

H-AES: Towards Automated Essay Scoring for Hindi

Shubhankar Singh¹, Anirudh Pupneja², Shivaansh Mital³, Cheril Shah⁴, Manish Bawkar⁵,
Lakshman Prasad Gupta⁶, Ajit Kumar³, Yaman Kumar³, Rushali Gupta⁷, Rajiv Ratn Shah³

¹ Manipal University Jaipur,

² BITS Pilani, K K Birla Goa Campus,

³ Indraprastha Institute of Information Technology, Delhi,

⁴ Pune Institute of Computer Technology,

⁵ Sardar Vallabhbhai National Institute of Technology, Surat,

⁶ University of Allahabad,

⁷ Banaras Hindu University

Abstract

The use of Natural Language Processing (NLP) for Automated Essay Scoring (AES) has been well explored in the English language, with benchmark models exhibiting performance comparable to human scorers. However, AES in Hindi and other low-resource languages remains unexplored. In this study, we reproduce and compare state-of-the-art methods for AES in the Hindi domain. We employ classical feature-based Machine Learning (ML) and advanced end-to-end models, including LSTM Networks and Fine-Tuned Transformer Architecture, in our approach and derive results comparable to those in the English language domain. Hindi being a low-resource language, lacks a dedicated essay-scoring corpus. We train and evaluate our models using translated English essays and empirically measure their performance on our own small-scale, real-world Hindi corpus. We follow this up with an in-depth analysis discussing prompt-specific behavior of different language models implemented.

1 Introduction

Academic assessments have long used short-text response and essay writing tasks, which have resulted in numerous scoring approaches, prompt design, and assessment methods. These tasks help judge many aspects of a student’s language learning abilities and are commonly integrated with curricula and standardized tests worldwide in multiple languages. Automated essay scoring (AES) emulates human judgment when evaluating the quality of these written essays. Traditional essay scoring methods entail a vast corpus of written data that human scorers manually evaluate. Manual scoring is a laborious task which is challenging to scale much beyond the limited classroom settings (Kumar et al. 2019; Zhang 2013). AES utilizes Natural Language Processing (NLP) and Machine Learning (ML) techniques to evaluate these essays in a more efficient and scalable manner. Commonly used in standardized tests like the GRE and the TOEFL (Attali and Burstein 2006), many organizations and education councils have turned to use AES to reduce workload (Singla et al. 2022).

Commonly, AES is facilitated by a vast set of training data which are scored using expert-designed evaluation rubrics

(Weinberger et al. 2011). The scores assigned can be broadly categorized into two types: *Holistic scores* - a single discrete score value from a given whole number range and *Trait-based scores* - multiple score values assigned based on multi-dimensional criteria such as relevance to prompt, argument quality, and coherence (Ke and Ng 2019; Bamdev et al. 2022). A majority of work done on the task has entailed holistic scoring (Ke and Ng 2019), leveraging the Automated Student Assessment Prize (ASAP) corpus¹, released on Kaggle in 2012. Eight essay prompts are included in the ASAP corpus covering a broad range of topics. Over time, it has become a widely used corpus for holistic scoring. Most approaches to the AES task are specific to each prompt.

Although short-text response writing is a standard assessment task for students worldwide in different languages, most research has focused on the English language domain, with relatively few focusing on other languages and multi-lingual approaches. Research in AES for Indic languages is negligible. Hindi is the most spoken language in India, with around 528 million native speakers, according to the 2011 language census (INDIA 2011). With millions in the country writing their school-level, graduation, and public-sector examinations in Hindi, there is a dire need to automate the scoring processes. In addition, with the rapidly expanding middle class in India, the number of Hindi telecallers is also increasing at an equally rapid pace. New-age unicorn startups like Apna which hire blue and white collar telecallers for the South Asian market, have instituted automatic scoring as the first filtering step. On the other hand, Hindi NLP methods for automatic scoring are still in the early development stages. The Hindi NLP space’s predicament is a veritable lack of data resources. An approach to address this issue is to use English data translated to Hindi, which has been commonly used to train large language models (Kakwani et al. 2020; Conneau et al. 2020; Khanuja et al. 2021). The need for feature engineering is removed by neural approaches which discover both simple and complex features on their own.

Since a dedicated corpus for AES in Hindi currently does not exist, we use a machine-translated version of the ASAP corpus to train, validate and test our approaches. Our study

¹<https://www.kaggle.com/c/asap-aes>

implements recent approaches used for English AES to the Hindi language domain. Further, to validate that despite translation, models trained on the translated data perform adequately in the natural Hindi language settings, we collect a small-scale natural Hindi corpus and test all our models on this corpus. We employ both classical feature-based models and advanced end-to-end language models based on LSTMs and transformer architectures.

The following three sections explain the related material, the corpus used, and the methodologies employed to perform the task on different models. In Section-5, we explain our experimental procedures, provide empirical results and comparisons with relevant English AES benchmark results, and follow it up with a detailed discussion. Section-6 concludes this study and discusses the limitations and future directions that can be taken to build more accurate and robust AES systems in Hindi and other languages. The dataset, rubric and the source code are made publicly available².

2 Related Work

Research for AES in English has spanned decades (Shermis and Burstein 2013; Zupanc and Bosnic 2016; Ke and Ng 2019). Many studies have treated AES as a regression and text classification problem, while a few apply ranking-based approaches (Bamdev et al. 2020, 2022).

