

ELTSA-CUJAE at Rest-Mex 2025: A Novel Ensemble Learning Approach with Transformers for Mining Mexican Tourist Reviews

Miguel Angel Rivero-Tapia^{1,*†}, Alfredo Simon-Cuevas² and Ray Maestre Peña^{3,†}

¹Technological University of Havana “José Antonio Echeverría”, CUJAE, Cuba

Abstract

This paper describes the results of the team’s participation in the Rest-Mex 2025 competition, focused on the analysis of tourist reviews related to Mexico’s Magical Towns. The proposed solution addresses three main tasks: Sentiment polarity classification, establishment type identification, and detection of the associated Magical Town for each review. To this end, an ensemble of Transformer-based models was implemented, combined using the Zimmerman-Zysno operator.

The proposed approach proves to be highly effective, outperforming individual models due to its ability to capture different nuances of language in the reviews. The results suggest that this strategy can be a valuable tool for opinion analysis in the tourism sector, enabling businesses and governments to monitor visitor satisfaction, optimize marketing strategies, and improve destination management. This work not only contributes to advances in natural language processing for domain-specific tasks but also offers a reproducible framework for similar applications in other contexts.

Keywords

Sentiment Analysis, Tourist reviews, Transformer-based models, Zimmerman-Zysno operator, Text classification

1. Introduction

Sentiment analysis is the field dedicated to the computational study of opinions, sentiments, evaluations, attitudes, and emotions expressed by individuals toward entities such as products, services, organizations, people, issues, events, and their characteristics or attributes [1]. Opinions refer to a person’s viewpoint on a given topic, while sentiment or polarity refers to the emotion experienced by that person regarding the topic, usually identified as positive or negative.

The computational analysis of tourist reviews has become an essential tool for strategic decision-making in the tourism sector [2, 3, 4, 5]. In this context, the present work introduces an innovative solution developed for the Rest-Mex 2025 competition [6, 7], specifically focused on the analysis of opinions related to Mexico’s Magical Towns. Unlike past shared task [8, 9, 10], this edition of the challenge introduces significant complexities by requiring not only the classic sentiment polarity classification and establishment type identification but also the accurate recognition of the specific Magical Town mentioned in each review—a task particularly challenging due to the cultural and linguistic traits of these destinations.

Transformer models have revolutionized the computational approach to text classification in domain-specific scenarios [11, 12]. Architectures such as BERT [13], RoBERTa [14], and their specialized variants have demonstrated exceptional capabilities in capturing both emotional nuances and the linguistic particularities of domains where formal and informal language, regionalisms, and implicit cultural references coexist. These models, pre-trained on large corpora and later fine-tuned for specific tasks,

IberLEF 2025, September 2025, Zaragoza, Spain

*Corresponding author.

† These authors contributed equally.

✉ miguelangelriverotapia@gmail.com/ (M. A. Rivero-Tapia); asimon@ceis.cujae.edu.cu (A. Simon-Cuevas); raymaestrex@gmail.com (R. M. Peña)

ORCID 0009-0003-8289-6567 (M. A. Rivero-Tapia); 0000-0002-6776-9434 (A. Simon-Cuevas); 0009-0006-8179-5811 (R. M. Peña)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

significantly outperform traditional approaches by learning contextualized representations that allow them to accurately interpret complex expressions such as irony, contrast, or nuanced evaluations.

The strategic combination of multiple Transformer architectures has emerged as an effective approach, overcoming the limitations of individual models by capturing different levels of meaning and context in reviews. This paradigm has established new benchmarks in classification accuracy, especially in languages like Spanish, where specialized resources are more limited than in English.

Unlike traditional approaches that use lexical models or monolithic architectures, the proposed solution is based on a Transformer ensemble specifically adapted to capture the subtleties of Spanish-language tourism content. The implemented solution strategically combines several Transformer models using the Zimmerman-Zysno operator [15], a fuzzy fusion method developed by German researchers Bernd Zimmerman and Peter Zysno. This operator is notable for its ability to aggregate probabilistic value judgments in a non-additive manner, balancing between strict conjunctions (AND) and permissive disjunctions (OR) through a compensation parameter (γ). Unlike conventional averaging methods, the Zimmerman-Zysno operator preserves each model’s relative uncertainty and dynamically adjusts their influence in the final prediction, which is critical for domains with subjective language such as tourism [15], [16].

