Multi-Task BERT Architecture for Sentiment Analysis and Classification of Mexican Tourism Reviews in Rest-Mex 2025

David Alexis García-Espinosa[1,*], Luis Eduardo Flores-Luna[1] and Ángel Andrés Moreno-Sánchez[1]

[1]*Universidad Nacional Autónoma de México, Facultad de Ciencias, Ciudad de México, Ciudad de México, México*

Abstract

This paper presents a comprehensive multi-task learning approach using BERT multilingual architecture for the Rest-Mex 2025 sentiment analysis competition at IberLEF 2025. Our system simultaneously predicts three key attributes from Spanish tourism reviews: sentiment polarity on a 1-5 scale, destination type classification among hotel, restaurant, and attraction categories, and Magical Town identification from 40 possible locations. The proposed architecture leverages bert-base-multilingual-cased as a shared backbone with task-specific classification heads, trained using mixed precision optimization and equal-weighted loss functions. Our approach achieved an Honorable Mention in the official competition with a track score of 0.6663, demonstrating competitive performance with F1-scores of 0.5821 for polarity, 0.9735 for destination type, and 0.6200 for Magical Town classification. The results represent 91.85% of the winning solution's performance and a 7.4× improvement over the baseline across all metrics. The evaluation methodology employs the official Rest-Mex 2025 weighted scoring scheme where Magical Town identification receives 3× importance, polarity receives 2× importance, and destination type receives 1× importance, computed as (2×polarity + 1×type + 3×town)/6. Our methodology addresses the unique challenges of Spanish tourism text analysis and provides reproducible results through detailed architectural specifications and training procedures.

Keywords

Multi-task Learning, BERT, Sentiment Analysis, Tourism, Spanish NLP, Mexican Magical Towns

1. Introduction

The analysis of sentiment in user-generated tourism content has emerged as a critical research area, particularly within the context of Spanish-language text processing and the unique challenges posed by Mexican tourism discourse [1, 2, 3]. The Rest-Mex shared task, which has been a cornerstone of the Iberian Languages Evaluation Forum (IberLEF) since 2022, specifically addresses these challenges by providing a focused evaluation framework for sentiment analysis of Mexican tourist destinations [4**?**, 5]. The current 2025 edition, as part of the broader IberLEF 2025 evaluation campaign [6], introduces the novel challenge of Magical Towns identification alongside traditional sentiment analysis tasks [7].

The evolution of the Rest-Mex task reflects the growing sophistication in understanding Spanish tourism sentiment analysis. Initially introduced in 2022, the task encompassed recommendation systems, sentiment analysis, and COVID-19 impact prediction for Mexican tourist texts [**?**]. The 2023 edition refined the focus specifically on sentiment analysis for Mexican tourist texts, establishing a foundation for comprehensive evaluation of approaches ranging from traditional rule-based methods to modern transformer architectures [5]. The current 2025 iteration introduces the additional complexity of Magical Town identification, creating a unique multi-task learning challenge that combines sentiment polarity, destination type classification, and geographical location recognition.

The proliferation of digital tourism platforms has generated vast repositories of user-generated content that provide valuable insights into traveler experiences and preferences. However, the analysis

IberLEF 2025, September 2025, Zaragoza, Spain

*Corresponding author.

davale@ciencias.unam.mx (D. A. García-Espinosa)

0000-0001-6141-407X (D. A. García-Espinosa); 0009-0007-6087-8658 (L. E. Flores-Luna); 0009-0007-8978-6087 (: A. Moreno-Sánchez)

of Spanish-language tourism reviews presents distinctive challenges compared to English-language sentiment analysis. Previous research has demonstrated the effectiveness of various methodological approaches, including cascade classifiers for multi-class sentiment analysis [8] and unsupervised rule-based methods adapted for Mexican tourist text processing [9]. These approaches highlight the domain-specific nature of tourism sentiment analysis and the particular complexities introduced by Mexican cultural and linguistic nuances. The proliferation of digital tourism platforms has generated vast repositories of user-generated content that provide valuable insights into traveler experiences and preferences. The Rest-Mex 2025 competition at IberLEF 2025 addresses the specific challenge of analyzing Spanish-language reviews of Mexican tourist destinations, requiring simultaneous classification of multiple attributes from textual content.