Classical machine learning techniques like linear regression (Miltsakaki and Kukich 2004) support vector regression (Persing, Davis, and Ng 2010; Persing and Ng 2015; Cozma, Butnaru, and Ionescu 2018), and sequential minimal optimization (SMO) (Vajjala 2018) are used typically for regression-based AES. SMO (Vajjala 2018), logistic regression (Nguyen and Litman 2018) and Bayesian network classification (Rudner and Liang 2002) utilize classification approaches. For ranking, SVM ranking (Yannakoudakis, Briscoe, and Medlock 2011) and LambdaMART (Chen and He 2013) have been used. These approaches use custom linguistic features such as errors in grammar (Attali and Burstein 2006), readability features (Zesch, Wojatzki, and Scholten-Akoun 2015), length and syntactic features etc. These techniques are widely popular and have been used in the production of evaluation systems, like E-Rater (Attali and Burstein 2006) used in high-stake examinations like GRE and TOEFL.

Progress in Deep Neural Networks, like Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Long Short Term Memory networks (LSTMs), allow for better performance in AES systems (Taghipour and Ng 2016; Alikaniotis, Yannakoudakis, and Rei 2016; Dong, Zhang, and Yang 2017; Wang et al. 2018; Tay et al. 2018; Song et al. 2020a; Mathias and Bhattacharyya 2020). RNNs and LSTMs are typically used for language processing due to their ability to process sequential data. LSTMs have demonstrated better performance than RNNs for longer sequences due to their ability to retain long-term dependencies. These methods do not require any handcrafted components, leading to their widespread adoption. In the case of

end-to-end Deep Neural Networks, complex features are automatically discovered, saving the time spent on designing them manually. Out of these models, SKIPFLOW (Tay et al. 2018) is able to compete with state-of-the-art by capturing coherence, flow, and semantic relatedness over time, which the authors call neural coherence features.

Lately, pre-trained language models, built using transformer architecture, such as GPT, BERT (Devlin et al. 2018), and XLNet (Yang et al. 2019) have pushed forward the limits of language understanding and generation attributed to their excellent generalization and representation abilities. These models have achieved better results in text classification and regression tasks. Although initial studies using transformers for AES (Rodriguez, Jafari, and Ormerod 2019; Uto, Xie, and Ueno 2020; Mayfield and Black 2020) fail to demonstrate a significant advantage over other deep learning approaches, work from (Cao et al. 2020; Yang et al. 2020; Wang et al. 2022) leverage such pre-trained language models, particularly BERT, and have given state-of-the-art results outperforming various traditional and deep-learning architectures designed for AES. Cao et al. (2020) proposes two self-supervised tasks and a domain adversarial training technique for training optimization. This is the first work that uses pre-trained language models to outperform LSTM-based methods significantly. R²BERT (Yang et al. 2020) combines regression and ranking to fine-tune BERT and obtain the new state-of-the-art. Wang et al. (2022) developed a novel multi-scale essay representation approach based on BERT, employing multiple losses and transfer learning.

Attempts have been made in the past to develop AES for languages other than English. These include research in AES for Chinese (Schultz 2013; Song et al. 2020b), Arabic (Alghamdi et al. 2014; Azmi, Al-Jouie, and Hussain 2019), Japanese (Hirao et al. 2020), Swedish (Östling et al. 2013a,b), German (Zesch, Wojatzki, and Scholten-Akoun 2015; Ludwig et al. 2021), Portuguese (Amorim and Veloso 2017; Fonseca et al. 2018) and more. The lack of dedicated essay corpora as comprehensive as the ones in existence for English is a common feature among most of these studies. While few studies mentioned have explicitly created their organic corpora to work on, many have devised alternate solutions such as scraping the internet for essays, distilling essays from articles and combining datasets used for other tasks.

The progress made in language processing for Indic Languages is gradual but on course. A few notable developments include Bhattacharyya (2010), Arora (2020), Kakwani et al. (2020), Ramesh et al. (2021), and more. Desai and Dabhi (2021) presents a comprehensive report on the advancements in Hindi NLP while Harish and Rangan (2020) offers an in-depth survey on regional Indic language processing. Developments in large pre-trained multilingual models like mBERT (Devlin et al. 2018), XLM-RoBERTa (Conneau et al. 2020), DistilMBERT (Sanh et al. 2019), IndicBERT (Kakwani et al. 2020) etc. have included Hindi and a variety of other regional Indic languages. However, as mentioned in the previous section, the lack of annotated data is a challenge that many in the Hindi NLP space are trying to overcome.

²<https://github.com/midas-research/hindi-aes>

3 Corpora

The Automated Student Assessment (ASAP) corpus has benchmarked a variety of state-of-the-art English language models (Ke and Ng 2019) ever since its release. Due to the ASAP corpus’ ubiquity in AES research and its depth in terms of volume and a mix of narrative, expository and source-dependent response prompts, we decided to focus our analysis on a Hindi-translated version of the ASAP dataset itself (ASAP-Hindi). Translation has seen its way through the NLP discourse, especially in the development of multilingual corpora and models. The use of translated data, however, challenges the validity and quality of the corpus itself, with difficulty in determining the quality gain or loss in specific language-related attributes after translation is applied. We verified random subsets from the translated corpus with the help of both skilled bilingual speakers and expert academics. A common observation was that modern neural machine-translation engines have a proclivity to correct spelling mistakes, but the knowledge distilled and a majority of syntactic and semantic features are retained. Prompts 7 and 8 on the ASAP-Hindi dataset have a distinctly atypical scoring nature to the other prompts and real-world scoring rubrics. The first six prompts help generalize results to a range of AES contexts.