This article is structured with a section presenting the unique corpus of over 200,000 TripAdvisor reviews used in the competition, followed by a detailed description of the proposed Transformer ensemble system. Subsequently, experimental results validating the approach’s effectiveness are analyzed. The paper concludes with a discussion of the broader implications of this research and future directions for automated analysis of tourism-related content.

2. Task Description and Dataset

The Rest-Mex 2025 dataset represents a significant advancement in the analysis of tourist perceptions, focusing exclusively on Mexico’s Magical Towns. This corpus contains 208,051 reviews, equivalent to 70% of the total available data, while the remaining 30% is reserved for evaluation. Each entry includes six key fields: the review title, the full review text, sentiment polarity (on a scale from 1 to 5), the evaluated Magical Town, the corresponding region, and the type of establishment (hotel, restaurant, or tourist attraction).

The polarity distribution shows a characteristic pattern of this type of analysis, with a marked predominance of positive reviews. This distribution is presented in Table 2.

Table 1

Sentiment polarity distribution

Class	Instances
1	5,441
2	5,496
3	15,519
4	45,034
5	136,561

As shown in Table 2, the most positive reviews (class 5) account for 65.64% of the dataset, while the negative classes (1 and 2) together represent only 10,937 instances (5.26%). This natural imbalance in tourist review data poses significant challenges for predictive modeling.

On the other hand, the distribution of establishment types is presented in Table 3:

Restaurants appear as the most reviewed category (41.68%), highlighting the importance of the gastronomic experience in Mexico’s Magical Towns. Tourist attractions and hotels follow with 69,921 and 51,410 reviews respectively.

The geographic distribution reveals a notable concentration in coastal and colonial destinations, as seen in Table 4:

Table 2

Sentiment polarity distribution

Class	Instances
1	5,441
2	5,496
3	15,519
4	45,034
5	136,561

Table 3

Establishment type distribution

Type	Instances
Hotel	51,410
Restaurant	86,720
Attraction	69,921

Table 4

Geographic distribution of Magical Towns

Town	Count
Tulum	45,345
Isla Mujeres	29,826
San Cristóbal de las Casas	13,060
⋮	⋮
Real de Catorce	760
Tapalpa	725

This strategic selection of 40 towns allows the study of specific regional patterns, ranging from the Yucatán Peninsula to the central highlands, which is particularly valuable for Subtask 3 that requires comparative analysis across regions. Destinations such as Valladolid (11,637), Bacalar (10,822), and Palenque (9,512) represents the southeast, while Valle de Bravo (5,959) and Teotihuacán (5,810) cover the country’s central area.

One of the key innovations of this edition is the precise geographic labeling, which enables detailed analysis by locality and region. The corpus is encoded in UTF-8 to preserve the richness of the Spanish language.

3. Evaluation Metrics

The participating systems are evaluated using standard metrics such as Precision, Recall, and F1-score, applied to each proposed subtask. For sentiment polarity classification, the per-class F-measure is used, where $F_i(k)$ represents the F1-score for class i obtained by system k , with $C = \{1, 2, 3, 4, 5\}$ as the set of possible categories.

The popularity resolution (Res_P) is computed as the average of F-measures across all classes, using the following formula:

$$\text{Res}_P(k) = \frac{\sum_{i=1}^{|C|} F_i(k)}{|C|} \quad (1)$$

For thematic resolution (Res_T), performance is assessed in three specific categories: Attraction (F_A), Hotel (F_H), and Restaurant (F_R), by averaging their respective F-measures:

$$\text{Res}_T(k) = \frac{F_A(k) + F_H(k) + F_R(k)}{3} \quad (2)$$

In the task of Magical Town identification (Res_{MT}), the metric is the average F-measure obtained for each town in the list MTL:

$$\text{Res}_{MT}(k) = \frac{\sum_{i=1}^{|\text{MTL}|} F_{\text{MTL}_i}(k)}{|\text{MTL}|} \quad (3)$$

