VerbaNexAI at ASQP-PT 2025: Robust Detection of Tourism Aspects Using Pretrained Models and BIO Tagging in Portuguese

Melissa Moreno Novoa[1], Jairo Serrano[1], Juan Carlos Martinez-Santos[1] and Edwin Puertas[1]

[1]*Universidad Tecnologica de Bolivar, Cartagena, Colombia*

Abstract

This paper presents the system developed for the ASQP-PT 2025 shared task, explicitly addressing the Aspect Term Extraction (ATE) subtask in Portuguese tourism reviews. We used a pre-trained BERT model, which was fine-tuned with TripAdvisor reviews annotated with aspect categories, opinion terms, and sentiment polarity. Our contribution's novelty lies in designing and implementing a domain-adapted pipeline that integrates grammatical filtering, precise BIO tag generation with token alignment, and stratified evaluation using cross-validation and per-class metrics. Experimental results show that our system achieves an F1 score of 0.6108 on the ATE test set, demonstrating its ability to extract explicit aspect terms despite the challenges posed by colloquial and noisy language. Moreover, the pipeline's modular architecture allows for future extensions toward whole-aspect-opinion-category-polarity prediction. This work represents a technically robust and linguistically grounded contribution to sentiment analysis in underrepresented languages, offering a reproducible and practical solution for real-world applications in the tourism domain.

Keywords

Aspect Term Extraction, ATE, Portuguese, Tourism, BERT fine-tuning

1. Introduction

Tourism is a social and economic phenomenon that generates millions of opinions on social media and specialized platforms such as TripAdvisor. These reviews provide valuable information on lodging, food services and other tourist-related services, which are essential for agencies, operators, and destination managers to identify trends, pinpoint areas for improvement, and design strategies based on actual user feedback [1]. In this context, sentiment analysis becomes a key tool for interpreting user opinions and detecting emotional patterns. However, the growing volume and speed at which people generate these opinions pose significant challenges for manual analysis [2], highlighting the need for automatic systems capable of accurately and scalably processing large volumes of text. Thus, natural language processing (NLP) focused on sentiment analysis directly improves the quality of tourism services and strengthens data-driven decision-making processes [3].

Although sentiment analysis has seen considerable progress in other languages, particularly English, several limitations remain in Portuguese. Most efforts in NLP applied to the Portuguese language have focused on document-level sentiment analysis or tasks such as Aspect-Based Sentiment Analysis (ABSA), including the subtasks of Aspect-Based Sentiment Analysis in Portugues ABSAPT-2022 [4] and ABSAPT-2024 [5] within IberLEF. However, these approaches only partially address the complexity of sentiment analysis: they extract aspect terms or classify sentence polarity, but do not comprehensively link each aspect to its category, opinion terms, and the resulting polarity [6].

Although resources for Aspect Sentiment Quad Prediction (ASQP) already exist in English - for example, new data sets that include aspects, categories, opinions, and polarity in a single step [7] - there are still few ASQP data sets available in Portuguese. This gap underscores the need to explore

IberLEF 2025, September 2025, Zaragoza, Spain
memoreno@utb.edu.co (M. M. Novoa); jserrano@utb.edu.co (J. Serrano); jcmartinezs@utb.edu.co (J. C. Martinez-Santos); epuerta@utb.edu.co (E. Puertas)
0000-0002-1748-2208 (M. M. Novoa); 0000-0001-8165-7343 (J. Serrano); 0000-0003-2755-0718 (J. C. Martinez-Santos); 0000-0002-0758-1851 (E. Puertas)

a specific corpus for the subtasks of Aspect Term Extraction (ATE), Opinion Term Extraction (OTE), Aspect Category Detection (ACD), and Aspect Sentiment Quad Prediction (ASQP) in user-generated texts in Portuguese, to lay the groundwork for future studies and evaluations.

Portuguese reviews on TripAdvisor often exhibit colloquial variability that significantly hinders automatic processing: words with elongated vowels such as "péssimooo" to emphasize dissatisfaction or "legalzz" to express enthusiasm; abbreviations like "qnd" instead of "quando", "tb" for "também", or "mto" in place "of muito"; as well as emoticons and symbols that introduce noise into the text. These nonstandard forms prevent tokenizers from reliably recognizing lexical roots and syntactic markers, complicate text normalization, and obstruct accurately identifying relevant nominal fragments (aspects) and correctly classifying their categories and polarities.