We also built our own Hindi language corpus consisting of 126 real-world essays written by students between 18-20 years old. The submissions were taken as an online essay writing competition, consisting of a single prompt with “The essence of Travel in one’s life” as the central theme. The average essay length is 224 words per essay. The writers of these essays are bilingual, with proficiency in both English and Hindi, with Hindi being their first language. Although much smaller in scale, the use of this corpus helps further corroborate our methods and gives a better understanding of how our methods work on real-world data.

The scoring for the responses was in accordance with a comprehensive rubric that we prepared which took into account both subjectivity scores and attribute-based aspects. The subjectivity scores for each essay capture a general idea and take into account the evaluator’s general perception of the essay. Using an attribute-based approach, scores are provided for factors such as length sufficiency, coherence, relevance to prompt, argument quality, and vocabulary, which are more objective in nature. We combine these scores to obtain a final holistic score within a [0-12] range.

This process is in accordance with a variety of previous scoring techniques in both English and AES in other languages. An expert panel of three Hindi academics performed the scoring following the rubric provided. Scorers 1-2, 1-3, and 2-3 have high inter-rater reliability (Cohen 1968) of 0.831, 0.798, and 0.867, respectively. These scores are comparatively higher than the ones calculated for the ASAP dataset. To further reduce the cognitive bias, we average the three scores to obtain the final score. Our final dataset contains responses from the eight prompts of the ASAP-Hindi corpus and our organic prompt.

Figure 1 displays the prompt used for the organic corpus. Its translation in English is as follows: “*Travel takes us out of our comfort zone and inspires us to see, taste and try new*

	Prompt	Essay Count	Score Range
Hindi Translated ASAP Corpus	P1	1783	2-12
	P2	1800	1-6
	P3	1726	0-3
	P4	1772	0-3
	P5	1805	0-4
	P6	1800	0-4
	P7	1569	0-30
	P8	723	0-60
Organic Corpus	Travel	126	0-12

Table 1: Description of the prompts used for Hindi-AES

यात्रा हमें हमारे आराम क्षेत्र से बाहर ले जाती है और हमें नई चीजों को देखने, स्वाद लेने और कोशिश करने के लिए प्रेरित करती है। यह हमें लगातार नए परिवेश के अनुकूल होने और तलाशने, विभिन्न लोगों के साथ जुड़ने, रोमांच को अपनाने और दोस्तों और प्रियजनों के साथ नए और सार्थक अनुभव साझा करने की चुनौती देती है। यात्रा हमें मानवता के बारे में सिखाती है और हमें विभिन्न दृष्टिकोणों और जीवन के तरीकों के लिए सराहना, समझ और सम्मान देती है। यह हमारे जीवन में सकारात्मक बदलाव भी लाती है और हमें जीवित और सक्रिय रखती है।

इन बातों को ध्यान में रखते हुए एक निबंध लिखें कि कैसे यात्रा नए अनुभव लाती है और आपके जीवन पर सकारात्मक प्रभाव डालती है।

Figure 1: Prompt for the organic corpus

things. It challenges us to constantly adapt and explore new environments, connect with different people, embrace adventure, and share new and meaningful experiences with friends and loved ones. Travel teaches us about humanity and gives us an appreciation, understanding, and respect for different perspectives and ways of life. It also brings positive changes in our lives and keeps us alive and active. Keeping these things in mind, write an essay on how travel brings new experiences and positively impacts your life.”

4 Methodology

In our study we employ various methods ranging from classical machine learning approaches employing feature-extraction, to advanced fine-tuned transformer-based pre-trained models. We implement four models for classical regression and classification approaches; Linear Regression, Support Vector Regression (SVR), Random Forest and XGBoost. These leverage features like essay length, average sentence length, average word length, readability scores, semantic overlap, and vocabulary size. For neural network-based architecture, we implement a Bidirectional Long Short-Term Model (BiLSTM), a Convolutional Neural

Network (CNN) Model, a CNN + LSTM coupled with an attention mechanism, and the popular SKIPFLOW model. Lastly, we fine-tune multiple pre-trained multilingual language models. A wide variety of approaches are used to effectively benchmark corresponding approaches from English AES research, in the Hindi language domain.

4.1 Feature Extraction

Our features are largely inspired by the different dimensions of essay quality mentioned in Ke and Ng (2019), readability metrics presented in Sinha et al. (2012) and the attributes presented in the ASAP++ dataset (Mathias and Bhattacharyya 2018). Due to the tendency of neural translation engines to correct spelling, we do not explicitly employ a spelling-based feature in our feature selection. The extracted features are leveraged by our Linear Regression, SVR, Random Forest and XGBoost models. We extract and use the following six features:

- **Essay Length:** The essay’s length is determined by the total number of words it contains.
- **Average Sentence Length:** The average sentence length is calculated by adding up the lengths of all sentences and then dividing them by the number of sentences.
- **Average Word Length:** The average word length is calculated by adding up the lengths of all words and then dividing them by the number of words.
- **Readability Scores:** Sinha et al. (2012) presents this as a readability metric for Hindi and Bangla. Structural features such as Average Sentence Length (ASL), Average Word Length (AWL), Number of PolySyllabic Words (PSW), Number of Jukta-Akshars (JUK), and more are examined. They utilize Spearman’s rank correlation coefficient to analyze these structural features. Our readability scores are calculated using the following formula, based on the results of their regression analysis:

$$-2.34 + 2.14 * AWL + 0.01 * PSW$$

- **Vocabulary and OOV words:** An essay’s score highly depends on its vocabulary size. Here in our experiments, we have simply taken into consideration the count of unique and out-of-vocabulary words (OOVs) in the essay to get the vocabulary size. Machine translation phonetically renders such OOVs into the target language (Hindi). Although if the OOV word is in frequent occurrence, we discard its use.
- **Semantic Overlap and Coherence:** An essay’s score depends on how well it is connected and coherent. To determine this we have used mBERT and the mechanism used in SKIPFLOW for generating neural coherence features (Tay et al. 2018). We calculated the semantic similarity of two sentences (using mBERT) at a distance of four sentences. As a result, we can capture the coherence and semantic overlap of the essay. Finally, after capturing the scores of the sentence pairs we averaged them out.