Finally, the overall sentiment score combines the three subtasks, assigning greater weight to polarity (Res_P) and Magical Town resolution (Res_{MT}), as expressed in the following formula:

$$\text{Sentiment}(k) = \frac{2 \cdot \text{Res}_P(k) + \text{Res}_T(k) + 3 \cdot \text{Res}_{MT}(k)}{6} \quad (4)$$

This weighting reflects the relative importance of each subtask in the comprehensive evaluation of the system.

4. System Description

Sentiment analysis applied to the tourism domain can be approached as a text classification task. This technique is part of Natural Language Processing (NLP) and aims to determine the category to which a given text fragment belongs [17].

The most common methodology for solving text classification problems is based on supervised learning. In this paradigm, a structured dataset is used, composed of pairs of elements: on the one hand, the inputs (texts), and on the other, the outputs (corresponding labels or categories). The goal is to train a model capable of automatically predicting the category of a new text, even if it was not part of the training data. It is important to note that the possible classifications are limited to the labels present in the initial dataset.

4.1. Pre-processing

Sentiment analysis requires pre-processing to prepare the text data and facilitate the extraction of relevant information from reviews. While pre-processing tasks vary depending on the objective, certain steps are crucial and indispensable for training Transformer models, whereas others can be omitted.

Transformer models have been trained on millions of texts that retain information such as capitalization and punctuation. Therefore, it is unnecessary to convert text to lowercase or remove punctuation marks, as the model has learned to leverage these features to generate high-quality representations. Similarly, removing stop words (such as articles and prepositions) is also unnecessary, since they are already included in the training corpus and are properly handled by the tokenizer [18].

The first step in pre-processing is removing special characters, as they are not understood by the models. Then, for the model training phase, the text was converted into a specific format interpretable by the language model.

4.2. Model Training and Ensemble

Text classification is performed using an ensemble of Transformer models, where predictions are optimally combined via the Zimmermann-Zysno operator [15]. This operator is a nonlinear fuzzy fusion method that aggregates probabilistic judgments from multiple models, solving the problem of weighting each model's relative influence on the final decision. Mathematically, the operator is defined for each class j as:

$$Z_j(s_1, s_2, \dots, s_n) = \left(\prod_{i=1}^n s_i \right)^{1-\gamma} \cdot \left(1 - \prod_{i=1}^n (1 - s_i) \right)^{\gamma} \quad (5)$$

where $s_i \in [0, 1]$ represents the probability assigned by the i -th model to class j , and $\gamma \in [0, 1]$ is the compensation parameter that regulates the balance between conjunction (t-norm) and disjunction (t-conorm) operations. The term γ is computed adaptively as:

$$\gamma = \frac{T(s_1, \dots, s_n)}{T(s_1, \dots, s_n) + T(1 - s_1, \dots, 1 - s_n)}, \quad \text{with} \quad T(s_1, \dots, s_n) = \prod_{i=1}^n s_i. \quad (6)$$

This formulation allows the operator to act as an adaptive hybrid. When $\gamma \rightarrow 0$, conjunction dominates (equivalent to the product of probabilities), requiring high consensus among models to assign confidence to a prediction. When $\gamma \rightarrow 1$, disjunction prevails (similar to a compensated sum), where partial disagreement between models does not invalidate the prediction. For intermediate values ($\gamma \approx 0.5$), the operator balances both extremes through a weighted geometric mean, ideal for balanced distributions.