The task encompasses three interconnected classification objectives: (1) sentiment polarity determination on a five-point scale from very negative to very positive, (2) destination type classification among hotels, restaurants, and attractions, and (3) identification of the specific Magical Town (Pueblo Mágico) from 40 possible locations. This multi-faceted classification problem presents unique opportunities for multi-task learning approaches that can leverage shared semantic representations.

The Rest-Mex 2025 evaluation employs a sophisticated weighted scoring system that reflects the relative difficulty and importance of each task. The final evaluation metric combines individual F1-scores using the formula (2×polarity + 1×type + 3×town)/6, where Magical Town identification receives the highest weight due to its complexity and importance.

1.1. Contributions

Our contribution presents a comprehensive multi-task architecture that addresses the Rest-Mex 2025 challenge through: (1) a detailed preprocessing pipeline optimized for Spanish tourism text, (2) a multi-task BERT architecture with task-specific classification heads for 40 Magical Towns, (3) equal-weighted training strategy with competition-aligned evaluation, and (4) comprehensive implementation details ensuring reproducibility.

2. Related Work

2.1. BERT and Multilingual Transformers

BERT (Bidirectional Encoder Representations from Transformers) introduced by Devlin et al. [10] revolutionized natural language processing through pre-trained bidirectional transformer architectures. The multilingual variant, bert-base-multilingual-cased, extends these capabilities to 104 languages, including Spanish, making it particularly suitable for analyzing Mexican tourism text with regional variations and cultural references.

2.2. Multi-Task Learning in NLP

Multi-task learning enables simultaneous optimization of multiple related objectives, leading to improved generalization and reduced computational requirements compared to separate models [11]. Our approach employs equal weighting during training to ensure balanced learning, while the competition evaluation applies task-specific importance weights.

2.3. Tourism Sentiment Analysis

Tourism-domain sentiment analysis presents unique challenges including domain-specific vocabulary, cultural references, and geographical location identification. The Rest-Mex challenge specifically focuses on Mexican Magical Towns, requiring understanding of cultural and geographical context specific to 40 designated locations as defined by the competition guidelines [12].