Based on these observations, our work focuses on the ATE subtask within the ASQP-PT 2025 [8] framework, which is a shared task organized under the IberLEF [9] initiative. This study proposes and evaluates a robust pipeline for aspect-term extraction in Portuguese tourism reviews as a foundation for future full-stimulation quad prediction. In later stages, we propose a pipeline that will serve as a solid foundation for extending the analysis to complete sentiment quad prediction, ensuring that each extracted aspect term acts as a reliable anchor for the assignment of category and polarity.

2. Related Work

Analyzing the millions of reviews generated on platforms such as TripAdvisor allows one to identify the strengths and weaknesses of tourist destinations, providing critical information on culture, gastronomy, and services that help to consolidate a positive image of a destination [10]. In this context, ASQP emerges as an advanced task within ABSA, whose objective is to extract the aspect term, the opinion term, the polarity of the feelings and the aspect category simultaneously [11]. ASQP has gained traction in multiple languages; in particular, Portuguese studies were scarce until 2020 [12]. Between 2020 and 2025, approaches ranging from rule- and lexicon-based methods to transformer-based models have been introduced, along with newly annotated corpora in domains such as tourism, politics and e-commerce, yielding significant improvements in aspect extraction and classification metrics.

To organize and compare the various ABSA approaches in Portuguese, Table 1 groups the studies according to the type of subtask and the technique used without implying a strict chronological progression. -ABSA (E2E -ABSA) schemes based on fine-tuning of lightweight models with automatic annotation; classical ATE methods using Conditional Random Fields (CRF) with post-processing; hybrid proposals that combine ATE with Sentiment Orientation Extraction (SOE) through transformer ensembles and conditional text generation; and frameworks that, building on ATE, apply Aspect Sentiment Classification (ASC) via BERT variants and data augmentation strategies employing ChatGPT. It also includes extensions to non-tourism domains that contrast rule-based and heuristic named entity recognition (NER) techniques with generative models. Thus, the table provides a thematic overview of the most relevant methods in the field.

3. Data

The dataset employed in the present study consists of thousands of reviews of hotel establishments located in Paris, Las Vegas, and New York City, extracted from the TripAdvisor platform[20]. Based on the existing annotations from the ABSAPT-2024 challenge, we expanded the corpus by incorporating new annotations covering aspect terms, associated opinion terms, aspect categories, and their corresponding polarities. Each review may contain multiple quadruples, understood as combinations of aspect, sentiment term, category, and polarity, all explicitly identifiable within the text.

We stratified data segmentation to preserve the balance between categories and polarities, allocating 60% for training. Figure 1 provides a comprehensive and precise representation of the distribution of aspects, opinions, polarities, and categories within the dataset. These data are the starting point for modeling and evaluating the ABSA system. The remaining 40% was reserved for testing, subdivided

Table 1
Summary of Articles on ABSAPT

First Author	Task Summary	Technique Used
Pereira et al. [13]	End-to-End ABSA (E2E-ABSA)	Fine-tuning with PTT5, FLAN-T5, mT0 using instruction tuning; GPT-3.5 for automatic annotation
Machado et al. [14]	Aspect Term Extraction (ATE)	Conditional Random Fields (CRF) with BIO tags; POS features; post-processing with lemmatization
Resplande Sant et al. [4]	ATE and Sentiment Orientation Extraction (SOE)	Ensemble of Transformers (RoBERTa, mDeBERTa); conditional text generation with PTT5 Large
Thurow Bender et al. [5]	ATE and Aspect Sentiment Classification (ASC)	BERTimbau Large for ATE (token classification); BERT fine-tuned with ChatGPT-augmented data for ASC
Seno et al. [15]	Aspect Detection and Polarity Classification in political comments	ChatGPT vs. rule-based methods, NER heuristics, and BERT; best performance in polarity classification (PC) with ChatGPT
Zhao et al. [16]	Unified multi-task extraction of aspect terms, ASTE in single sentences	Dependency-Enhanced GCN (DE-GCN) with location-aware graphs and span-sharing joint extraction
Lin et al. [17]	Cross-lingual ABSA with adaptive rebalancing for class-imbalance issues	Equi-XABSA framework using dynamic weighted loss and anti-decoupling mechanisms
Aziz et al. [18]	Unified ABSA framework covering multiple subtasks (ATE, ASC)	BERT encoding with bi-affine attention and multi-layer GCN for modeling aspect–opinion relations
Li et al. [19]	Implicit aspect-level sentiment analysis	Generative T5 model with integrated GNN and multi-prompt fusion strategy

into 20% for ATE and OTE, and 20% for ACD and ASQP. Consequently, the primary objective of this study is to evaluate the precision, recall, and F1 score of the predictions, considering as correct only those that exactly match all four components.