In practice, the ensemble operates on the probabilistic outputs of each Transformer model for the k possible classes (e.g., polarities 1-5 or establishment types). Each model M_i generates a vector $P_{M_i} = [p_{i1}, p_{i2}, \dots, p_{ik}]$, where $p_{ij} \in [0, 1]$ is the probability assigned to class j and $\sum_{j=1}^k p_{ij} = 1$. The aggregation process consists of three stages:

1. **Per-class computation:** For each class j , the Zimmermann-Zysno operator is applied to the probabilities $p_{1j}, p_{2j}, \dots, p_{nj}$ from the n models, yielding an aggregated probability vector $Z = [Z_1, Z_2, \dots, Z_k]$.
2. **Normalization:** To ensure $\sum_{j=1}^k Z_j = 1$, softmax normalization is applied:

$$Z_j^{(\text{norm})} = \frac{e^{Z_j}}{\sum_{m=1}^k e^{Z_m}}. \quad (7)$$

3. **Final prediction:** The selected class is $\arg \max_j Z_j^{(\text{norm})}$.

The interpretation of γ is crucial in the applied context. For instance, a low value ($\gamma = 0.2$) favors fuzzy intersection (AND), useful when individual model accuracy is high but their errors are independent, reducing false positives at the cost of discarding minority predictions. An intermediate value ($\gamma = 0.5$) implements an optimal trade-off between sensitivity and specificity, recommended for uniformly distributed data. A high value ($\gamma = 0.8$) prioritizes fuzzy union (OR), capturing weak but consistent signals from at least one model, which is relevant in tasks with high semantic noise, such as polarity classification in tourism reviews, where models may diverge in subjective interpretations but provide complementary information.

Consider three models (M_1, M_2, M_3) and two classes (A and B) with $\gamma = 0.5$, where the original outputs are $P_{M_1} = [0.7, 0.3]$, $P_{M_2} = [0.6, 0.4]$ and $P_{M_3} = [0.8, 0.2]$. The aggregation for class A yields:

$$\begin{aligned} Z_A &= (0.7 \cdot 0.6 \cdot 0.8)^{0.5} \cdot (1 - (1 - 0.7)(1 - 0.6)(1 - 0.8))^{0.5} \\ &= (0.336)^{0.5} \cdot (0.976)^{0.5} \approx 0.58, \end{aligned}$$

while for class B we obtain $Z_B \approx 0.13$. After normalization, the final vector $Z^{(\text{norm})} \approx [0.61, 0.39]$ predicts class A. This example illustrates how γ controls the weighting between agreements and disagreements. This approach explains why high γ values (e.g., $\gamma = 0.8$) may be optimal for minority classes or scenarios with semantic noise, where partially consistent signals should be preserved. The solution workflow is shown in Figure 1.

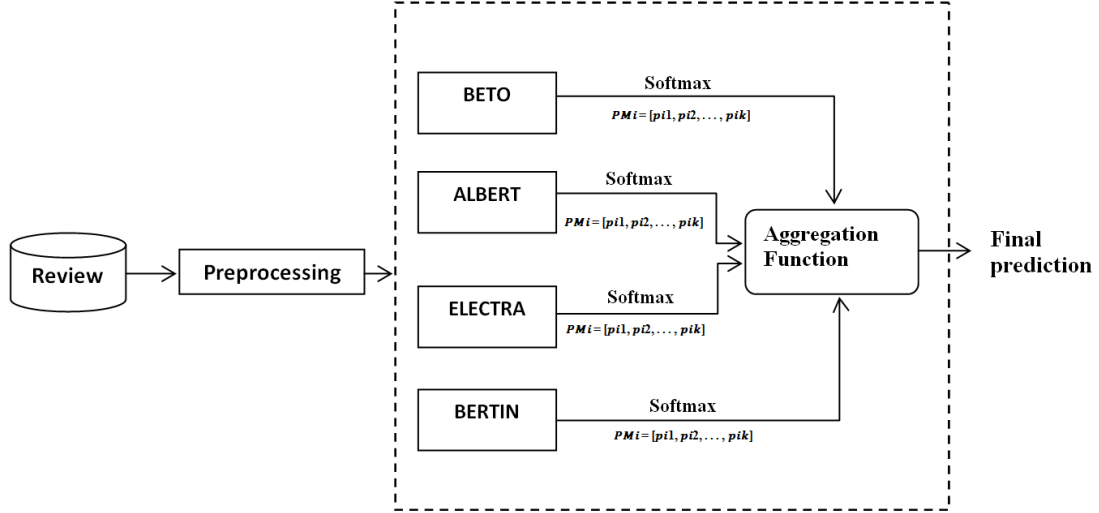


Figure 1: Workflow of the ensemble strategy using Transformer models combined via the Zimmerman-Zysno operator.

