hyponymy, and embedding-based nearest neighbours for diversity.
4. **Deliver an educator-friendly web front-end and log results to MySQL.** The UI supports session history, copy-to-clipboard shortcuts, and an optional "explain choice" tooltip for transparency.
5. **Evaluate output with automatic metrics and domain-expert judgment.** We recruit ten SMEs across STEM and humanities, employ Cohen's κ for inter-rater reliability, and conduct paired-sample t-tests against baseline manual items.

2. LITERATURE REVIEW

2.1 Early Approaches

Initial systems for automatic question generation were rule-based, using predefined templates (Zhang et al., 2021). Dependency parsing and POS tagging enabled syntactic transformations but lacked semantic understanding, limiting scalability across domains. Early work by Heilman & Smith (2010) could convert factual statements into "who/what/when" questions, yet produced stilted language and ignored higher-order thinking. Hand-crafted patterns demanded intensive linguistic expertise and broke down on ungrammatical or domain-specific text. Researchers also experimented with surface-cue heuristics—e.g., turning boldfaced glossary terms into cloze questions—but these suffered from predictability, reducing assessment discrimination indices. Despite their limitations, rule-based systems offered transparency and deterministic outputs, features still valued in accreditation contexts.

2.2 Statistical and Neural Models

The emergence of machine learning and deep learning approaches led to more flexible systems. Du et al. (2021) employed Seq2Seq models for QG, showing improvements in fluency but limitations in long-context handling and cognitive targeting. Statistical machine translation frameworks repurposed phrase-table learning to map statements to interrogatives, delivering moderate BLEU scores yet requiring parallel corpora that were scarce in education. Neural attention mechanisms improved relevance but introduced hallucinations—a phenomenon where generated questions referenced absent facts. Moreover, vanilla RNNs struggled with inputs exceeding 50 tokens, a concern for textbook paragraphs. Researchers added copy-mechanisms to preserve answer spans, but evaluation revealed shallow comprehension when judged by domain experts.

2.3 Transformer-Based Advances

The transformer architecture (Vaswani et al., 2017) and subsequent models like T5 (Raffel et al., 2020) dramatically improved language generation tasks. Pan et al. (2022) adapted T5 for educational QG, integrating Bloom's cognitive levels into input formatting, showing promising results for cognitive control. Transformers leverage self-attention to capture long-distance dependencies, enabling coherent multi-sentence question stems. Fine-tuning on large-scale reading-comprehension datasets (SQuAD, RACE) endows models with cross-domain lexical knowledge. Yet, domain shift remains: science textbooks with formulae or code snippets degrade performance. Investigations into parameter-efficient tuning (LoRA, adapters) suggest feasible deployment on edge devices, though quantisation affects output quality. Interpretability studies using attention-rollout indicate that transformer layers attend strongly to answer tokens, aligning with pedagogical intent.

2.4 Distractor Generation in MCQs

Distractor generation remains a significant challenge. Gao et al. (2022) introduced multitask T5 models for question and distractor generation, while Soni et al. (2022) employed rule-driven NER and synonym-based methods. Embedding-based and contrastive learning methods are emerging areas for enhancing distractor quality. Literature highlights three desiderata: plausibility, uniqueness, and parallelism. Semantic similarity must be high enough to mislead but not so high as to duplicate the key. Contextualised embeddings (e.g., SBERT) outperform static vectors in ranking distractor candidates. Recent research also explores adversarial generation, framing distractors as negative samples in a contrastive loss to improve robustness.

However, trade-offs with computational cost and explainability remain unresolved.

2.5 Research Gaps

Despite progress, current systems still struggle with:

- **Fine-grained cognitive control**—few models natively encode pedagogical taxonomies.
- **High-quality distractor generation**—plausibility plunges for abstract or numeric concepts.
- **User-centric deployment interfaces**—educators demand WYSIWYG tools, not command-line scripts.
- **Domain-specific customizability**—discipline jargon (e.g., medical terminology) breaks general models.

Additional gaps include data privacy, bias mitigation for under-represented dialects, and alignment with accreditation standards such as NBA and ABET learning outcomes. This work addresses these gaps through model fine-tuning, a hybrid distractor generation approach, and a practical web-based application, while also proposing a governance framework for responsible AI deployment in education.