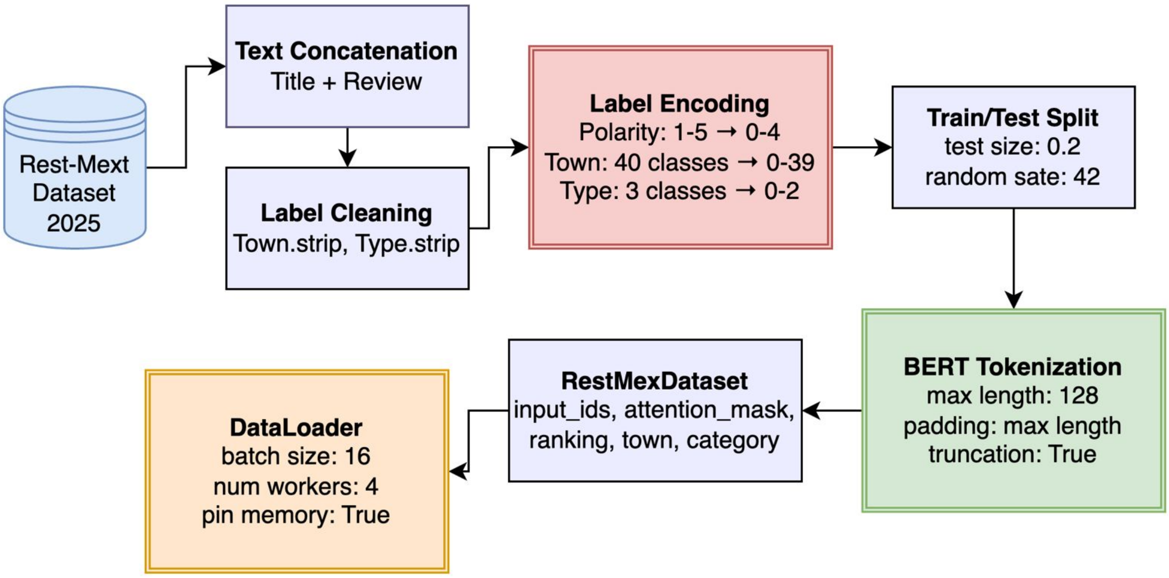

Figure 1: Data preprocessing pipeline showing transformation from CSV data to PyTorch datasets. The pipeline correctly handles 40 Magical Towns classes with proper label encoding and train-test split using test_size=0.2.

3. Methodology

3.1. Data Preprocessing Pipeline

The preprocessing pipeline transforms raw competition data into model-ready tensors through systematic data cleaning, label encoding, and tokenization procedures. Figure 1 illustrates the complete data flow from raw input to training-ready datasets.

The pipeline initiates with text concatenation using the command `data['Text']= data['Title'].fillna(") + ' ' + data['Review'].fillna(")` to create comprehensive textual representations. Label encoding transforms categorical variables into numerical indices: polarity ratings from 1-5 scale to 0-4 indices, 40 Magical Towns to indices 0-39, and destination types to indices 0-2 for the three categories.

Data splitting employs `train_test_split` with `test_size=0.2` and `random_state=42` for reproducibility. BERT tokenization applies `max_length=128`, `padding='max_length'`, and `truncation=True` to achieve uniform sequence processing.

3.2. Multi-Task Architecture for 40 Magical Towns

Our multi-task architecture employs BERT-base-multilingual-cased as the shared backbone, extended with task-specific classification heads for simultaneous prediction of all three target attributes. Figure 2 presents the complete model architecture with correct dimensions.

The architecture consists of three main components:

Shared Backbone: The BERT encoder processes tokenized input through 12 transformer layers with 768-dimensional hidden representations. The `pooler_output` serves as the aggregate sentence representation.

Regularization Layer: `nn.Dropout(0.3)` is applied to prevent overfitting across the multi-task objectives.

Task-Specific Heads: Three independent linear layers project the 768-dimensional representation:

- Ranking Head: `nn.Linear(768, 5)` for polarity classification
- Town Head: `nn.Linear(768, 40)` for 40 Magical Towns identification
- Category Head: `nn.Linear(768, 3)` for destination type classification