4.3. Training and Evaluation

To train the models, a representative subset of the available data was used: 15% of the opinions for the sentiment polarity classification (Task 1) and establishment type identification (Task 2), and 40% for the detection of the specific “Pueblo Mágico” (Task 3). This decision was primarily due to computational constraints, as training Transformer models on full datasets requires significant computational capacity, both in hardware and processing time.

The training and ensemble process consists of three fundamental stages. In the first stage, pre-trained Spanish Transformer models, including BETO [13], ALBERT [18], BERTIN [14], and ELECTRA [19], were selected and fine-tuned using the aforementioned data subsets. Each model was individually adapted to the three tasks. In the second stage, probabilistic predictions were generated for the test reviews by each model. Finally, in the third stage, these predictions were combined using the Zimmerman-Zysno operator, optimizing the γ parameter to balance precision and robustness.

4.4. Evaluation and Discussion

The experimental evaluation of the solution was carried out on the set of Spanish-language tourist reviews provided by the Rest-Mex 2025 competition. Each review contains a unique ID, a title, and the opinion text. It is worth noting that this dataset is one of the few available in Spanish for evaluating this type of solution. The results were computed using the equations described in the previous section.

The evaluation process considers four pre-trained Transformer models: BETO [13], ALBERT [18], BERTIN [14], and ELECTRA [19]. A fine-tuning process was applied to each model. In addition, each model was trained on the dataset reviews for 2 epochs using a batch size of 16. The choice of this batch size is based on several factors: large sizes may cause fluctuations in gradient updates; while they accelerate convergence, they can also lead to overfitting, harming the model’s generalization.

Table 5 shows the hyperparameter configuration used in the experiments.

To evaluate the performance of the Zimmerman-Zysno operator, experiments were conducted using different values of the γ parameter (0.2, 0.5, and 0.8), which allows adjusting the balance between strict conjunctions and permissive disjunctions in the prediction fusion process. These values were selected to explore a representative range, spanning from a more conservative approach ($\gamma = 0.2$) to a more flexible one ($\gamma = 0.8$). Additionally, an ensemble based on average voting across the four Transformer models (BETO, ALBERT, BERTIN, and ELECTRA) was implemented as a baseline, in order

Table 5

Hyperparameter configuration used during model training.

Hyperparameter	Value
Seed	42
Learning rate	2e-5
Epochs	2
Batch size	16
Maximum sequence length	512

to contrast the proposed method’s performance with traditional model combination techniques. The detailed results of this comparison are presented in Table 6, thus enabling a quantitative evaluation of the Zimmerman–Zysno approach in the context of tourism sentiment analysis.

Table 6

Comparison of different solution variants.

Metrics	Average Voting	ELTSA $\gamma = 0.2$	ELTSA $\gamma = 0.5$	ELTSA $\gamma = 0.8$
Track Score	0.562	0.537	0.558	0.565
Macro F1 (Polarity)	0.528	0.493	0.525	0.537
Macro F1 (Type)	0.908	0.879	0.902	0.901
Macro F1 (Town)	0.468	0.452	0.466	0.472
Accuracy (Polarity)	60.45%	53.68%	59.66%	62.73%
Accuracy (Type)	91.13%	88.13%	90.46%	90.44%
Accuracy (Town)	53.48%	51.90%	53.31%	54.21%

The results demonstrate that the Zimmerman–Zysno operator with $\gamma = 0.8$ achieves the best overall performance across most key metrics, particularly in polarity classification and the identification of Mexican Magical Towns. It outperforms both the classic average voting ensemble and other configurations of the operator. Although average voting showed a slight advantage in the task of establishment type classification, the versatility of the Zimmerman–Zysno approach with $\gamma = 0.8$ is reflected in its robustness to handle data imbalance and capture complex linguistic nuances. In Table 7, we compare the performance of our best solution against the top result reported in the Rest-Mex 2025 competition.