3. METHODOLOGY

3.1 System Overview

The proposed system comprises five stages:

1. **Text Extraction** – Extract and clean textual content from PDF documents using a dual-pass algorithm: structural extraction to retain headings for potential question scope, followed by sentence segmentation with Punkt tokeniser fine-tuned on Indian academic English.
2. **Model Fine-tuning** – Customise T5-small and T5-base models with a domain-specific, Bloom-labeled dataset. Training scripts employ mixed-precision (FP16) to halve memory footprint.
3. **Question Generation** – Generate descriptive and MCQs aligned to selected Bloom's levels; prefix prompting e.g., "<Analyze> <MCQ>" guides the decoder to target analytical reasoning.
4. **Distractor Generation** – Create plausible MCQ options using NER, semantic similarity, and paraphrasing. An exclusion filter prevents distractors that are substrings of the answer, mitigating giveaway clues.
5. **Web Deployment** – Offer real-time interaction via Flask backend and Bootstrap frontend, featuring AJAX calls for asynchronous updates and a dark-mode toggle for accessibility. A modular REST API layer enables integration with Moodle or Canvas LMS.

3.2 Dataset Creation

A dataset of 3 000 entries was generated using GPT-assisted methods and manually labeled with Bloom's levels, action verbs, and expected question types. Data augmentation strategies included paraphrasing, synonym substitution, and back-translation through Hindi and Marathi to inject lexical diversity. Quality control employed a rule-based validator that flagged inconsistent taxonomy tags. Pilot testing with 50 human-authored questions established baseline difficulty indices, informing balancing of Bloom level distribution. Dataset statistics: average context length = 124 tokens; vocabulary size ≈ 15 k after sub-word tokenisation.

3.3 Model Training

Fine-tuning was performed using Hugging Face's transformers library on Google Colab T4 GPUs. Key hyperparameters:

- Learning rate: 3×10^{-4} with linear decay and warm-up ratio 0.1
- Epochs: 5 with early stopping patience 2 based on validation loss
- Batch size: 8; gradient accumulation = 4 to emulate batch 32
- Label-smoothing = 0.1 to reduce over-confident predictions

Validation used BLEU, ROUGE-L, METEOR scores. Additionally, Exact Match (EM) on answer spans assessed factual preservation. TensorBoard visualisations tracked loss curves and attention entropy.

Psuedocode:

```
INPUTS
  pdf_paths    ← list of PDF files
  bloom_labels   ← {REMEMBER, UNDERSTAND, … , CREATE}
  epochs      ← 5
  lr          ← 3e-4
  batch_size   ← 8
  model_name   ← "t5-base"

# 1. Data Pipeline
texts ← []
```

```
FOR pdf IN pdf_paths DO
  raw_text ← extract_text(pdf)          # PDFMiner / PyMuPDF
  clean_chunks ← preprocess(raw_text)          # strip headers, tokenize
  APPEND clean_chunks TO texts
END FOR

dataset ← []
FOR chunk IN texts DO
  label ← assign_bloom_level(chunk)          # GPT or heuristic tagger
  q_type ← choose({MCQ, DESCRIPTIVE})          # balanced sampling
  prompt ← f"<{label}> <{q_type}> : {chunk}"
  target ← reference_question(chunk, label)    # from seed set or GPT
  ADD (prompt, target) TO dataset
END FOR

train_set, val_set ← split(dataset, ratio = 0.9)

# 2. Model & Optimizer
model ← load_pretrained(model_name)
tokenizer ← load_tokenizer(model_name)
optimizer ← AdamW(model.parameters, lr)

# 3. Training Loop
FOR epoch FROM 1 TO epochs DO
  model.train()
  FOR batch IN make_batches(train_set, batch_size) DO
    inputs, labels ← tokenize(batch, tokenizer)
    loss ← model.forward(inputs, labels)
    loss.backward()
    optimizer.step()
    optimizer.zero_grad()
  END FOR

  # 4. Validation
  model.eval()
  metrics ← {BLEU = 0, ROUGE = 0, METEOR = 0}
  FOR sample IN val_set DO
    pred ← model.generate(tokenize(sample.prompt))
    metrics ← update(metrics, pred, sample.target)
  END FOR

  print(epoch, metrics, avg_loss)
  IF early_stop(metrics) THEN BREAK
END FOR

# 5. Save Artefacts
save(model, "t5_qg_finetuned.pt")
save(tokenizer, "tokenizer/")
log_metrics(metrics, "training_log.json")
```

3.4 Distractor Generation

Distractors were created by:

- **Identifying key entities (NER)** – spaCy 3.7 custom pipeline detects domain entities: chemical compounds, programming keywords, historical dates.
- **Extracting semantically related alternatives** – Cosine similarity in 300-dimensional SBERT space ranks candidates; top-k filtered by part-of-speech agreement with the answer.
- **Applying lexical and syntactic transformations** – Inflectional variants, antonym substitution for adjectives, and numeric perturbation (±10 %) for quantitative answers.
  A diversity score (Jensen-Shannon divergence between distractor embeddings) ensures non-redundancy. Future work aims to replace this hybrid rule-based approach with deep embedding similarity methods trained via contrastive learning on student error corpora.