4.2 Neural Approaches

Deep neural networks like Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Long Short-Term Memory networks (LSTMs) have facilitated the development of end-to-end machine learning workflows which do not rely on handcrafted features. We implement four popular approaches for AES using neural methods: BiLSTM, CNN, CNN + LSTM + Attention Mechanism, and SKIPFLOW. The training pipeline consists of an initial padding process for every sentence to equalize lengths, followed by a tokenization process using the IndicNLP Tokenizer. Word embeddings for the tokenized sentences were obtained using fastText(Wiki) pre-trained word vectors for Hindi (Bojanowski et al. 2017). These word embeddings are passed to the following model architectures:

- **BiLSTM:** A Bidirectional LSTM, or BiLSTM, is a sequence processing model that consists of two LSTMs processing the sequence in both forward and backward directions. We use a BiLSTM over a normal LSTM as essays are dynamically and artistically written literary devices including references to both previous and forthcoming statements. Word embeddings are passed to a model with three BiLSTM layers and two dense layers to obtain the final score. We use Rectified Linear Unit (ReLU) activation functions between hidden layers and apply Batch Normalization after each BiLSTM layer. The first dense layer is followed by a Dropout layer.
- **CNN:** When looked at from a different perspective, short-text responses and essays have many literary intricacies and relations that are more concentrated in short windows. Thus, using a CNN seems intuitive in capturing these short dependencies over a fixed window size. We use a CNN with three 1-Dimensional convolutional hidden layers separated by pooling layers, Swish (Ramachandran, Zoph, and Le 2017) non-linearities, and Batch Normalisation layers.
- **CNN + LSTM + Attention:** We use the model proposed by Dong, Zhang, and Yang (2017) to score the essays based on sentence representations using a model with CNN, LSTM, and Attention layers. Initially, a 1-Dimensional convolutional layer is used to extract features from a text in a way that is intuitively similar to window-based feature extraction. LSTM layers are used after that to process sequences. Finally, an attention layer is used to pool the outputs.
- **SKIPFLOW:** We implement the SKIPFLOW model introduced by Tay et al. (2018). SKIPFLOW proposes a novel method to calculate the textual coherence in the essays by modeling the relationships between snapshots of the hidden representations of a long short-term memory (LSTM) network as it reads. Subsequently, the semantic relationships between multiple snapshots are used as auxiliary features for prediction.

4.3 Fine-Tuned Transformers

A transformer is a deep neural architecture that uses an attention mechanism to handle sequences of ordered

data. Transformers include several layers each containing a multi-head self-attention network and a position-wise feed-forward network (Vaswani et al. 2017). Unlike LSTMs, Transformers use non-sequential processing, they do not process words in series but rather as a whole, allowing Transformers to take lesser time steps to process input in comparison to LSTMs and RNNs. Transformers also do not suffer from long-term dependency problems credit to their self-attention mechanism.

Multilingual language models use BERT-based architecture and its variants and replace the corpus being used for the task of pre-training in English with the one containing text from multiple languages. Along with this, a few multilingual language models also introduce an additional pre-training task, for instance, Translation Language Modeling (TLM) in the case of XLM-R and MuRIL. IndicBERT and MuRIL are the two prominent Multilingual Language Models which specifically focus on regional Indic languages.

We fine-tune prominent pre-trained multilingual transformer language models: Multilingual BERT (mBERT), DistilmBERT, XLM-Roberta, MuRIL, and IndicBERT. The tokenized prompt and essay are fed to the transformer model as input. The input is tokenized with a [CLS] token at the beginning. [SEP] tokens are added at the end of both the prompt and the essay to differentiate them. The prompt and the essay sequences are fed to the transformer encoder and become hidden layer sequences. These sequences are passed to a simple feed-forward network, consisting of two hidden layers, attached to the base transformer architecture and a final numerical score is obtained.

5 Experiments

This section describes the experimental procedure, evaluation metric, empirical results, and comparison with published results on prominent English AES models.

5.1 Experimental Setup

We conduct prompt-specific experimentation in accordance with Taghipour and Ng (2016). While it might seem ideal to train prompts together, it’s important to note that each prompt might include genres that are in sharp contrast to one another, such as narrative or argumentative essays. Prompts can also be scored differently according to the level of the students and the scoring rubrics. This makes training prompts together extremely challenging.

For feature-based methods, the text was pre-processed to filter out stopwords, named-entities and mentions (denoted by ‘@’ symbols in the ASAP Dataset). The custom feature scores were normalized to improve the stability of the model. The neural LSTM and CNN models were all trained for 100 epochs with a learning rate of 1e-4. For fine-tuning the pre-trained multilingual transformer models we use the AdamW optimizer (Loshchilov and Hutter 2017), which is a stochastic optimization method that modifies the typical implementation of weight decay in Adam (Kingma and Ba 2014), by decoupling weight decay from the gradient update. We set our learning rate to 5e-5 which decreases linearly to 0. We train all models on a 6-8 epoch range, depending on the base transformer model’s ability to learn.

