

MedHastra@DravidianLangTech 2026: Piecewise Style Classification for Telugu Prompt Recovery Using XLM-RoBERTa

Shruti Chandrasekar and Vedajanaani R S and P. Vijayalakshmi

SSN College of Engineering, Chennai, India

{shruti2210139, vedajanaani2310594, vijayalakshmip}@ssn.edu.in

Abstract

We present a system for the DravidianLangTech @ ACL 2026 shared task on Telugu Prompt-Style Recovery (B et al., 2026). The task requires classifying Telugu text into one of nine communicative styles: Formal, Informal, Optimistic, Pessimistic, Humorous, Serious, Inspiring, Authoritative and Persuasive. Our approach fine-tunes the multilingual XLM-RoBERTa base model with a piecewise segment comparison strategy that evaluates distinct stylistic markers across sentence segments, enabling richer contextual discrimination between visually similar styles. Evaluated on the official test set, our system achieves a Macro F1 score of 0.1205, Accuracy of 0.1196, Precision of 0.1205 and Recall of 0.1231. We analyze the challenges of stylistic ambiguity in low-resource Telugu NLP and discuss directions for future improvement.

1 Introduction

Identifying the communicative style of a text is a fundamental yet challenging problem in natural language understanding. Beyond lexical meaning, style encodes pragmatic intent, rhetorical tone, affective stance and social register, all of which are crucial for building language technologies that interact naturally with users (Biber, 1988; Pang and Lee, 2008). Prior work on stylistic variation and style transfer has demonstrated that modeling style requires capturing subtle linguistic and discourse-level cues rather than relying solely on surface sentiment signals (Jhamtani et al., 2017; Pan et al., 2026).

For morphologically rich and agglutinative languages such as Telugu, stylistic analysis presents additional complexity. Telugu exhibits extensive inflectional morphology and flexible word order, which can obscure surface-level stylistic markers. Furthermore, the scarcity of labeled corpora and the prevalence of code-mixing with English in so-

cial media introduce additional variability (Kusampudi et al., 2021; Chakravarthi et al., 2022). These characteristics make nuanced style detection significantly more challenging than in well-resourced languages such as English.

The DravidianLangTech @ ACL 2026 shared task on Prompt Recovery for LLM in Telugu addresses the problem of identifying the communicative style of Telugu text generated under style-specific prompts. The dataset contains social media excerpts annotated across nine stylistic categories: Formal, Informal, Optimistic, Pessimistic, Humorous, Serious, Inspiring, Authoritative and Persuasive. Unlike traditional sentiment classification, these categories frequently overlap in semantic and pragmatic dimensions. For instance, Inspiring and Optimistic texts both exhibit positive affect, while Authoritative and Formal styles share structural rigidity and controlled lexical choice. Similarly, Serious and Pessimistic texts may overlap in tonal gravity despite differing in communicative intent. Such inter-class proximity necessitates modeling approaches capable of capturing fine-grained contextual signals.

In this work, we present our submission to the shared task. We fine-tune XLM-RoBERTa (Conneau et al., 2020), a multilingual transformer model pretrained on diverse languages, for nine-way style classification. To enhance sensitivity to localized stylistic cues, we introduce a piecewise segmentation strategy that decomposes each input into shorter contextual units and aggregates their representations. Our design is motivated by prior work showing that transformer architectures can benefit from modified attention mechanisms and segment-level modeling when handling longer or discourse-sensitive texts (Beltagy et al., 2020). Our analysis includes per-class performance evaluation and detailed error analysis to better understand stylistic confusion patterns in low-resource Telugu text.

2 Related Work

2.1 Computational Modeling of Style

Computational approaches to stylistic analysis have evolved from rule-based and feature-engineered systems to deep neural architectures. Early work in stylistic classification focused on authorship attribution and genre identification using surface-level lexical and syntactic features (Stamatatos, 2009).

With the advent of neural models, style modeling has been reframed as a representation learning problem, where latent stylistic properties are captured through distributed embeddings. Neural approaches to non-parallel style transfer and controlled text generation have emphasized disentangling content from stylistic attributes (Shen et al., 2017; Hu et al., 2017). These works highlight the inherent difficulty of separating semantics from stylistic variation.