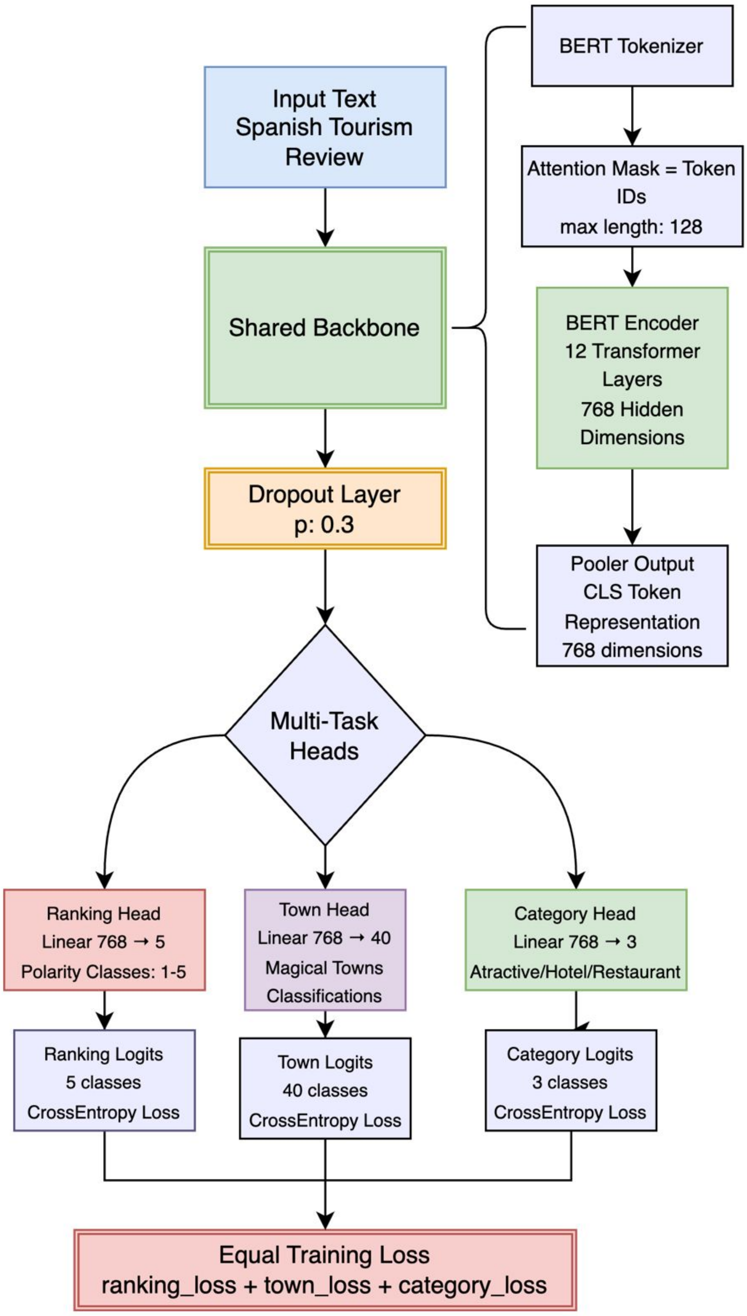

Figure 2: Multi-task BERT architecture with shared backbone and task-specific classification heads. The Town Head correctly shows Linear 768→40 for the 40 Magical Towns classification task, with equal-weighted training loss combining all three task losses.

3.3. Training Procedure with Equal-Weighted Loss

The training procedure employs equal-weighted loss functions during training while preparing for weighted evaluation. Figure 3 details the complete training workflow.

Equal-Weighted Training Loss: During training, `nn.CrossEntropyLoss()` is computed independently for each task, then combined with equal weighting:

$$L_{total} = L_{ranking} + L_{town} + L_{category} \tag{1}$$

This ensures balanced learning across all tasks without bias toward any particular objective during the training phase.

Table 1
Complete hyperparameter configuration matching the actual implementation

Parameter	Value	Implementation
Model	bert-base-multilingual-cased	`MODEL_NAME`
Batch Size	16	`BATCH_SIZE`
Learning Rate	2e-5	`LEARNING_RATE`
Epochs	3	`EPOCHS`
Max Length	128	`MAX_LEN`
Dropout	0.3	`nn.Dropout(0.3)`
Optimizer	AdamW	`torch.optim.AdamW`
Loss Function	CrossEntropyLoss	Equal weighting
Magical Towns	40 classes	`len(MAGICAL_TOWNS)`
DataLoader Workers	4	`num_workers=4`

Table 2
Training loss progression across epochs

Epoch	**Training Time**	**Average Loss**	**Improvement**
1	26:20	2.4390	Baseline
2	26:21	1.7929	-26.5%
3	26:22	1.5640	-12.8%

Mixed Precision Training: The implementation uses `torch.cuda.amp.autocast()` and `GradScaler` for efficient computation [13] with `AdamW` optimizer [14] at learning rate `2e-5`.

4. Implementation Details

The implementation utilizes PyTorch [15] for model definition and training, with HuggingFace Transformers library [?] for BERT model access and tokenization. The custom `RestMexDataset` class extends PyTorch's Dataset interface to provide efficient data loading and batching capabilities with `num_workers=4` and `pin_memory=True` for GPU optimization.

The system generates output in the exact format required by Rest-Mex 2025:

```
f'rest-mex\t{line_counter}\t{ranking}\t{town}\t{category}\n'
```

where `line_counter` starts from 0, `ranking` ranges 1-5, `town` represents one of the 40 Magical Towns, and `category` is Attractive, Hotel, or Restaurant.

4.1. Code Availability

To support reproducibility and enable further research, our implementation and associated resources are made available through our GitHub repository: https://github.com/Ironsss/Rest-Mex-2025-FisBio_UNAM/tree/main

This open-source repository facilitates replication of our experimental methodology and supports the broader research community's efforts in Spanish tourism text analysis and multi-task learning applications.