For our experiments on all prompts (including our organic prompt corpus) we use a 60/20/20 split for train, validation and test sets. Our normalization process keeps all score ranges within [0,1]. To calculate the Quadratic Weighted Kappa (QWK) scores, the scores are re-scaled to the original prompt-specific scale for prediction.

5.2 Evaluation Metric

We use the Quadratic Weighted Kappa (QWK), which is a common measure for evaluating and comparing AES methods (Cohen 1968). A key reason for its prevalence is its particular sensitivity to differences in scores and its ability to take into account chance agreements. QWK score generally ranges from 0 to 1. When the agreement is lower than expected by chance, the score becomes negative. QWK Score is calculated as follows. Initially, a weight matrix W is created in accordance to Equation 1:

$$W_{i,j} = \frac{(i-j)^2}{(N-1)^2} \quad (1)$$

If N is the total number of possible ratings, i and j represent the reference rating (given by a human annotator) and hypothesis rating (awarded by an AES system), The outer product of the reference and hypotheses rating histogram vectors results in an expected count matrix E . This matrix is normalized so that the elements in E and O have the same sum. Finally, given the matrices O and E , the QWK score is calculated according to Equation 2:

$$\kappa = 1 - \frac{\sum_{i,j} w_{i,j} O_{i,j}}{\sum_{i,j} w_{i,j} E_{i,j}} \quad (2)$$

5.3 Results and Comparison

In this section, we present the results of our experimentation. Table 2 (rows 1-13) reports the empirical results (QWK scores) on all our approaches on the ASAP Hindi Dataset and our organic prompt. We have also provided an average for results on the ASAP-Hindi set for a better comparison of the models across prompts. As a baseline for AES in Hindi does not exist, Table 3 (rows 1-4) presents published results on prominent models for AES in English. These results provide a benchmark upper limit for us to compare our scores, across all types of models. According to Tables 2 and 3, all 13 models implemented were able to learn the task and perform competitively when compared to results on the English models.

Results from the feature-based (Table 2, rows 1-4) approaches were also competitive, with Linear Regression and XGBoost outperforming all other models on prompt 1 and prompts 2 and 8 respectively. Averages for the feature-extraction models were not far off from the Ease SVR average benchmark in Table 3, with the lowest and highest averages falling 0.078 and 0.011 points short of the benchmark, respectively.

Although the neural models (Table 2, rows 5-8) did not outperform other models on any prompt, their averages were higher than the feature-based models and lower than the fine-tuned transformers, providing decent transitional results

ID	Hindi AES Models	P1	P2	P3	P4	P5	P6	P7	P8	Average	Organic
1	SVR	0.799	0.612	0.605	0.657	0.797	0.630	0.400	0.380	0.610	0.579
2	Linear Regression	0.800	0.614	0.588	0.624	0.768	0.605	0.680	0.635	0.664	0.681
3	Random Forest	0.705	0.608	0.621	0.685	0.791	0.665	0.652	0.560	0.661	0.762
4	XGBoost	0.794	0.667	0.573	0.676	0.792	0.653	0.713	0.641	0.688	0.827
5	CNN	0.571	0.513	0.529	0.614	0.657	0.703	0.521	0.426	0.566	0.762
6	BiLSTM	0.631	0.517	0.612	0.703	0.643	0.713	0.607	0.443	0.608	0.842
7	CNN + LSTM + Attention	0.723	0.597	0.677	0.711	0.781	0.791	0.701	0.593	0.696	0.827
8	SKIPFLOW LSTM (Tensor)	0.742	0.621	0.695	0.731	0.804	0.777	0.717	0.619	0.713	0.812
9	mBERT	0.683	0.652	0.711	0.775	0.828	0.785	0.781	0.548	0.720	0.852
10	DistilmBERT	0.661	0.592	0.698	0.766	0.825	0.793	0.785	0.596	0.714	0.784
11	XLm-RoBERTa	0.758	0.585	0.692	0.809	0.834	0.822	0.794	0.639	0.741*	0.831
12	MuRIL	0.620	0.412	0.528	0.756	0.812	0.713	0.547	0.327	0.589	0.528
13	IndicBERT	0.651	0.489	0.659	0.751	0.799	0.784	0.708	0.412	0.656	0.796

Table 2: Experiment results of all models in terms of QWK on ASAP-Hindi corpus and the organic prompt. The bold number is the best performance for each prompt. The best average QWK is annotated with *.

ID	English AES Models	P1	P2	P3	P4	P5	P6	P7	P8	Average
1	EASE (SVR)	0.781	0.621	0.630	0.749	0.782	0.771	0.727	0.534	0.699
2	CNN + LSTM + Attention	0.822	0.682	0.672	0.814	0.803	0.811	0.801	0.705	0.764
3	SKIPFLOW LSTM(Tensor)	0.832	0.684	0.695	0.788	0.815	0.810	0.800	0.697	0.764
4	R ² BERT	0.817	0.719	0.698	0.845	0.841	0.847	0.839	0.726	0.791

Table 3: Published results on prominent models for AES in English. All results in terms of QWK score on the original ASAP Corpus.

from classical to advanced end-to-end methods. SKIPFLOW (Table 2, row-8) gave the highest average score amongst these, followed closely by the CNN + LSTM + Attention average. The SKIPFLOW average fell 0.051 points short off it’s original implementation in English.