More recently, transformer-based encoders such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) have demonstrated strong performance across a wide range of text classification tasks. Their contextualized representations enable modeling subtle pragmatic and rhetorical cues, making them particularly suitable for fine-grained stylistic classification. However, distinguishing closely related stylistic categories remains challenging due to overlapping semantic signals.

2.2 Emotion and Pragmatic Signal Modeling

Style classification intersects with emotion recognition and pragmatic intent detection. Emotion classification datasets such as GoEmotions (Demszky et al., 2020) demonstrate that fine-grained affective categories often exhibit substantial inter-class confusion. Similarly, research on politeness and formality detection (Danescu-Niculescu-Mizil et al., 2013; Pavlick and Tetreault, 2016) shows that pragmatic signals are highly context-dependent and socially grounded. These findings suggest that effective style modeling requires sensitivity to discourse-level and contextual cues rather than isolated lexical indicators.

2.3 Low-Resource and Indian Language NLP

For Indian languages, recent efforts have focused on developing large-scale corpora and pretrained multilingual models. IndicNLP Suite (Kakwani et al., 2020) provides monolingual corpora, evaluation benchmarks and pretrained language models for multiple Indian languages. Despite these ad-

Split	# Instances	# Classes
Train	3000	9
Dev	300	9
Test	301	9

Table 1: Dataset split statistics for the Telugu prompt-style recovery task.

vances, fine-grained stylistic classification remains underexplored for Dravidian languages such as Telugu. The limited availability of annotated stylistic corpora and the prevalence of informal, code-mixed social media text pose additional challenges for robust style modeling.

3 Task Description and Dataset

The DravidianLangTech @ ACL 2026 shared task provides a labeled dataset of Telugu text excerpts annotated with one of nine communicative style categories: Formal, Informal, Optimistic, Pessimistic, Humorous, Serious, Inspiring, Authoritative and Persuasive. Each instance consists of an original text and a corresponding style-transformed version (referred to as the *changed-style* text). The task requires predicting the stylistic label of the transformed text.

In our experiments, the input to the model is the changed-style text and the output is a predicted style label from the predefined nine-class taxonomy.

The label distribution reflects the stylistic taxonomy defined in the task guidelines. Several style categories exhibit semantic proximity, such as Optimistic vs. Inspiring and Formal vs. Authoritative. This inter-class similarity introduces ambiguity and increases classification difficulty, even for pretrained multilingual transformer models.

Additionally, the dataset reflects characteristics of social media language, including informal punctuation, truncated constructions, orthographic variation and occasional code-switching between Telugu and English. These linguistic phenomena further complicate stylistic modeling in a low-resource setting.

4 Methodology

4.1 Backbone Model: XLM-RoBERTa

We adopt xlm-roberta-base (Conneau et al., 2020) as our backbone encoder. XLM-RoBERTa is a multilingual transformer model pretrained using masked language modeling on large-scale CommonCrawl data covering 100 languages. Its

cross-lingual pretraining enables robust contextual representations for morphologically rich and low-resource languages such as Telugu.

Given an input segment s , the encoder produces a contextualized representation $h_{[\text{CLS}]} \in \mathbb{R}^d$ corresponding to the special classification token. We append a task-specific linear classification head to map this representation to nine stylistic categories:

$$\ell = Wh_{[\text{CLS}]} + b \quad (1)$$

where $W \in \mathbb{R}^{9 \times d}$ and $b \in \mathbb{R}^9$ are learnable parameters and ℓ denotes the logits over the nine style classes. The model is trained using the cross-entropy loss function applied to the predicted logits in Equation 1.

4.2 Piecewise Segment Aggregation

We hypothesize that stylistic cues are often localized within specific portions of a text rather than uniformly distributed across the entire sequence. For example, humorous texts may contain localized punchlines, while authoritative texts may begin with imperative constructions. Modeling the input as a single flat sequence may dilute these localized stylistic signals.

To address this, we introduce a piecewise segmentation strategy. Each input text x is decomposed into k non-overlapping segments:

$$x = \{s_1, s_2, \dots, s_k\} \quad (2)$$

Segmentation is performed using punctuation-based heuristics adapted for Telugu script and common social media conventions. Each segment s_i is independently encoded by XLM-RoBERTa to produce a logit vector ℓ_i as defined in Equation 1.