4. Architecture

We structured the architecture of the aspect term extraction and Beginning-Inside-Outside (BIO) classification system into five sequential phases, ranging from ingesting the input data to generating results with exact positional references. Each phase is described in the following, following the structure presented in Figure 2.

4.1. Pre-Processing

In this initial stage, the repository is mounted, and the file containing the raw reviews is loaded. Each text then undergoes a cleaning process that removes special characters, punctuation marks, and digits and converts all content to lowercase to standardize the representation. Finally, we leverage spacy capabilities to reconstruct the reviews using their noun chunks, allowing us to retain primarily nominal entities and reduce lexical noise.

4.2. Feature Extraction

With the preprocessed text, we extract the essential features the model requires. First, we reorganize the dataset columns to obtain the aspect term and its corresponding start and end positions within

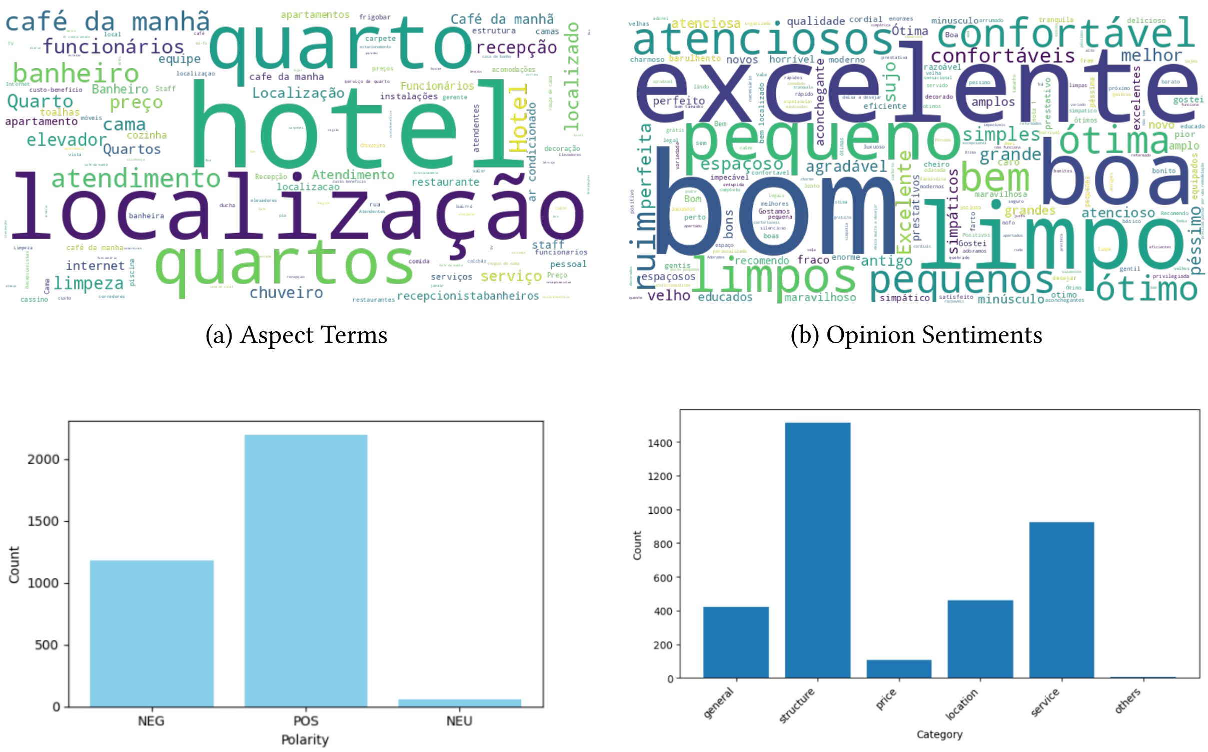

(a) Aspect Terms

(b) Opinion Sentiments

(c) Sentiment Polarity

(d) Aspect Categories

Figure 1: Visual summary data train aspect-based sentiment analysis.

the original string. We then convert the list of associated opinion terms into a single text string. To generate the BIO labels, we tokenize the preprocessed text using spacy and assign each token a B-TERM, I-TERM, or O label depending on whether it matches the position of the aspect term, as illustrated in Table 2. Finally, we use a configured BertTokenizerFast to align these word-level tokens with the BERT subtokens, ensuring a precise correspondence between the labels and the model input.