On average, the fine-tuned highly multilingual transformers (Table 2, rows 9-11) gave higher results across all prompts. The fine-tuned mBERT model gives the maximum score for prompt 3. The fine-tuned XLM-R model outperforms all other models on four of the eight ASAP-Hindi prompts (Prompts 4, 5, 6, 7), giving the maximum average QWK as well thereby, establishing a state-of-the-art for AES in Hindi. Compared to R²BERT’s average the fine-tuned XLM-R is 0.050 points short. Results on the mBERT model closely follow the results on the XLM-R model. The fine-tuned Indic transformers (Table 2, rows 12 and 13) did not perform comparably well. While IndicBERT did try to compete and learn, MuRIL’s behavior during the training process was highly inconsistent, resulting in unpredictable results. This behavior could be attributed to a variety of factors which will be discussed in the next sub-section.

Results on the organic prompt were favourable (in comparison to results on the ASAP-Hindi set) with the fine-tuned mBERT model giving the highest QWK score. It is important to note that in contrast to the other prompts, the QWK

scores obtained for organic prompt showed a slightly higher variance during training. It is likely that this is a result of the organic set’s smaller magnitude compared to the ASAP-Hindi Dataset.

5.4 Analysis and Discussion

A general observation that is consistent with both previously established English AES results and our study, is that results on prompts 4,5, and 6 are higher than the other prompts for the ASAP dataset. Prompts requiring source-dependent responses perform better during training as compared to narrative, persuasive or expository prompts possibly due to a general consistency in syntax, coherence, and availability of source material. Such prompts are coupled with more balanced real-world rubrics allowing for consistency all throughout the writing and the scoring process, making such prompts ideal for generalization. In contrast generalization of ideas is slightly more difficult on prompts that allow for persuasive and expository discussions due to variance in human thought and cognition. Prompts with rubrics that are not consistent with real-world rubrics such as prompt-8, give results much worse than the aforementioned source-dependent prompts.

The more discrete nature of feature-based approaches allows for consistent performance, but their failure to under-

stand nuance and interpret content limits their potential for improvement when compared to large pre-trained language models. Given the considerably larger scale (parameters and pre-training data) of XLM-R and mBERT they significantly outperform the other models. Although DistilmBERT is approximately 40% lighter than mBERT, it performs almost equally well and was the fastest to train among all large language models.

A result observed was the relatively poor performance of the Indic language models, especially MuRIL. Several plausible explanations might explain this finding. IndicBERT's performance can be attributed to its base architecture being an ALBERT model which is almost 90% lighter than the BERT-base. IndicBERT may have been disadvantaged by the lack of trainable parameters because of this, resulting in it not competing with mBERT and XLM-R. MuRIL, however, does have the parameter size to compete. Patil et al. (2021) provides compelling evidence as to why MuRIL might fail to perform on syntactically-complex tasks such as AES, it includes the following points: Both IndicBERT and MuRIL perform masked word-level language modeling and do not have a sentence level pre-training task. An important aspect of AES is the morpho-syntactical relationship between ideas and sentences. mBERT, XLM-R, and DistilmBERT are highly multilingual language models pre-trained on more than 100 languages. It may provide them with linguistic and typological generalizations needed to model morpho-syntax more effectively than Indic models, which are trained on a small number of Indic languages and English. Another factor to add to this is that monolingual and multilingual large language models show syntactic localization across their layers which makes them perform better at complex tasks with long-range syntactical dependencies. In comparison, Indic language models IndicBERT and MuRIL show little localization, with MuRIL showing the least localization across all layers amongst all models. The syntactic directness of source-dependent responses is possibly why MuRIL still remains competitive on prompts 4, 5, and 6.

6 Conclusion and Future Work

In this study, we implement and analyze various methods for AES in Hindi ranging from classical feature-based Machine Learning (ML) to advanced end-to-end models to set benchmarks and the state-of-the-art for AES in Hindi. We also introduce a single-prompt corpus of student-written essays in Hindi to further substantiate our findings from the Hindi-Translated ASAP corpus. The results of our experiments shed new light on AES research. We obtain competitive results when compared to the benchmark and the state-of-the-art methods in English AES. We also try to explain and analyze the results obtained using our models, attributing to the idiosyncrasies of the models as well as the nature of the prompts on which they are tested.

In view of the fact that AES in Hindi is particularly unexplored, multiple future research directions are possible. We plan to extend our work by scaling the organic corpus (both in terms of essays per prompt and the type of prompts) and proposing architectures that learn the syntactical features and nuances of Hindi by leveraging the trait-based scores

that are included with the corpus. It is reasonable to use a Hindi-translated version of the ASAP dataset as precedent, but a comprehensive Hindi corpus is essential for AES in Hindi. A larger corpus facilitates more nuanced learning, which enables models to generalize from a wider spectrum of results.

We obtain the most favorable results on our fine-tuned pre-trained multilingual transformer language models and to push these results further, we plan to try different training optimization methods including domain adversarial training, multi-scale essay representation approaches and more. Using such training optimization methods might improve performance on Hindi which is more morpho-syntactically complex than English. For the same reason, introducing linguistic knowledge to segment at a more reasonable scale may bring further improvement. We also hope to push the results on the Indic language models IndicBERT and MuRIL through these optimization strategies. Due to the prominence of regional languages in Indian communities, a multilingual essay evaluation approach will allow for more diversity in essay writing and large-scale examinations.

Acknowledgements

Rajiv Ratn Shah was partly supported by the Infosys Center for Artificial Intelligence and the Center of Design and New Media at IIT Delhi, India.