To compute the final prediction, we aggregate segment-level logits using element-wise max pooling:

$$\ell^* = \max(\ell_1, \ell_2, \dots, \ell_k) \quad (3)$$

The predicted label is then obtained as:

$$\hat{y} = \arg \max(\ell^*) \quad (4)$$

The aggregation step in Equation 3 allows the model to prioritize the most stylistically salient segment, thereby increasing sensitivity to localized rhetorical and pragmatic cues. Compared to mean aggregation, max pooling preserves strong segment-level signals that may otherwise be attenuated.

System	F1	Acc	Prec	Rec
XLM-R + Piecewise	0.1205	0.1196	0.1205	0.1231

Table 2: Performance of XLM-R with Piecewise pooling

4.3 Training Procedure

We fine-tune the model using the AdamW optimizer with a learning rate of 2×10^{-5} and a batch size of 16. Training is conducted for up to 10 epochs with early stopping based on macro F1 score on the validation set (patience = 2). Mixed precision (FP16) training is employed when GPU resources are available to improve computational efficiency.

Each segment is truncated or padded to a maximum length of 128 tokens. Hyperparameters are selected based on development set performance. For final evaluation, the model is trained on the combined training and development splits and predictions are generated for the held-out test set.

5 Results and Analysis

5.1 Overall Performance

Table 2 presents the official evaluation metrics of our system on the shared task test set. Our model achieves a Macro F1 score of 0.1205, Accuracy of 0.1196, Precision of 0.1205 and Recall of 0.1231.

The Macro F1 score of 0.1205 is slightly above the random baseline of $1/9 \approx 0.111$, indicating that the model captures limited but non-trivial stylistic signal across classes. The proximity between Macro F1 and Precision suggests relatively balanced per-class behavior, without extreme bias toward particular categories. The slightly higher Recall (0.1231) indicates a marginal tendency toward over-prediction rather than overly conservative classification.

These results highlight the inherent complexity of nine-way stylistic classification in a low-resource setting. Unlike sentiment analysis, stylistic categories in this task are not mutually exclusive in terms of lexical or syntactic patterns, making fine-grained discrimination substantially more challenging.

5.2 Discussion

Several factors likely contribute to the observed performance. First, many of the stylistic categories exhibit substantial semantic and pragmatic overlap. For instance, *Inspiring* and *Optimistic* texts

frequently share positive and forward-looking language, while *Authoritative* and *Formal* styles both employ structured syntax and controlled lexical choice. Such inter-class proximity reduces separability in embedding space.

Second, Telugu presents additional challenges due to rich morphology and flexible word order. Subtle stylistic cues may be expressed through inflectional variation rather than distinct lexical markers, which can be difficult for transformer models to isolate in limited data regimes.

Third, the presence of code-mixing between Telugu and English introduces orthographic and subword segmentation variability. Although XLM-RoBERTa supports multilingual input, inconsistent transliteration and informal punctuation may disrupt coherent representation learning.

5.3 Error Analysis

Qualitative analysis on the development set reveals consistent confusion patterns across semantically related styles. The model most frequently confuses *Inspiring* with *Optimistic*, suggesting difficulty in distinguishing motivational framing from general positive sentiment. Similarly, *Authoritative* and *Formal* are often interchanged due to shared structural characteristics.

Humorous instances are particularly challenging, as irony and exaggeration rely heavily on cultural context and discourse-level interpretation. In several cases, the model correctly captures sentiment polarity but fails to identify comedic intent.

The piecewise segmentation strategy provides limited but observable benefits for classes where stylistic signals are localized within specific clauses (e.g., punchlines in humorous text). However, for styles characterized by globally consistent tone—such as *Formal* or *Serious*—segment-level aggregation does not substantially improve discriminative capacity.

Overall, the results indicate that fine-grained stylistic modeling in Telugu requires deeper discourse modeling, improved handling of pragmatic nuance and potentially contrastive learning strategies tailored to overlapping stylistic categories.