Table 2
BIO tagging example for ATE

Token	BIO Tag
Avaliacao	B-TERM
detalhada	I-TERM
e	O
internet	B-TERM
do	O
ibis	B-TERM

4.3. Regularization

Since the O label is overwhelmingly predominant compared to B-TERM and I-TERM, we randomly split the dataset into 80% for training and 20% for testing. We then compute class weights using only the labels from the training set. We incorporated these weights into the loss function via a custom Trainer. Hence, errors in the much less frequent B-TERM and I-TERM classes are more important.

4.4. Classifiers

We conducted the classification phase by fine-tuning the BERT base Portuguese case, known as BERTimbau [21], for a token classification task with three labels. We define training arguments with a batch size of 8, a learning rate of 2×10^{-5}, a warm-up of 200 steps, and checkpoint saving at the end of each epoch. We apply 3-fold cross-validation (K-Fold) to ensure model robustness and include a callback mechanism with three-epoch iterations to prevent overfitting.

4.5. Evaluation

After each validation fold, we compute the accuracy, precision, recall, and F1 scores (macro- and per-label) on the validation set. We select the model that maximizes the macro F1 score and apply it to the entire set of tests. Based on the predicted BIO labels, we reconstruct the aspect spans within the original text and compare them with the reference annotations. Finally, we exported all results to an Excel file, including text, tokens, gold and predicted labels, extracted terms, and their positions, enhancing interpretability.

5. Experimental Results

During training, the model produced standard warnings about overlapping faces with respect to initialized weights, which did not prevent convergence. In each fold, we calculated general metrics (accuracy, precision, recall, and macro-F1) and label-specific metrics (O, B-TERM, and I-TERM) to assess the model's ability to identify aspect terms and their continuations correctly.

Table 5 reports the most representative metrics for the first two evaluated folds. As observed, although the overall accuracy metrics remain above 87%, the performance in the B-TERM and I-TERM classes presents low F1 scores, particularly due to still limited precision. However, the recall for both labels is high, indicating that the model successfully identifies a significant proportion of relevant terms, albeit with some confusion in their exact boundaries.

The confusion matrix obtained in Figure 3 shows that although the model achieves solid performance in classifying the majority class O - tokens that do not represent aspect terms - it faces challenges in correctly identifying the minority classes B-TERM and I-TERM. Specifically, out of 91,821 tokens truly labeled as 'O', 86,384 were correctly classified, while 3,695 and 1,742 were incorrectly assigned to 'B-TERM' and 'I-TERM', respectively. This indicates a tendency of the model to produce false positives in the detection of aspect terms. On the other hand, out of 915 true 'B-TERM' tokens, the model correctly identified 709, with an approximate recall of 77.5%, but with low precision (16%) due to the large number of false positives originating from the 'O' class. Similarly, for the 'I-TERM' class, out of 528 true tokens, only 330 were correctly classified, with an approximate recall of 62.5% and equally low precision (16%).

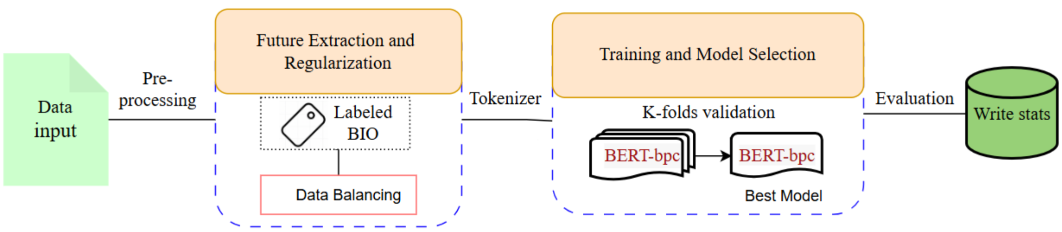

Figure 2: Pepeline

These results demonstrate that the model behaves consistently when identifying tokens of the O class (noun chunks) yet still faces challenges in precisely delimiting compound terms (B-TERM, I-TERM). The high recall in these labels suggests that the model can serve as a foundation for post-processing systems aimed at refining initial predictions. Additionally, we explored the model to identify noun phrases that could serve as initial candidates for aspect terms. Although the model enabled the construction of