References

- Alghamdi, M.; Alkanhal, M.; Al-Badrashiny, M.; Al-Qabbany, A.; Areshey, A.; and Alharbi, A. 2014. A hybrid automatic scoring system for Arabic essays. *Ai Communications*, 27(2): 103–111.
- Alikaniotis, D.; Yannakoudakis, H.; and Rei, M. 2016. Automatic Text Scoring Using Neural Networks. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 715–725. Berlin, Germany: Association for Computational Linguistics.
- Amorim, E.; and Veloso, A. 2017. A multi-aspect analysis of automatic essay scoring for Brazilian Portuguese. In *Proceedings of the Student Research Workshop at the 15th Conference of the European Chapter of the Association for Computational Linguistics*, 94–102.
- Arora, G. 2020. iNLTK: Natural language toolkit for indic languages. *arXiv preprint arXiv:2009.12534*.
- Attali, Y.; and Burstein, J. 2006. Automated Essay Scoring With e-rater® V.2. *The Journal of Technology, Learning and Assessment*, 4(3).
- Azmi, A. M.; Al-Jouie, M. F.; and Hussain, M. 2019. AAEE—Automated evaluation of students' essays in Arabic language. *Information Processing & Management*, 56(5): 1736–1752.
- Bamdev, P.; Grover, M. S.; Kumar, Y.; Vafayee, P.; Hama, M.; and Shah, R. R. 2020. Classification, Regression or Ordinal Regression? Assessing oral proficiency of non-native English speakers and interpreting results. WiML, NeurIPS 2020.
- Bamdev, P.; Grover, M. S.; Singla, Y. K.; Vafae, P.; Hama, M.; and Shah, R. R. 2022. Automated Speech Scoring System Under The Lens. *International Journal of Artificial Intelligence in Education*, 1–36.

- Bhattacharyya, P. 2010. IndoWordNet. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valletta, Malta: European Language Resources Association (ELRA).
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, 5: 135–146.
- Cao, Y.; Jin, H.; Wan, X.; and Yu, Z. 2020. Domain-Adaptive Neural Automated Essay Scoring. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, 1011–1020. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380164.
- Chen, H.; and He, B. 2013. Automated essay scoring by maximizing human-machine agreement. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1741–1752.
- Cohen, J. 1968. Weighted kappa: nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4): 213.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8440–8451. Online: Association for Computational Linguistics.
- Cozma, M.; Butnaru, A. M.; and Ionescu, R. T. 2018. Automated essay scoring with string kernels and word embeddings. *arXiv preprint arXiv:1804.07954*.
- Desai, N. P.; and Dabhi, V. K. 2021. Taxonomic survey of Hindi Language NLP systems. *CoRR*, abs/2102.00214.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dong, F.; Zhang, Y.; and Yang, J. 2017. Attention-based Recurrent Convolutional Neural Network for Automatic Essay Scoring. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, 153–162. Vancouver, Canada: Association for Computational Linguistics.
- Fonseca, E.; Medeiros, I.; Kamikawachi, D.; and Bokan, A. 2018. Automatically grading brazilian student essays. In *International Conference on Computational Processing of the Portuguese Language*, 170–179. Springer.
- Harish, B. S.; and Rangan, R. K. 2020. A comprehensive survey on Indian regional language processing. *SN Applied Sciences*, 2: 1–16.
- Hirao, R.; Arai, M.; Shimanaka, H.; Katsumata, S.; and Komachi, M. 2020. Automated Essay Scoring System for Non-native Japanese Learners. In *Proceedings of the 12th Language Resources and Evaluation Conference*, 1250–1257.
- INDIA, P. 2011. *Census of India 2011 Provisional Population Totals*. Office of the Registrar General and Census Commissioner New Delhi.
- Kakwani, D.; Kunchukuttan, A.; Golla, S.; Gokul, N.; Bhattacharyya, A.; Khapra, M. M.; and Kumar, P. 2020. IndicNLP Suite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 4948–4961.
- Ke, Z.; and Ng, V. 2019. Automated Essay Scoring: A Survey of the State of the Art. In *IJCAI*, volume 19, 6300–6308.
- Khanuja, S.; Bansal, D.; Mehtani, S.; Khosla, S.; Dey, A.; Gopalan, B.; Margam, D. K.; Aggarwal, P.; Nagipogu, R. T.; Dave, S.; et al. 2021. Muril: Multilingual representations for indian languages. *arXiv preprint arXiv:2103.10730*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kumar, Y.; Aggarwal, S.; Mahata, D.; Shah, R. R.; Kumaraguru, P.; and Zimmermann, R. 2019. Get it scored using autosas—an automated system for scoring short answers. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 9662–9669.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ludwig, S.; Mayer, C.; Hansen, C.; Eilers, K.; and Brandt, S. 2021. Automated essay scoring using transformer models. *Psych*, 3(4): 897–915.
- Mathias, S.; and Bhattacharyya, P. 2018. ASAP++: Enriching the ASAP Automated Essay Grading Dataset with Essay Attribute Scores. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki, Japan: European Language Resources Association (ELRA).
- Mathias, S.; and Bhattacharyya, P. 2020. Can Neural Networks Automatically Score Essay Traits? In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, 85–91.
- Mayfield, E.; and Black, A. W. 2020. Should you fine-tune BERT for automated essay scoring? In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, 151–162.
- Miltsakaki, E.; and Kukich, K. 2004. Evaluation of text coherence for electronic essay scoring systems. *Natural Language Engineering*, 10(1): 25–55.
- Nguyen, H.; and Litman, D. 2018. Argument mining for improving the automated scoring of persuasive essays. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Östling, R.; Smolentzov, A.; Tyrefors Hinnerich, B.; and Höglin, E. 2013a. Automated Essay Scoring for Swedish. In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, 42–47. Atlanta, Georgia: Association for Computational Linguistics.
- Östling, R.; Smolentzov, A.; Tyrefors Hinnerich, B.; and Höglin, E. 2013b. Automated Essay Scoring for Swedish. In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, 42–47. Atlanta, Georgia: Association for Computational Linguistics.
- Patil, R.; Dhillon, J.; Mahurkar, S.; Kulkarni, S.; Malhotra, M.; and Baths, V. 2021. Vyākaraṇa: A Colorless Green Benchmark for Syntactic Evaluation in Indic Languages. In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, 153–165. Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Persing, I.; Davis, A.; and Ng, V. 2010. Modeling organization in student essays. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, 229–239.
- Persing, I.; and Ng, V. 2015. Modeling argument strength in student essays. In *Proceedings of the 53rd Annual Meeting of*