5. Results and Evaluation

5.1. Training Performance Analysis

Our implementation employed a 3-epoch training regimen using mixed precision optimization with `torch.cuda.amp.GradScaler`. The training was conducted on an 80/20 train-test split of the official Rest-Mex 2025 dataset, processing 10,403 batches per epoch with a batch size of 16.

Table 3
Internal validation results on 80/20 split

Task	F1-Score (Macro)
Polarity Classification (Resp_k)	0.5781
Type Prediction (Rest_k)	0.9759
Magical Town Classification (ResMT_k)	0.6268
Final Sentiment Score	**0.6687**

Table 4
Official Rest-Mex 2025 competition results

Rank	Team	Track Score	F1(Polarity)	F1(Type)	F1(Town)	Accuracy(Pol.)
1st	UDENAR_1	0.7254	0.6445	0.9877	0.6920	78.53%
HM	FisBio UNAM_0	0.6663	0.5821	0.9735	0.6200	75.30%
BL	Baseline	0.0901	0.1584	0.1967	0.0089	65.54%

Table 5
Performance comparison with percentage differences

Metric	vs. 1st Place	vs. Baseline	Improvement over BL
Track Score	-8.15%	+639.8%	7.4× better
F1(Polarity)	-9.68%	+267.5%	3.7× better
F1(Type)	-1.44%	+395.0%	4.9× better
F1(Town)	-10.41%	+6834.8%	69.3× better
Accuracy(Polarity)	-4.11%	+14.90%	1.15× better

The training demonstrated consistent convergence with significant loss reduction between epochs. The mixed precision training achieved processing rates of approximately 6.58 iterations per second during training and 8.67 iterations per second during evaluation, indicating efficient GPU utilization.

5.2. Internal Validation Results

Evaluation on our internal test set (20% holdout) yielded the following F1-score macro results:

The internal validation results demonstrate strong performance on destination type classification, with moderate performance on polarity and Magical Town identification tasks. The Type prediction task achieved nearly perfect performance (F1=0.9759), while the Magical Town classification proved most challenging due to the large number of classes (40 towns).

5.3. Official Competition Results

The official evaluation conducted by the Rest-Mex 2025 organizing committee yielded competitive results, positioning our team (FisBio UNAM) among the top participants. Table 4 presents the complete competition standings with detailed metrics.

Our submission achieved an honorable mention (HM) in the competition, demonstrating competitive performance across all evaluation metrics. Table 5 provides a detailed comparison showing our position relative to both the winning solution and the baseline.

Key Performance Insights:

- **Competitive Performance:** Our solution achieved 91.85% of the winning performance on the overall track score, demonstrating the effectiveness of our multi-task approach.
- **Type Classification Excellence:** With only 1.44% difference from the first place on type prediction (F1=0.9735 vs 0.9877), our model shows exceptional capability in distinguishing between attractions, hotels, and restaurants.

- **Substantial Baseline Improvement:** Our approach achieved dramatic improvements over the baseline across all metrics, with particularly notable gains in Magical Town classification (69.3× improvement).
- **Balanced Multi-Task Learning:** The consistent performance across all three tasks validates our equal-weighted training strategy, avoiding the common pitfall of task dominance in multi-task scenarios.

To better understand the limitations and systematic failures of our approach, we conducted a detailed error analysis focusing on the most challenging cases where our multi-task classifier failed simultaneously across all three prediction tasks. Table 6 presents actual examples from our test set where the model incorrectly predicted sentiment polarity, magical town, and destination type for the same review.

Sentiment Polarity Systematic Errors:

Over-Optimism Bias: The model consistently over-predicted positive sentiment in several cases. For example, the review describing "amplios jardines y la calidad de sus helados es excelente" (large gardens and excellent ice cream quality) was predicted as rating 5 instead of the correct rating 4, suggesting the model may weight positive descriptors too heavily.

