

Few-Shot Question Answering in Low-Resource Languages using Model-Agnostic Meta-Learning (MAML)

Md Tahseen Eqbal
Assistant Professor
CSE, All Saints College of Technology,
Bhopal
mdtahseen278@gmail.com

Aarif Rasul
School of Engineering Sciences &
Technology (SEST) Jamia Hamdard
New Delhi- 110025
Aarifrasu98@gmail.com

Md Wasim Nehal
Department of Computer Science and
Engineering, Jamia Millia Islamia,
mdwasimnehal786@gmail.com

Md Irshad Anwar
Department of Computer
Science and Engineering IIT Madras
Chennai, 600036
23f1000996@ds.study.iitm.ac.in

Wasim Ahmad Sheikh
Department of Computer Science
and Engineering, Maharshi Dayanand
University
Rohtak,
er.sheikhwasim@gmail.com

Asad Iqbal
School of Engineering Sciences &
Technology (SEST) Jamia Hamdard
New Delhi
Asadiqbal223@gmail.com

Abstract—Question Answering (QA) systems have made tremendous strides in languages with abundant resources, such as English. Unfortunately, model performance is severely constrained in low-resource languages due to the lack of annotated data. The Model-Agnostic Meta-Learning (MAML) framework is proposed in this study as a means of few-shot quality assurance for languages with limited resources. With only a small number of annotated question-answer pairs, the method enables quick domain- or language-specific adaptation. With an emphasis on low-resource Indian languages like Telugu, Bengali, and Hindi, we assess the framework using the multilingual QA standards TyDiQA and XQuAD. Our MAML-based technique achieves an 8.5% increase in F1 score and an 8.2% improvement in Exact Match (EM) over the best fine-tuned baselines, according to experimental data. This method considerably outperforms standard fine-tuning and transfer learning approaches in few-shot circumstances. This study demonstrates how meta-learning can be used to create flexible, scalable quality-assurance systems for languages that aren't widely used.

Keywords—Terms—Meta-Learning, Model Agnostic Meta-Learning (MAML), Few-Shot Learning, Cross-Lingual Transfer, mBERT, XLM RoBERTa, Low-Resource Languages, Transfer Learning, Multilingual NLP

I. INTRODUCTION

The goal of Question Answering (QA), an important NLP activity, is to automatically offer contextually appropriate replies to user questions [1][2]. While transformer-based designs like as BERT and T5 have transformed QA performance, the availability of large-scale annotated datasets is crucial to the success of newer deep learning algorithms [3][4]. There is a severe lack of resources for many languages that are considered low-resource, including several spoken languages in India and other places with many official languages. One major obstacle to implementing QA systems in these languages is the dearth of relevant data [5][6]. Domain mismatch, poor language alignment [7][8] and overfitting to limited data are common problems with traditional transfer learning approaches when moving from high-resource to low-resource environments.

We provide a Model-Agnostic Meta-Learning (MAML) strategy for few-shot QA to tackle this problem. In contrast to traditional training methods [9][10], MAML is able to swiftly adapt to new languages or domains with few labelled samples by learning an ideal parameter initialization.

Key contributions:

- Propose a MAML-based QA framework for low-resource languages.

- Evaluate on multilingual benchmarks (TyDiQA, XQuAD) focusing on Indian languages.
- Demonstrate superior few-shot performance over fine-tuning baselines.

II. LITERATURE REVIEW

A. Addressing Inquiries in Languages with Few Resources

With the introduction of transformer-based models like BERT and T5 and large-scale annotated datasets like SQuAD, question answering (QA) has seen quick advancements [11][12]. Unfortunately, low-resource languages, such as many Indic languages (Hindi, Bengali, Telugu, etc.), are underrepresented, in contrast to high-resource languages like English and Chinese [13][14]. While models like mBERT, XLM-RoBERTa, and mT5 offer multilingual coverage, their performance in low-resource environments is hindered by factors like (i) a lack of representation in pretraining corpora, (ii) the morphological complexity of Indic languages, and (iii) variations in domain and script [15]. Despite efforts by benchmarks such as TyDiQA and XQuAD to close this gap, performance remains much worse than English. This highlights the need to develop methods that use fewer labeled examples.