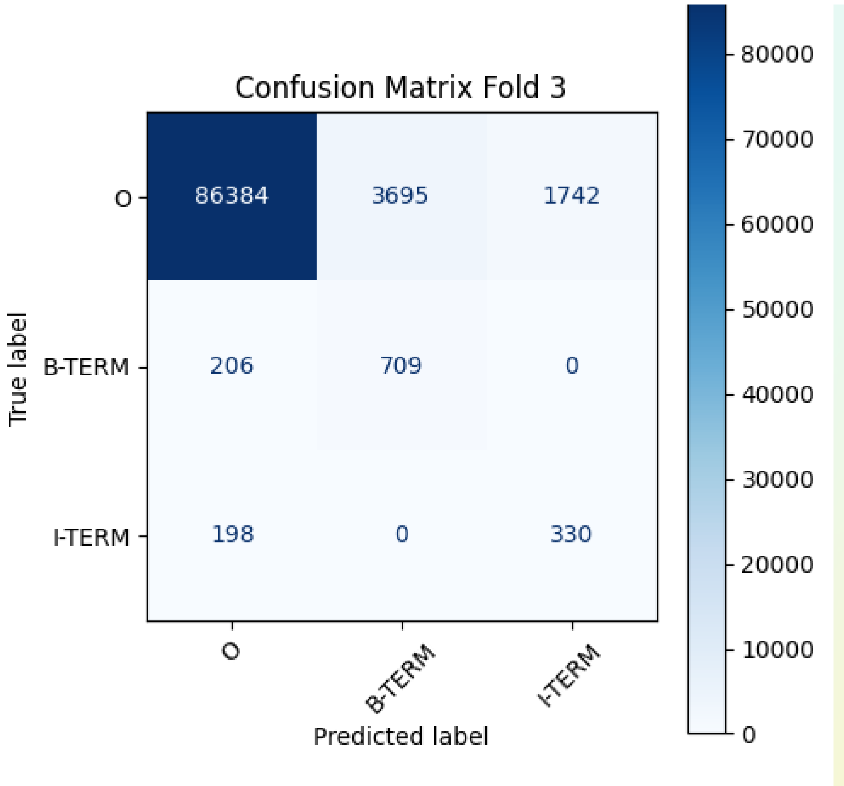

Figure 3: Confusion matrix of the model for Aspect Term classification (BIO labels) – Fold 3.

Table 3
ATE — Task 1: Performance comparison

Team VerbanexAI	Precision	Recall	F1_score
Train	0.4135	0.8697	0.4559
Test	0.6125	0.6091	0.6108

structured lists, including noun chunks, head nouns, their grammatical functions, and related words, the manual analysis revealed unreliable coverage.

For example, when applied to an extended review containing multiple evaluative expressions, spaCy identified phrases such as "infraestruturas bastante degradadas", "casa de banho partilhada", or "a cortina", but also registered ambiguous or irrelevant chunks such as "que", "piso", "alguém" or "convívio", many of which do not represent analyzable aspects nor are associated with consistent semantic judgments.

It suggests that although spaCy processing provides a practical lexical foundation, its direct applicability for automatic BIO label generation requires subsequent filtering based on grammatical rules or supervised training, especially in contexts involving complex and lengthy opinions, such as TripAdvisor reviews in Portuguese.

Thus, once we completed the training and cross-validation process, without applying noun fragments but instead using preprocessed data as outlined in Section 4.1, the model with the best macro-F1 performance was selected and applied to the complete test set. Table 3 summarizes the system's final performance on the aspect term extraction task (ATE). During training, the model achieved a precision of 0.4135, a recall of 0.8697, and an F1 score of 0.4559, highlighting its ability to cover many relevant terms, although with moderate precision.

We implemented a validation and inference pipeline on an unseen test set to evaluate the model's generalization capability. First, we cleaned and normalized each review by removing punctuation via regular expressions to ensure consistency with the training format. Next, we tokenized the preprocessed texts once using BertTokenizerFast, preserving their associated BIO labels, and fed them directly

Table 4
Comparative performance on the official ranking

Team	Precision	Recall	F1_score
Task best system	0.6898	0.7305	0.7096
VerbanexAI	0.6125	0.6091	0.6108

into the trained model, which returned a sequence of predicted labels (B-TERM, I-TERM, or O). No hyperparameter search was performed.

Subsequently, we applied a decoding step to reconstruct aspect terms from the labelled sequences. This step grouped consecutive tokens labelled as B-TERM and I-TERM, determining their exact character-level positions within the original text. In this way, compound terms, even when they spanned multiple tokens, could be accurately recovered.