Table 7

Comparison between our solution and the top-performing result in the Rest-Mex 2025 competition.

Metric	ELTSA $\gamma = 0.8$	UDENAR	Difference
Track Score	0.565	0.725	-0.160
Macro F1 (Polarity)	0.537	0.644	-0.107
Macro F1 (Type)	0.901	0.987	-0.086
Macro F1 (Town)	0.472	0.692	-0.220
Accuracy (Polarity)	62.73	78.52	-15.79
Accuracy (Type)	90.44	98.82	-8.38
Accuracy (Town)	54.21	77.03	-22.82

Although our solution shows competitive performance in the classification of establishment types, there is a significant gap in the polarity and geographical location tasks, especially in the identification of *Pueblos Mágicos*. These differences may be attributed to limitations in the size of the data subset used for model training. Table 8 presents a comparison between our best-performing solution and the average result of the Rest-Mex 2025 competition. This comparison provides context for assessing our approach in relation to other participating systems, highlighting its competitive advantages.

The results show that our solution consistently outperforms the average of the competition across all

Table 8

Comparison of the solution with the average result of the competition.

Metric	ELTSA $\gamma = 0.8$	Average Result	Baseline
Track Score	0.565	0.476	0.009
Macro F1 (Polarity)	0.537	0.441	0.158
Macro F1 (Type)	0.901	0.788	0.196
Macro F1 (Town)	0.472	0.394	0.008
Accuracy (Polarity)	62.73	61.36	65.53
Accuracy (Type)	90.44	80.38	41.86
Accuracy (Town)	54.21	48.26	21.79

evaluated metrics, with particularly notable improvements in the classification of establishment types and the overall *Track Score*. This superior performance validates the effectiveness of the Zimmerman-Zysno operator for integrating multiple Transformer models, especially in tasks where tourism-related language presents high variability. The positive differences in polarity and geographic location reinforce that our approach achieves a balance between generalization and adaptation to specific domains, even under limited computational resources.

5. Conclusions

This work explored the potential of Transformer models combined through the Zimmerman-Zysno operator for sentiment analysis in tourist reviews related to Mexico’s Magical Towns. The proposed solution proved effective, outperforming the average performance of the Rest-Mex 2025 competition across all evaluated metrics—particularly in the classification of establishment types and in the overall *Track Score*.

Although a gap remains compared to the top-performing result of the competition—especially in the identification of *Pueblos Mágicos*—the Zimmerman-Zysno approach with $\gamma = 0.8$ stood out for its ability to handle data imbalance and capture complex linguistic nuances. This validates its applicability in highly subjective domains such as tourism.

This work not only contributes to the advancement of Natural Language Processing for Spanish in specific contexts, but also offers a reproducible framework for future opinion analysis applications.

6. Acknowledgments

This work has been partially supported by the National Program of Science and Technology PN223LH004: Automatics, Robotics, Artificial Intelligence, of the Ministry of Science, Technology and Environment, under grant project PN223LH004-038.

Declaration on Generative AI

We declare that the present manuscript has been written entirely by the authors and that no generative artificial intelligence tools were used in its preparation, drafting, or editing.

References

- [1] B. Liu, L. Zhang, A survey of opinion mining and sentiment analysis, in: Mining text data, Springer, 2012, pp. 415–463.
- [2] B. A. Sparks, V. Browning, The impact of online reviews on hotel booking intentions and perception of trust, *Tourism Management* 32 (2011) 1310–1323.