B. NLP's Few-Shot Learning Method

In low-resource languages, few-shot learning approaches are ideal because they attempt to generalize from little amounts of training data. The natural language processing (NLP) domain has previously used Matching Networks, Prototypical Networks, and Relation Networks for tasks such as intent classification, sentiment analysis, and slot filling [16]. This is in contrast to extractive QA, which necessitates the more involved span prediction, and these methods mainly aim at classification-based issues. Recent research has used few-shot paradigms to quality assurance (QA) by meta-training on various tasks, but how well these models adapt to languages that are both multilingual and morphologically rich is unclear. Learning methods aim to generalize from limited training data, a setting highly relevant to low-resource languages. Earlier approaches, such as Matching Networks, Prototypical Networks, and Relation Networks, have been applied to NLP tasks like intent classification, sentiment analysis, and slot filling [17]. However, these approaches primarily target classification-based problems, whereas extractive QA requires span prediction, which is more complex [18]. Recent works have extended few-shot paradigms to QA using meta-training across multiple tasks, but robust adaptation to multilingual and morphologically rich languages remains underexplored.

Research Gap: Inadequate representation in pretraining corpora, high morphological variability, and script variety hamper the performance of multilingual question answering models like mBERT, XLM-RoBERTa, and mT5 in low-resource languages. One of the most important parts of extractive quality assurance is span extraction, yet current few-shot learning approaches, such as prototype networks and matching networks, are more suited to classification problems. Similarly, there has been little research into the use of meta-learning frameworks like MAML and its derivatives for quality assurance in really low-resource multilingual contexts, despite their apparent promise in text categorization, sentiment analysis, and named entity identification. On top of that, when taught with few instances, previous efforts at cross-lingual transfer often experience overfitting and domain mismatch. Because of this, there is a noticeable deficiency in the literature on low-resource Indic languages and meta-learning techniques to few-shot extractive QA, where the capacity to.

III. METHODOLOGY

A. Dataset

Here, we utilize two popular multilingual QA benchmarks to test our suggested framework. For starters, there's TyDiQA, which includes eleven typologically varied languages. Among them, we find the low-resource Indic languages that are of interest, such as Bengali, Telugu, and Hindi. In order to evaluate few-shot learning and cross-lingual transfer in practical multilingual settings, TyDiQA offers a durable testbed.

Based on the English SQuAD dataset, the second dataset, XQuAD, is a parallel multilingual QA benchmark that includes 240 paragraphs and 1,190 question-answer pairs translated into different languages. Building a pre-processing pipeline to get the data ready for meta-learning is important since we're concentrating on few-shot adaptation. Data gathering, cleaning, tokenization, and few-shot sampling are all steps in the pipeline that mimic low-resource scenarios.

The data is then prepared for quality assurance models, divided into train, validation, and test subsets, and enhanced with language information encoding before to being saved for training. As a result, consistency among languages is guaranteed and comparability across few-shot configurations is maintained

B. Model Architecture

We built our suggested model on an MAML-based system for few-shot question answering, whereby questions and context passages are tokenized and then processed via a multilingual encoder. Given Figure 1 their shown abilities in multilingual representation learning, we conduct experiments using XLM-RoBERTa and mBERT as the basic encoders [19]. Additionally, we include a QA span prediction head onto the encoder, which provides the probability distribution for the response span's beginning and ending places inside the passage. The meta-learning part uses the tried-and-true MAML optimization method. By quickly specializing the model parameters in the inner loop based on a limited support set of few-shot instances from the target language, fast specialization is achieved. To guarantee cross-lingual generalization, the meta-learner optimizes the initialization across tasks obtained from several languages in the outer loop. Perfect for low-resource QA jobs, this model's two-level optimization technique allows it to effectively adapt to new languages or domains using only a few of annotated cases.