The decoding function iterates over each token along with its predicted label, recording each term's start and end position. When a B-TERM label is detected, we initiate a new term; if followed by an I-TERM label, we extend the term. Upon encountering an O label or reaching the end of the sequence, the term is finalized and stored. The extracted terms include both the textual content and their positions. Therefore, we stored the reconstructed information in an organized structure, and each review is associated with its original text and a list of identified aspect terms. Each term includes the lexical content and its coordinates within the text. This final set of predictions is exported to an Excel file, facilitating manual inspection and further analysis.

In the final evaluation of the test set, our system achieved a precision of 0.6125, a recall of 0.6091, and an F1 score of 0.6108. Although the best performing task system reached a precision of 0.6898, a recall of 0.7305, and an F1 score of 0.7096, see Table 4, our approach still demonstrates competitive aspect detection and maintains balanced performance when generalizing to unseen data. These results underscore the practical applicability of the model for Portuguese-language sentiment analysis in the tourism domain, particularly given that aspect terms may span multiple words and be expressed in varied ways, and pave the way for future enhancements.

Table 5
Fold-wise performance of the BIO model for aspect term extraction

Fold	**Accuracy**	**F1 Macro**	**B-TERM F1**	**B-TERM Recall**	**I-TERM F1**	**I-TERM Recall**	**O F1**
Fold 1	0.9150	0.4711	0.2284	0.8534	0.2298	0.6863	0.9550
Fold 2	0.8795	0.4407	0.2072	0.9568	0.1801	0.9259	0.9349
Average	0.8973	0.4559	0.2178	0.9051	0.2050	0.8061	0.9450

6. Future Work

We recommend prioritizing linguistic data augmentation strategies for future iterations, such as synonym replacement or back-translation, especially given the limited coverage observed for infrequent aspects. These techniques enhance the model's generalization capabilities when encountering lexical variations and less common expressions.

Although adjustments to class weights were implemented to address the observed class imbalance, the results obtained reveal ongoing precision challenges, particularly in correctly identifying aspect terms within minority classes ('B-TERM' and 'I-TERM'). This suggests that class weighting alone did not fully resolve the problem, indicating that additional factors, such as semantic similarity between labeled and unlabeled terms or significant structural class imbalance, require more sophisticated strategies. Therefore, future research should evaluate complementary methods, synthetic oversampling techniques, or hybrid approaches incorporating linguistic rules to mitigate observed false positives.

The visualization of the predicted aspects (see Figure 4) clearly highlights the challenges of the current system. The word cloud generated from these predictions demonstrates a need for refinement in aspects such as lexical unification, error reduction in aspect detection, and consistent handling of capitalization and segmentation.

Finally, given the model's sensitivity to specific parameters during training, we recommend conducting a more exhaustive hyperparameter search, including learning rate, batch size, and number of training epochs. Furthermore, we suggest exploring multitask learning approaches capable of simultaneously detecting aspects and opinion terms, or reframing the problem as a multi-label classification task to capture overlapping sentiment expressions. In the long term, we advise extending the study towards a more comprehensive framework encompassing the remaining subtasks of aspect sentiment quad prediction to enable a more robust evaluation of the model's performance in real-world scenarios.

Figure 4: Map of aspects extracted from the predictions made on the analyzed reviews.

7. Conclusion

The results of this study demonstrate that combining rigorous linguistic preprocessing with the fine-tuning of a BERT model for Portuguese enables accurate and robust identification of aspect terms in tourism reviews. This approach validates the potential of contextual language models to address specialized tasks in underrepresented languages. It underscores the importance of adapting natural language processing techniques to the lexical and syntactic particularities of the tourism domain.

This work's main contribution lies in the design and implementation of a comprehensive pipeline for aspect term extraction in Portuguese capable of handling the informal and noisy language frequently found on platforms such as TripAdvisor. Our proposal integrates strategies such as grammatical filtering, precise BIO tag generation with token alignment, and detailed evaluation through cross-validation and per-class metrics. Together, these components result in an adaptable, transparent, and high-performance system for the ATE subtask of the ASQP-PT 2025 challenge.

Furthermore, the model effectively performs on the training data and maintains a balanced performance when generalizing to unseen texts. It is a practical tool for opinion analysis systems and personalized recommendation services in the tourism sector. Its architecture also allows for future extensions to more complex tasks, such as predicting whole aspect–opinion–category–polarity quadruples.