- the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 543–552.
- Ramachandran, P.; Zoph, B.; and Le, Q. V. 2017. Searching for activation functions. *arXiv preprint arXiv:1710.05941*.
- Ramesh, G.; Doddapaneni, S.; Bheemaraj, A.; Jobanputra, M.; AK, R.; Sharma, A.; Sahoo, S.; Diddee, H.; J, M.; Kakwani, D.; Kumar, N.; Pradeep, A.; Nagaraj, S.; Deepak, K.; Raghavan, V.; Kunchukuttan, A.; Kumar, P.; and Khapra, M. S. 2021. Samanantar: The Largest Publicly Available Parallel Corpora Collection for 11 Indic Languages. *arXiv:2104.05596*.
- Rodriguez, P. U.; Jafari, A.; and Ormerod, C. M. 2019. Language models and automated essay scoring. *arXiv preprint arXiv:1909.09482*.
- Rudner, L. M.; and Liang, T. 2002. Automated essay scoring using Bayes' theorem. *The Journal of Technology, Learning and Assessment*, 1(2).
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108.
- Schultz, M. T. 2013. The intellimetric automated essay scoring engine—a review and an application to chinese essay scoring. *Handbook of automated essay scoring: Current applications and future directions*, 89–98.
- Shermis, M. D.; and Burstein, J. 2013. *Handbook of automated essay evaluation*. NY: Routledge.
- Singla, Y. K.; Krishna, S.; Shah, R. R.; and Chen, C. 2022. Using sampling to estimate and improve performance of automated scoring systems with guarantees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 12835–12843.
- Sinha, M.; Sharma, S.; Dasgupta, T.; and Anupam, B. 2012. New Readability Measures for Bangla and Hindi Texts. 1141–1150.
- Song, W.; Song, Z.; Liu, L.; and Fu, R. 2020a. Hierarchical Multi-task Learning for Organization Evaluation of Argumentative Student Essays. In *IJCAI*, 3875–3881.
- Song, W.; Zhang, K.; Fu, R.; Liu, L.; Liu, T.; and Cheng, M. 2020b. Multi-stage pre-training for automated Chinese essay scoring. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 6723–6733.
- Taghipour, K.; and Ng, H. T. 2016. A neural approach to automated essay scoring. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, 1882–1891.
- Tay, Y.; Phan, M.; Tuan, L. A.; and Hui, S. C. 2018. Skipflow: Incorporating neural coherence features for end-to-end automatic text scoring. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Uto, M.; Xie, Y.; and Ueno, M. 2020. Neural automated essay scoring incorporating handcrafted features. In *Proceedings of the 28th International Conference on Computational Linguistics*, 6077–6088.
- Vajjala, S. 2018. Automated assessment of non-native learner essays: Investigating the role of linguistic features. *International Journal of Artificial Intelligence in Education*, 28(1): 79–105.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wang, Y.; Wang, C.; Li, R.; and Lin, H. 2022. On the Use of Bert for Automated Essay Scoring: Joint Learning of Multi-Scale Essay Representation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3416–3425. Seattle, United States: Association for Computational Linguistics.
- Wang, Y.; Wei, Z.; Zhou, Y.; and Huang, X.-J. 2018. Automatic essay scoring incorporating rating schema via reinforcement learning. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, 791–797.
- Weinberger, A.; Dreher, H.; Al-Smadi, M.; and Guetl, C. 2011. Analytical assessment rubrics to facilitate semi-automated Essay grading and feedback provision. In *Proceedings of the Australian Technology Network Assessment Conference 2011*, 170–177. Curtin University.
- Yang, R.; Cao, J.; Wen, Z.; Wu, Y.; and He, X. 2020. Enhancing automated essay scoring performance via fine-tuning pre-trained language models with combination of regression and ranking. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 1560–1569.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R. R.; and Le, Q. V. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Yannakoudakis, H.; Briscoe, T.; and Medlock, B. 2011. A new dataset and method for automatically grading ESOL texts. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, 180–189.
- Zesch, T.; Wojatzki, M.; and Scholten-Akoun, D. 2015. Task-independent features for automated essay grading. In *Proceedings of the tenth workshop on innovative use of NLP for building educational applications*, 224–232.
- Zhang, M. 2013. Contrasting automated and human scoring of essays. *R & D Connections*, 21(2): 1–11.
- Zupanc, K.; and Bosnic, Z. 2016. Advances in the field of automated essay evaluation. *Informatica*, 39(4).