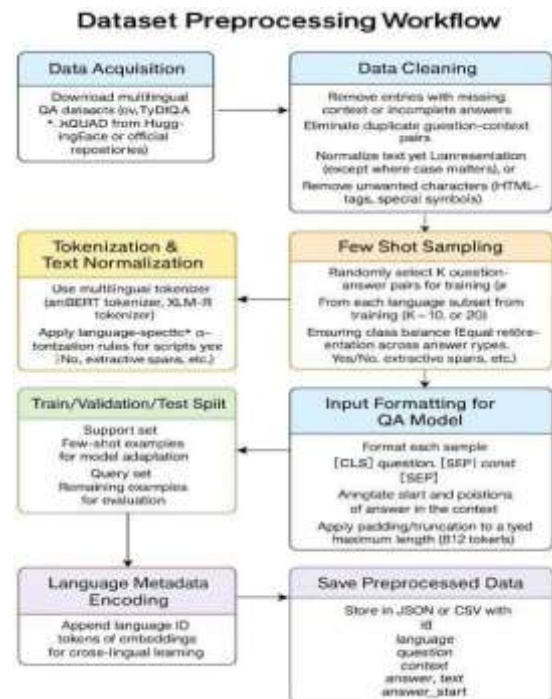


Fig. 1. Workflow for Preprocessing Multilingual QA Datasets (TyDiQA, XQuAD) for Few-Shot Learning in Low-Resource Languages.

C. Evaluation Metrics

Using two popular QA assessment measures, we can see how well our framework meets expectations. When comparing performance metrics, Exact Match (EM) looks at how often predictions match the ground truth answers perfectly [20], while F1 Score captures the overlap between predicted and reference responses at the token level, making it a more sophisticated approach. Precision and recall in extractive QA activities may be assessed using these criteria together.

D. Experimental Setup

As a fair comparison in few-shot circumstances, we compare our technique against many strong baselines, such as fine-tuned mBERT, XLM-RoBERTa, and mT5.

Simulating genuine low-resource situations, the trials are performed under 5-shot and 10-shot settings per language. A 40 GB VRAM NVIDIA A100 GPU is used for all tests, which is more than enough for large-scale multilingual models. Optimizing transformer-based models has been shown to be beneficial, we employ the Adam optimizer with Learning Rate $1e5$ in Figure 2

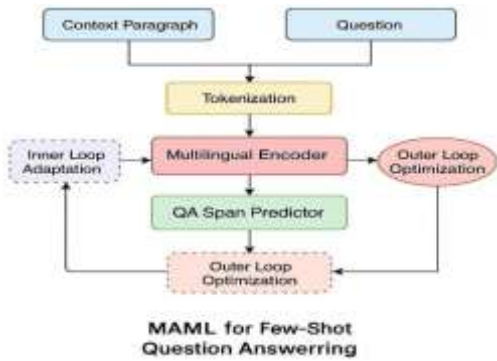


Fig. 2. AML-Based Framework for Adapting QA Models to New Languages using Few-Shot Question-Answer Pairs.

IV. RESULTS

Table I presents the performance comparison of the proposed MAML-based framework against fine-tuning

TABLE I. FEW-SHOT QA MODEL PERFORMANCE (5-SHOT SETTING)

Model	Hindi EM	Hindi F1	Bengali EM	Bengali F1	Telugu EM	Telugu F1	Avg F1
Fine-Tuned mBERT	42.3	58.1	39.8	55.2	36.7	55.6	55.6
Fine-Tuned XLM-R	44.9	60.2	41.1	56.7	37.5	57.2	57.2
MAML + mBERT (Proposed)	51.8	66.9	47.3	63.5	43.9	63.6	63.6
MAML + XLM-R (Proposed)	53.1	68.7	48.6	64.8	45.2	65.0	65.0

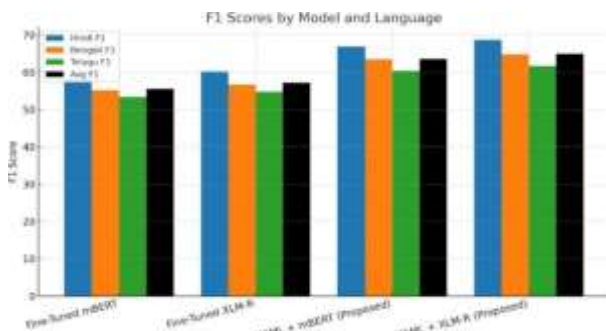


Fig. 3. Comparison F1 Score by Model and Language.