- [3] Á. Díaz-Pacheco, R. Guerrero-Rodríguez, M. Á. Álvarez-Carmona, A. Y. Rodríguez-González, R. Aranda, A comprehensive deep learning approach for topic discovering and sentiment analysis of textual information in tourism, *Journal of King Saud University - Computer and Information Sciences* 35 (2023) 101746. URL: <http://dx.doi.org/10.1016/j.jksuci.2023.101746>. doi:10.1016/j.jksuci.2023.101746.
- [4] R. Guerrero-Rodríguez, M. A. Álvarez Carmona, R. Aranda, A. P. López-Monroy, Studying online travel reviews related to tourist attractions using nlp methods: the case of guanajuato, mexico, *Current Issues in Tourism* 26 (2023) 289–304. URL: <https://doi.org/10.1080/13683500.2021.2007227>. doi:10.1080/13683500.2021.2007227. arXiv:<https://doi.org/10.1080/13683500.2021.2007227>.
- [5] E. Olmos-Martínez, M. Á. Álvarez-Carmona, R. Aranda, A. Díaz-Pacheco, What does the media tell us about a destination? the cancan case, seen from the usa, canada, and mexico, *International Journal of Tourism Cities* 10 (2023) 639–661. URL: <http://dx.doi.org/10.1108/IJTC-09-2022-0223>. doi:10.1108/ijtc-09-2022-0223.
- [6] M. A. Álvarez-Carmona, A. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, L. Bustio-Martínez, V. Herrera-Semenets, *Panorama de rest-mex en iberlef 2025: Investigando la evaluación del sentimiento en textos para pueblos mágicos mexicanos*, volume 75, 2025.
- [7] J. Á. González-Barba, L. Chiruzzo, S. M. Jiménez-Zafra, 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.
- [8] M. Á. Álvarez-Carmona, R. Aranda, S. Arce-Cárdenas, D. Fajardo-Delgado, R. Guerrero-Rodríguez, A. P. López-Monroy, J. Martínez-Miranda, H. Pérez-Espinosa, A. Rodríguez-González, Overview of rest-mex at iberlef 2021: Recommendation system for text mexican tourism, *Procesamiento del Lenguaje Natural* 67 (2021). doi:<https://doi.org/10.26342/2021-67-14>.
- [9] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, D. Fajardo-Delgado, R. Guerrero-Rodríguez, L. Bustio-Martínez, Overview of rest-mex at iberlef 2022: Recommendation system, sentiment analysis and covid semaphore prediction for mexican tourist texts, *Procesamiento del Lenguaje Natural* 69 (2022).
- [10] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, L. Bustio-Martínez, V. Muñis-Sánchez, A. P. Pastor-López, F. Sánchez-Vega, Overview of rest-mex at iberlef 2023: Research on sentiment analysis task for mexican tourist texts, *Procesamiento del Lenguaje Natural* 71 (2023).
- [11] A. Vaswani, et al., Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [12] 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, A. 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. URL: <https://www.sciencedirect.com/science/article/pii/S1319157822003615>. doi:<https://doi.org/10.1016/j.jksuci.2022.10.010>.
- [13] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, 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*, volume 1, 2019, pp. 4171–4186.
- [14] K. L. Tan, C.-P. Lee, K. S. M. Anbananthan, K. M. Lim, Roberta-lstm: a hybrid model for sentiment analysis with transformer and recurrent neural network, *IEEE Access* 10 (2022) 21517–21525.
- [15] H. J. Zimmermann, P. Zysno, Latent connectives in human decision making, *Fuzzy sets and systems* 4 (1980) 37–51.
- [16] M. Oliveira, Applied to neuro-fuzzy models for signal validation in Angra 1 nuclear power plant, Technical Report INIS-BR-12365, Universidade Federal do Rio de Janeiro (COPPE/UFRJ), 1999.
- [17] V. Dogra, S. Verma, Kavita, P. Chatterjee, J. Shafi, J. H. Choi, M. F. Ijaz, A complete process

of text classification system using state-of-the-art nlp models, Computational Intelligence and Neuroscience 2022 (2022) 1883698.

- [18] C. Sun, X. Qiu, Y. Xu, X. Huang, How to fine-tune bert for text classification?, in: China National Conference on Chinese Computational Linguistics, Springer, 2019, pp. 194–206.
- [19] H. Wang, X. Hu, H. Zhang, Sentiment analysis of commodity reviews based on albert-lstm, Journal of Physics: Conference Series 1651 (2020) 012022.

Online Resources

The results and official rankings of the shared task can be accessed through the following link:

- [Rest-Mex 2025 Results](#)