Nonetheless, limitations related to corpus size and the difficulty of capturing infrequent or implicit aspects are acknowledged. As a result, future work will explore active learning techniques, multilingual transfer learning, and linguistic data augmentation. In summary, this work provides a technically robust and scientifically relevant solution for aspect-based analysis in Portuguese, laying the groundwork for developing more inclusive, efficient, and practically applicable ABSA systems.

Declaration on Generative AI

During the preparation of this investigation, ChatGPT (OpenAI) was used for the revision of translations into English, as well as for grammatical and spelling correction. After using this tool, the content was reviewed and edited as necessary, and full responsibility for the content of the publication is assumed.

Acknowledgments

The authors express their gratitude to the Call 933 “Training in National Doctorates with a Territorial, Ethnic and Gender Focus in the Framework of the Mission Policy — 2023” of the Ministry of Science, Technology and Innovation (Minciencia). In addition, we thank the team of the Artificial Intelligence Laboratory VerbaNex (https://github.com/VerbaNexAI), affiliated with the Universidad Tecnológica de Bolívar (UTB), for their contributions to this project.

References

[1] M. Chu, Y. Chen, L. Yang, J. Wang, Language interpretation in travel guidance platform: Text mining and sentiment analysis of tripadvisor reviews, Frontiers in Psychology 13 (2022) 1029945. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC9607914/. doi:10.3389/FPSYG.2022.1029945.

[2] M. A. Álvarez-Carmona, R. Aranda, A. Y. Rodríguez-Gonzalez, D. Fajardo-Delgado, M. G. Sánchez, H. Pérez-Espinosa, J. Martínez-Miranda, R. Guerrero-Rodríguez, L. Bustio-Martínez, Ángel Díaz-Pacheco, Natural language processing applied to tourism research: A systematic review and future research directions, Journal of King Saud University - Computer and Information Sciences 34 (2022) 10125–10144. doi:10.1016/j.jksuci.2022.10.010.

[3] K. R. Chowdhary, Natural language processing, Fundamentals of Artificial Intelligence (2020) 603–649. URL: https://link.springer.com/chapter/10.1007/978-81-322-3972-7_19. doi:10.1007/978-81-322-3972-7_19.

[4] J. R. S. Gomes, E. A. S. Garcia, A. F. B. Junior, R. C. Rodrigues, D. F. C. Silva, D. F. Maia, N. F. F. da Silva, A. R. G. Filho, A. da Silva Soares, Deep learning brasil at absapt 2022: Portuguese transformer ensemble approaches, 2023. URL: https://arxiv.org/abs/2311.05051. arXiv:2311.05051.

[5] A. T. Bender, G. A. Gomes, E. P. Lopes, R. M. Araújo, L. A. de Freitas, U. B. Corrêa, Overview of absapt at iberlef 2024: Overview of the task on aspect-based sentiment analysis in portuguese, Procesamiento del lenguaje natural, ISSN 1135-5948, Nº. 73, 2024, págs. 315-322 (2024) 315–322. URL: https://dialnet.unirioja.es/servlet/articulo?codigo=9736870&info=resumen&idioma=ENGhttps://dialnet.unirioja.es/servlet/articulo?codigo=9736870&info=resumen&idioma=SPAhttps://dialnet.unirioja.es/servlet/articulo?codigo=9736870.

[6] R. R. L. Barbosa, S. Sánchez-Alonso, M. A. Sicilia-Urban, Evaluating hotels rating prediction based on sentiment analysis services, Aslib Journal of Information Management 67 (2015) 392–407. URL: https://sites.google.com/inf.ufpel.edu.br/asqp-pt-2025/description. doi:10.1108/AJIM-01-2015-0004.

[7] J. Zhou, H. Yang, Y. He, H. Mou, J. Yang, A unified one-step solution for aspect sentiment quad prediction, Proceedings of the Annual Meeting of the Association for Computational Linguistics (2023) 12249–12265. URL: https://arxiv.org/pdf/2306.04152. doi:10.18653/v1/2023.findings-acl.777.

[8] E. P. Lopes, G. Gomes, A. T. Bender, R. M. Araújo, L. A. de Freitas, U. B. Corrêa, Overview of asqp-pt at iberlef 2025: Overview of the task on aspect-sentiment quadruple prediction in portuguese, Procesamiento del Lenguaje Natural 75 (2025).