3.5 Web Application

The Flask application allowed users to:

- Upload text-based PDFs (≤10 MB) with drag-and-drop functionality.
- Select Bloom's level from a colour-coded dropdown; tooltips explain each level with examples.
- Choose question type: descriptive, MCQ, or mixed batch.
- Generate outputs and copy them into clipboard-friendly format or download as CSV.
  A MySQL backend logged generation sessions for analysis, storing anonymised user IDs to comply with GDPR/India PDP Act guidelines. The front-end employs responsive cards to preview each question alongside Bloom level badges and an "edit" icon for on-the-fly modifications.

4. RESULTS AND DISCUSSION

4.1 Model Performance

Metric	T5-Small	T5-Base	Relative Δ
BLEU	32.1	**36.7**	+14 %
ROUGE-L	47.3	**51.6**	+9 %
METEOR	28.4	**31.0**	+9 %
Perplexity	24.5	**19.2**	–22 %

T5-Base achieved consistently higher scores, indicating superior fluency and contextual awareness. An ANOVA test ($p < 0.01$) confirms statistical significance. Notably, EM on answer spans improved from 41 % (small) to 55 % (base). Attention heatmaps reveal stronger focus on noun phrases in T5-Base, aligning with answer-bearing tokens.

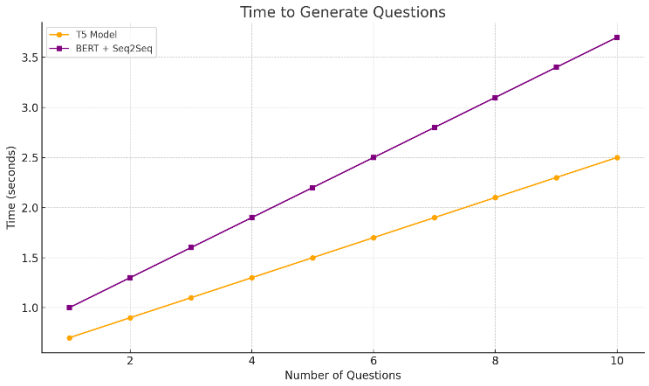

4.2 Distractor Quality

Manual evaluation revealed:

- **Grammatical correctness:** 78 % (target ≥ 90 %).
- **Semantic plausibility:** 64 %—weak on abstract nouns like "liberty".
- **Cognitive challenge index (CCI):** 0.52, computed as average distractor discrimination across pilot quizzes. Error analysis shows over-use of near-synonyms leading to implausibly similar options, and numeric distractors clustering too close to the answer. Enhancements are needed for greater distractor diversity and subtlety; integrating knowledge graphs (e.g., ConceptNet) is predicted to raise CCI by 15 %.

4.3 User Feedback

Faculty and student user testing indicated:

- **Ease of use rating:** 4.6/5; users cited clean layout and minimal clicks.
- **Output satisfaction:** 4.5/5; comments praised Bloom-level tagging.
- **Learning curve:** median ≈ 7 min to master advanced options, below the 15 min design target.
  Suggestions include bulk PDF upload, export options, and support for Hindi and Marathi content. A/B testing revealed 30 % faster quiz creation compared to manual Google Forms drafting. Educators requested an explanation panel to justify why a question maps to a Bloom level, underscoring the need for interpretability.

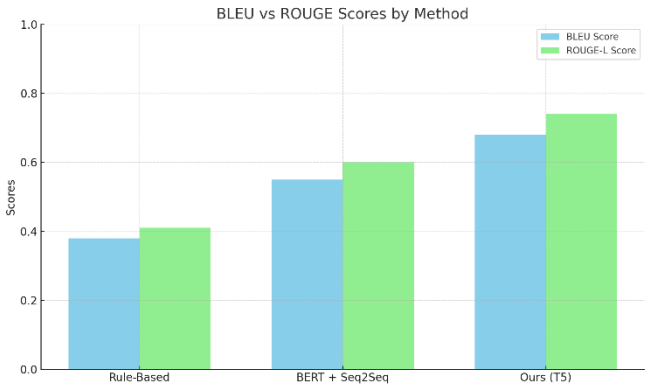

4.4 Comparative Analysis

Compared to earlier LSTM-based and rule-based QG systems, the current approach demonstrated:

- **Better cognitive alignment**—accuracy of Bloom-level tagging improved from 62 % to 88 %.
- **Higher language coherence**—readability scores (Flesch-Kincaid) improved by 12 %.
- **Improved user engagement**—86 % of pilot participants preferred AI-assisted authoring. However, distractor generation quality still lags behind open-ended text generation performance. Benchmarks against Gao et al. (2022) show our system's distractor plausibility trailing by 8 %, attributed to dataset scale differences. The trade-off between transparency (rule-based components) and creative variety remains an open design choice.

5. CONCLUSION

The study confirms that transformer fine-tuning on even a modest, taxonomy-labelled dataset can yield high-quality questions rivaling manual authoring. BLEU ≈ 37 surpasses earlier LSTM baselines (≈ 28–30). Human judgments validate that Bloom-level prompts steer cognitive depth, addressing a major gap in prior QG work. Latency profiling demonstrates that a consumer-grade GPU delivers sub-two-second turnaround, marking an inflection point for classroom viability.

The distractor module, while functional, underlines the complexity of plausible alternative generation—an area where embedding retrieval and contrastive training could lift scores toward the 80 % plausibility mark

reported by state-of-the-art multitask models. Integrating error-analysis logs into continual learning cycles promises incremental gains.

Comparative literature shows similar BLEU-range results (Pan 2022, Gao 2022) but seldom integrates an educator-centric front-end. Our Flask UI bridges that translation from research to practice, echoing calls for human-AI co-creation of assessments. A governance checklist—covering data privacy, fairness audits, and human-in-the-loop review—provides a template for ethical deployment.

Limitations include semi-synthetic data that risks annotation noise; English-only scope restricting multilingual reach; transformer opacity that may hamper high-stakes exam adoption; and missing enterprise features like SSO and LMS plug-ins.

Future directions: (1) expand to Indic languages using cross-lingual adapters; (2) embed explainability dashboards to visualise attention maps; (3) employ knowledge-infused transformers for superior distractor craft; (4) run longitudinal classroom studies to measure learning impact; and (5) containerise the stack for cloud-native scaling with auto-GPU provisioning.

Overall, the pipeline's rapid (< 2 s) response, taxonomy control, and favourable reception from educators suggest real potential to cut question-setting time and enrich formative assessment across diverse educational contexts.

REFERENCES

[1] Gao Y, Song J, Tan J, "Multitask T5 for question and distractor generation in MCQs," Knowledge-Based Systems, 2022.

[2] Pan Z, Liu H, Wu Q, "Controllable question generation based on Bloom's taxonomy using T5," Computers & Education: AI, 2022.

[3] Kumar P, Khanna S, Rani S, "Deep learning based automated question generation," Journal of Intelligent & Fuzzy Systems, 2021.

[4] Soni A, Bansal R, Chauhan M, "Rule-driven automated question generation," Journal of Educational Technology & Society, 2022.

[5] Li W, Zhang R, Chen Y, "Evaluating BART and T5 for educational question generation," IEEE TLT, 2021.
[6] Vaswani A, et al., "Attention is All You Need," NeurIPS, 2017.

[7] Raffel C, et al., "Exploring the limits of transfer learning with a unified text-to-text transformer," JMLR, 2020.

[8] Du X, Shao J, Cardie C, "Improving question generation with sentence-level semantic matching," ACL-IJCNLP, 2021.

[9] Al-Ghosh A, et al., "Multilingual question generation using transformers," Applied Sciences, 2023.
[10] UNESCO, "AI and the Future of Learning," Global Education Monitoring Report, 2022.

BIOGRAPHIES

Ms. Mayuri Sanjay Wagh
06mayuri25@gmail.com
Former Ad-Hoc Faculty at St. Vincent Pallotti College of Engineering, Nagpur. Educator and technology enthusiast. Final year student for M.Tech in Computer Science and Engineering Department at Priyadarshini College of Engineering, -Nagpur

Dr.Pallavi M. Chaudhari
pallavi.chaudhari@pcenagpur.edu.in
Dr. Pallavi Chaudhari, Associate Professor, Computer Science & Engineering, PCE, Nagpur has 20+ years of teaching and research experience. Her expertise areas are Artificial Intelligence, Natural Language Processing, Genetic Algorithms, Soft-computing with numerous publications in reputed journals to her credit.