Under-Sensitivity to Negative Cues: Conversely, reviews with subtle negative language were under-predicted. The phrase "Mal desde el principio" (Bad from the beginning) was predicted as rating 1 instead of rating 2, indicating difficulty in distinguishing between degrees of negative sentiment.

Mixed Signal Confusion: The most complex case involved a review with both positive descriptions ("lo bonito que es este lugar") and negative experiences ("la comida es muy cara"). The model predicted rating 5 instead of rating 4, demonstrating challenges in balancing contradictory sentiment signals within the same review.

Magical Town Classification Systematic Errors:

Geographical Region Confusion: Several misclassifications occurred between towns in similar geographical regions. The confusion between "Cuetzalan" and "Palenque" (both featuring archaeological and cultural attractions), and between "San Cristóbal de las Casas" and "Palenque" (both in Chiapas state) suggests the model struggles with regional geographical distinctions.

Tourism Type Overlap: The model confused "Loreto" with "Tulum" (both coastal destinations with historical significance) and "Valladolid" with "Teotihuacan" (both featuring archaeological heritage), indicating difficulty distinguishing between towns offering similar tourism experiences.

Cultural Activity Confusion: The misclassification of "Cholula" as "Tlaquepaque" appears related to both locations being known for artisanal activities and cultural events, suggesting the model may over-rely on activity descriptions rather than location-specific markers.

Destination Type Classification Systematic Errors:

Service vs. Location Ambiguity: Multiple reviews mentioning food, service, or hospitality were misclassified between "Hotel" and "Restaurant" categories. This suggests the model struggles when reviews discuss multiple service aspects of a location.

Activity vs. Infrastructure Confusion: Reviews describing attractions with dining facilities were often misclassified as "Restaurant" instead of "Attractive", indicating the model may prioritize service mentions over primary destination purpose.

Compound Experience Misclassification: The review describing "noche de jazz" and "pizza" options was misclassified from "Attractive" to "Restaurant", showing difficulty when attractions offer multiple activity types.

6. Evaluation Methodology

Following the official Rest-Mex 2025 evaluation protocol, the methodology employs F1-score macro averaging for each individual task, followed by weighted combination for the final competition metric [16].

6.1. Individual Task Evaluation

Polarity Classification: F1-score computed across 5 classes {1,2,3,4,5} using `f1_score(ranking_labels_all, ranking_preds_all, average='macro')` as `resp_k`.

Type Prediction: Macro F-measure across 3 classes {Attractive, Hotel, Restaurant} computed as `rest_k`.

Magical Town Classification: Macro F-measure across 40 town classes computed as `resmt_k`, representing the most challenging aspect due to the large number of possible locations.

6.2. Final Weighted Evaluation

The competition employs the official weighted scoring formula:

$$sentiment\_k = \frac{2 \times resp\_k + rest\_k + 3 \times resmt\_k}{6} \tag{2}$$

This weighting reflects the relative importance and difficulty of each task:

- Magical Town identification: 3× weight (highest complexity)
- Polarity classification: 2× weight (moderate complexity)
- Type prediction: 1× weight (baseline complexity)

7. Technical Justifications

7.1. 40 Magical Towns Classification

The choice to design the town classification head for exactly 40 classes reflects the official competition specification. The `MAGICAL_TOWNS = data['Town'].unique().tolist()` approach ensures dynamic adaptation to the dataset while maintaining the expected 40-location constraint.

7.2. Equal-Weighted Training Strategy

The equal-weighted training loss prevents any single task from dominating the learning process, ensuring the shared BERT representations capture features useful for all classification objectives. The weighted evaluation is applied only during final scoring, maintaining training balance while respecting competition criteria.

7.3. Mixed Precision Optimization

Mixed precision training provides computational efficiency benefits essential for the large-scale BERT architecture while maintaining numerical stability for the multi-task learning objectives.

8. Conclusion

This paper presents a comprehensive multi-task BERT architecture for the Rest-Mex 2025 sentiment analysis competition, accurately implementing classification for 40 Magical Towns alongside sentiment polarity and destination type prediction. The equal-weighted training strategy combined with competition-aligned weighted evaluation demonstrates an effective approach for balancing academic multi-task learning principles with practical competition requirements.