[9] González-Barba, J. Ángel, Chiruzzo, Luis, Jiménez-Zafra, S. María, Overview of IberLEF 2025: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2025), co-located with the 41st Conference of the Spanish Society for Natural Language Processing (SEPLN 2025), CEUR-WS. org, 2025.

[10] R. M. L. Freire, K. F. B. Maracajá, V. Valduga, A. B. F. M. do Nascimento, Análise da imagem afetiva dos turistas no destino vale dos vinhedos - rs, Turismo: Visão e Ação 25 (2023) 114–133. doi:10.14210/rtva.v25n1.p114-133.
[11] A. Wang, J. Jiang, Y. Ma, A. Liu, N. Okazaki, Generative data augmentation for aspect sentiment quad prediction, in: A. Palmer, J. Camacho-collados (Eds.), Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 128–140. URL: https://aclanthology.org/2023.starsem-1.12/. doi:10.18653/v1/2023.starsem-1.12.
[12] D. A. Pereira, A survey of sentiment analysis in the portuguese language, Artificial Intelligence Review 54 (2021) 1087–1115. doi:10.1007/s10462-020-09870-1.
[13] G. Pereira, L. Barbosa, J. Moreira, T. Melo, A. Silva, Enhancing aspect-based sentiment analysis for portuguese using instruction tuning, in: Anais do XXI Encontro Nacional de Inteligência Artificial e Computacional (ENIAC 2024), Sociedade Brasileira de Computação - SBC, 2024, pp. 990–1001. doi:10.5753/eniac.2024.245109.
[14] M. T. Machado, T. A. S. Pardo, Nilc at absapt 2022: aspect extraction for portuguese, in: Conference of the Spanish Society for Natural Language Processing - SEPLN, CEUR-WS, 2022.
[15] E. Seno, L. Silva, F. Anno, F. Rocha, H. Caseli, Aspect-based sentiment analysis in comments on political debates in Portuguese: evaluating the potential of ChatGPT, in: P. Gamallo, D. Claro, A. Teixeira, L. Real, M. Garcia, H. G. Oliveira, R. Amaro (Eds.), Proceedings of the 16th International Conference on Computational Processing of Portuguese - Vol. 1, Association for Computational Lingustics, Santiago de Compostela, Galicia/Spain, 2024, pp. 312–320. URL: https://aclanthology.org/2024.propor-1.32/.
[16] M. Zhao, J. Yang, F. Shang, Dependency-enhanced graph convolutional networks for aspect-based sentiment analysis, Neural Computing and Applications 35 (2023) 14195–14211. doi:10.1007/S00521-023-08384-5.
[17] J. Šmíd, P. Král, Cross-lingual aspect-based sentiment analysis: A survey on tasks, approaches, and challenges, Information Fusion 120 (2025) 103073. URL: https://www.sciencedirect.com/science/article/pii/S1566253525001460. doi:10.1016/J.INFFUS.2025.103073.
[18] K. Aziz, D. Ji, P. Chakrabarti, T. Chakrabarti, M. S. Iqbal, R. Abbasi, Unifying aspect-based sentiment analysis bert and multi-layered graph convolutional networks for comprehensive sentiment dissection, Scientific Reports 14 (2024) 1–22. URL: https://www.nature.com/articles/s41598-024-61886-7. doi:10.1038/S41598-024-61886-7;SUBJMETA=114,1305,2164,631;KWRD=DATA+MINING,MACHINE+LEARNING.
[19] X. Li, X. Wang, C. Yao, Y. Li, Graph-enhanced implicit aspect-level sentiment analysis based on multi-prompt fusion, Scientific Reports 15 (2025) 1–19. URL: https://www.nature.com/articles/s41598-025-02609-4. doi:10.1038/S41598-025-02609-4;SUBJMETA=2811,477,631;KWRD=HUMAN+BEHAVIOUR,PSYCHOLOGY.
[20] N. G. Ángel, N. N. Gómez, M. D. P. Esteban, I. M. Rubio, D. P. Góngora, Emprendimiento social en época de pandemia: Revisión bibliográfica, Miradas sobre el emprendimiento ante la crisis del coronavirus, 2022, ISBN 978-84-1377-995-9, págs. 528-533 (2022) 528–533. URL: https://dialnet.unirioja.es/servlet/autor?codigo=4224094.
[21] F. Souza, R. Nogueira, R. Lotufo, BERTimbau: pretrained BERT models for Brazilian Portuguese, in: 9th Brazilian Conference on Intelligent Systems, BRACIS, Rio Grande do Sul, Brazil, October 20-23 (to appear), 2020.