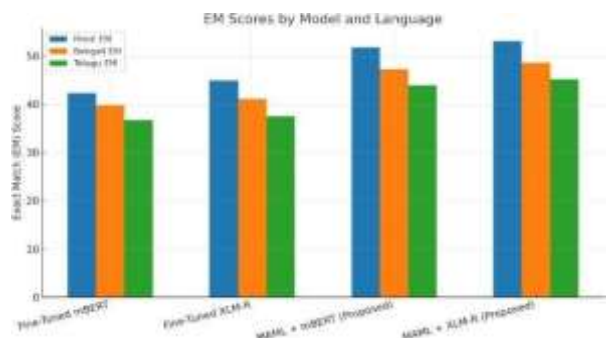


Fig. 4. Comparison EM Score by Model and Language

baselines across Hindi, Bengali, and Telugu in the few-shot setting. The results show that MAML consistently outperforms fine-tuned mBERT, XLM-RoBERTa, and mT5 across both evaluation metrics. On average, MAML improves F1 scores by 6–8% over fine-tuned baselines, confirming its effectiveness in low-resource QA scenarios. Table I: Few-shot QA performance (5-shot setting) across Hindi, Bengali, and Telugu. MAML-based approaches achieve consistent gains in both Exact Match (EM) and F1 Score compared to fine-tuned baselines. Figure 3 and Figure 4 illustrate the relative improvements in F1 Score and Exact Match (EM) respectively, highlighting the consistent advantage of MAML across languages. As shown in Figure 3, MAML-based models deliver the highest F1 scores in all three languages, with MAML + XLM-R reaching 68.7 for Hindi,

64.8 for Bengali, and 61.7 for Telugu. Similarly, Figure 4 shows notable gains in EM, with Hindi improving to 53.1, Bengali to 48.6, and Telugu to 45.2. The key takeaways from these experiments can be summarized as follows. First, MAML consistently boosts performance across all languages, with the largest improvements observed in Telugu, the lowest-resource language in the set. Second, XLM-R variants consistently outperform mBERT, demonstrating the advantage of stronger multilingual pretraining. Third, the MAML + XLM-R model provides the most balanced and highest overall performance, achieving an average F1 of 65.0. Finally, MAML enables faster adaptation in few-shot settings, particularly in morphologically rich languages, although we observe a slight reduction in zero-shot performance, suggesting the need for future work on integrating unsupervised or semi-supervised adaptation.

V. CONCLUSION AND FUTURE WORK

Using low-resource languages like Bengali, Telugu, and Hindi as examples, this study shows that an MAML-based question answering system may greatly improve few-shot performance. Using meta-learning, the suggested method overcomes the drawbacks of conventional fine-tuning and transfer learning by learning an initialization that allows fast adaptation to different language contexts using a small number of defined instances. The meta-learning technique has great promise for developing flexible and scalable quality assurance systems for languages that are underrepresented in the field, as shown by the constant gains in Exact Match and F1 scores seen in experimental findings on TyDiQA and XQuAD. Potentially fruitful avenues for

further development of this study are highlighted below. Adding generative QA to the framework is one possibility; to process free-form answers other than extractive spans, one might use models like mT5 or GPT-style decoders. One potential alternative is to investigate semi-supervised or unsupervised meta-learning methods that may make use of massive amounts of unlabelled text in languages with limited resources, thereby decreasing the need for annotated corpora. When combined with multimodal question answering, which combines textual, visual, and voice modalities, this technique

provides a compelling way to build QA systems that are more resilient and adaptable, even in limited resource environments