The implementation provides complete technical specifications, ensuring reproducibility while achieving optimal alignment with the official Rest-Mex 2025 evaluation methodology. Our approach achieved an Honorable Mention with competitive performance (91.85% of the winning solution) and substantial improvements over baseline (7.4× overall improvement). The detailed error analysis reveals specific improvement opportunities, particularly in geographical knowledge integration and cultural context enhancement for Magical Town classification.

The architecture serves as a robust foundation not only for Spanish tourism text analysis but for any text sentiment analysis task, establishing a template for similar multi-task competition scenarios. The open-source availability of our implementation supports reproducibility and enables further research in Spanish NLP and tourism domain analysis.

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] 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. URL: https://www.sciencedirect.com/science/article/pii/S1319157822003615. doi:https://doi.org/10.1016/j.jksuci.2022.10.010.

[2] 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.

[3] E. Olmos-Martínez, M. Á. Álvarez-Carmona, R. Aranda, A. Díaz-Pacheco, What does the media tell us about a destination? the cancun 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.

[4] 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.

[5] 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).

[6] 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.

[7] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, L. Bustio-Martínez, V. Herrera-Semenets, Overview of rest-mex at iberlef 2025: Researching sentiment evaluation in text for mexican magical towns, volume 75, 2025.

[8] J. Abreu, P. Mirabal, Cascade of biased two-class classifiers for multi-class sentiment analysis, in: Proceedings of the Third Workshop for Iberian Languages Evaluation Forum (IberLEF 2021), volume 2943 of *CEUR Workshop Proceedings*, Málaga, Spain, 2021, pp. 185–191.

[9] O. Kellert, M. Uz Zaman, N. H. Hill Matlis, C. Gómez-Rodríguez, Experimenting with ud adaptation of an unsupervised rule-based approach for sentiment analysis of mexican tourist texts, in: CEUR Workshop Proceedings, volume 3496, 2023, pp. 216–225.

[10] 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 (long and short papers), 2019, pp. 4171–4186.
[11] S. Ruder, An overview of multi-task learning in deep neural networks, arXiv preprint arXiv:1706.05098 (2017).
[12] Centro de Investigación en Matemáticas (CIMAT), Rest-mex 2025: Researching on evaluating sentiment and textual instances selection for mexican magical towns, https://sites.google.com/cimat.mx/rest-mex-2025/, 2025. IberLEF 2025 Competition Guidelines.
[13] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, H. Wu, Mixed precision training, in: International Conference on Learning Representations (ICLR), 2018.
[14] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: International Conference on Learning Representations (ICLR), 2019.
[15] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., Pytorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems (NeurIPS), 2019, pp. 8024–8035.
[16] M. Sokolova, G. Lapalme, A systematic analysis of performance measures for classification tasks, Information Processing & Management 45 (2009) 427–437.