6 Conclusion

We presented a system for Telugu prompt-style recovery that combines the multilingual representation power of XLM-RoBERTa with a piecewise segment comparison strategy. Our approach ad-

resses the challenge of localised stylistic markers by decomposing inputs into segments and aggregating predictions via max-pooling. The system achieves a Macro F1 of 0.1205 on the official test set of the DravidianLangTech @ ACL 2026 shared task. Future directions include: (1) training on augmented data generated by prompting LLMs in specific Telugu styles; (2) employing ensemble methods that combine XLM-RoBERTa with Telugu-specific language models such as MuRIL or IndicBERT; (3) incorporating contrastive learning objectives that explicitly push representations of different style classes apart in the embedding space; and (4) applying prompt-based fine-tuning approaches that leverage the stylistic label descriptions as part of the input. Despite modest performance, our work provides insight into the challenges of fine-grained stylistic modeling in low-resource Telugu NLP and highlights directions for future research.

7 Limitations

This work is limited by the size of the available training data and the inherent difficulty of stylistic annotation — a subjective task even for native speakers. Our piecewise segmentation heuristic is rule-based and may not always identify semantically meaningful boundaries, particularly in complex Telugu constructions. The results are specific to the provided dataset split and may not generalize to other Telugu style corpora without further adaptation.

Code Availability

The implementation of our system is publicly available at: [GitHub Repository](#).

8 AI Writing Assistance Disclosure

In accordance with the ACL 2023 Policy on AI Writing assistance, the authors disclose that AI-based language tools were used for minor editing assistance and formatting guidance during the preparation of this manuscript. All technical content, experimental design, implementation and analysis have been verified by the authors. The authors carefully reviewed and edited all generated text to ensure accuracy, clarity and originality.

References

- Premjith B, Jyothish Lal G, Bharathi Raja Chakravarthi, Saranya Rajiakodi, Durairaj Thenmozhi, Ratnavel Rajalakshmi, Rahul Ponnusamy, and Chinthala Bhuvanesh. 2026. Shared Task on Prompt Rrecovery for LLM in Telugu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *arXiv preprint arXiv:2004.05150*.
- Douglas Biber. 1988. *Variation across Speech and Writing*. Cambridge University Press, Cambridge.
- Bharathi Raja Chakravarthi, Ruba Asoka Chakravarthi, Vigneshwaran Muralidaran, Navya Jose, Shardul Suryawanshi, Sherly Elizabeth, and John McCrae. 2022. [Dravidiancodemix: Sentiment analysis and offensive language identification dataset for dravidian languages in code-mixed text](#). *Language Resources and Evaluation*, 56:1–42.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. [A computational approach to politeness with application to social factors](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–259, Sofia, Bulgaria. Association for Computational Linguistics.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. [Goemotions: A dataset of fine-grained emotions](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. Toward controlled generation of text. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 1587–1596, Sydney, Australia. JMLR.org.
- Harsh Jhamtani, Varun Gangal, Eduard Hovy, and Eric Nyberg. 2017. [Shakespeareizing modern language using copy-enriched sequence-to-sequence models](#). In *Proceedings of the Workshop on Stylistic Variation*, pages 10–19, Copenhagen, Denmark. Association for Computational Linguistics.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. [IndicInpsuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for indian languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4948–4961. Association for Computational Linguistics.
- Siva Subrahmanyam Varma Kusampudi, Preetham Sathineni, and Radhika Mamidi. 2021. [Sentiment analysis in code-mixed telugu-english text with unsupervised data normalization](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 753–760, Held Online. INCOMA Ltd.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized BERT pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Lei Pan, Yuanyuan Liang, Hongshen Xu, Da Ma, Su Zhu, and Shuai Fan. 2026. Unsupervised text style transfer via llms and mask-filling with multi-way interactions. In *Proceedings of Neural Information Processing*, pages 199–213, Singapore. Springer Nature Singapore.
- Bo Pang and Lillian Lee. 2008. [Opinion mining and sentiment analysis](#). *Foundations and Trends in Information Retrieval*, 2(1–2):1–135.
- Ellie Pavlick and Joel Tetreault. 2016. [An empirical analysis of formality in online communication](#). *Transactions of the Association for Computational Linguistics*, 4:61–74.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi S. Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Proceedings of the 31st Conference on Neural Information Processing Systems (NeurIPS)*, pages 6830–6841.
- Efstathios Stamatatos. 2009. [A survey of modern authorship attribution methods](#). *Journal of the American Society for Information Science and Technology*, 60(3):538–556.