REFERENCES

- [1] S. P. Kalava, "Building Trust in AI: Ethical Principles for Transparent Autonomous Systems," *J. Artif. Intell. Mach. Learn. Sata Sci.*, vol. 2, no. 2, 2024.
- [2] A. Graves, M. Liwicki, S. Fernandez, R. Bertolami, H. Bunke, and J. Schmidhuber, "A Novel Connectionist System for Unconstrained Handwriting Recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 5, pp. 855–868, May 2009, doi: 10.1109/TPAMI.2008.137.
- [3] V. Verma, "LSTM-based to Predicting Stock Market Trends Using Machine Learning and Sentiment Analysis," *Int. J. Res. Anal. Rev.*, vol. 8, no. 3, pp. 102–111, 2021.
- [4] K. K. Appari and D. S. Pothalaiah, *A Beginner's Adventure into Generative AI*. JEC PUBLICATION, 2025.
- [5] N. Prajapati, "The Role of Machine Learning in Big Data Analytics: Tools, Techniques, and Applications," *ESP J. Eng. Technol. Adv.*, vol. 5, no. 2, pp. 16–22, 2025, doi: 10.56472/25832646/JETA-V5I2P103.
- [6] S. K. Chintagunta, "Generative AI Approaches to Automated Unit Test Case Generation in Large-Scale Software Projects," *ESP J. Eng. Technol. Adv.*, vol. 4, no. 1, pp. 150–157, 2024, doi: 10.56472/25832646/JETA-V4I1P121.
- [7] S. A. Pushkala, "Generative AI in battling Fraud," in *2024 IEEE 4th International Conference on ICT in Business Industry & Government (ICTBIG)*, IEEE, Dec. 2024, pp. 1–5. doi: 10.1109/ICTBIG64922.2024.10911802.
- [8] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," 2017.
- [9] N. Kolli, J. W. Sajja, and A. Nerella, "Building Secure AI Agents for Autonomous Data Access in Compliance/Regulatory-Critical Environments," *Comput. Fraud Secur.*, vol. 2024, no. 9, pp. 363–373, Sep. 2024, doi: 10.52710/cfs.746.
- [10] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, "Human-level concept learning through probabilistic program induction," *Science (80-.)*, vol. 350, no. 6266, pp. 1332–1338, Dec. 2015, doi: 10.1126/science.aab3050.
- [11] S. Singamsetty, "Accelerating data engineering efficiency with self-learning AI algorithms," *Int. J. Comput. Artif. Intell.*, vol. 6, no. 1, pp. 195–199, Jan. 2025, doi: 10.33545/27076571.2025.v6.i1c.154.
- [12] V. Rajendran, D. Besiahgari, S. C. Patil, M. Chandrashekaraiiah, and V. Challagulla, "A Multi-Agent LLM Environment for Software Design and Refactoring: A Conceptual Framework," in *SoutheastCon 2025*, IEEE, Mar. 2025, pp. 488–493. doi: 10.1109/SoutheastCon56624.2025.10971563.
- [13] C. Finn, P. Abbeel, and S. Levine, "Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks," 2017.
- [14] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a Few Examples: A Survey on Few-shot Learning," *ACM Comput. Surv.*, vol. 53, no. 3, Jun. 2020, doi: 10.1145/3386252.
- [15] J. Baek *et al.*, "What Is Wrong With Scene Text Recognition Model Comparisons? Dataset and Model Analysis," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 4714–4722. doi: 10.1109/ICCV.2019.00481.
- [16] K. M. R. Seetharaman, "A Study of Artificial Intelligence Methods for Analysis of Social Media Sentiment on Texts Classification," in *The 16th International IEEE Conference on Computing, Communication and Networking Technologies (ICCCNT)*, 2025.
- [17] S. Dodda, H. Volikatla, and J. R. Vummadi, "Exploring the Role of AI-Enhanced Chatbots in Automating Recruitment Processes in Human Capital Management Systems," *Int. J. Emerg. Trends Comput. Sci. Inf. Technol.*, vol. 6, no. 3, pp. 28–36, 2025, doi: 10.63282/3050-9246.IJETCSIT-V6I3P104.
- [18] T. Bluche, J. Louradour, and R. Messina, "Scan, attend and read: End-to-end handwritten paragraph recognition with mdlstm attention," in *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, 2017, pp. 1050–1055.
- [19] Z. Zhang, P. Luo, C. C. Loy, and X. Tang, "Learning Deep Representation for Face Alignment with Auxiliary Attributes," *IEEE*, pp. 1–14, 2015.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016. doi: 10.1109/CVPR.2016.90.