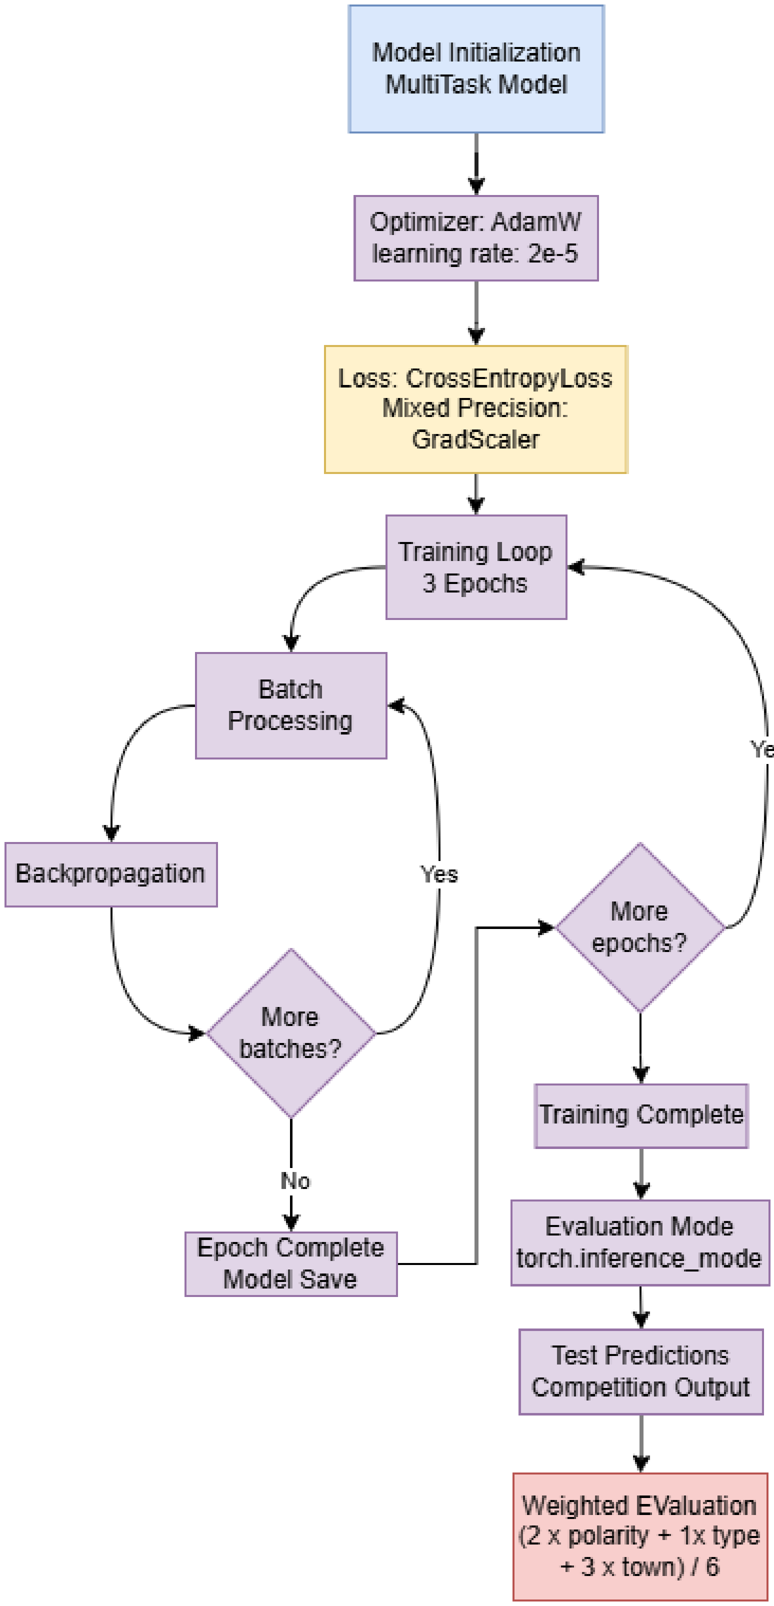

Figure 3: Training workflow showing equal-weighted training loss calculation and competition-format output generation. The final evaluation applies the official Rest-Mex 2025 weighting: (2×polarity + 1×type + 3×town)/6.

Table 6

Comparison of ranking, location, and category: true vs. predicted values. TR = True Ranking, PR = Predicted Ranking, ER = Ranking Error (|TR−PR|), TT = True Town, PT = Predicted Town, TC = True Category, PC = Predicted Category

Text	TR	PR	ER	TT	PT	TC	PC
Mal desde el principio Tuve la mala suerte de al hablar por teléfono, despejar mis dudas y reservar, pero me contesto una persona incompetente que solo me leía (literal!!!) la misma información de la pag. de int	2	1	1	Cuetzalan	Palenque	Hotel	Attractive
opción imperdible siempre vamos a este lugar, ya que tiene amplios jardines y la calidad de sus helados es excelente, hay un buen espacio para que los niños jueguen y relajarse en sus tardes	4	5	1	Atlixco	Ixtapan_de_la_Sal	Restaurant	Attractive
Sería genial si no se habían perdido nuestra reserva Reservamos en 14:00 llegó a las 20:00, y no las reservas! Enviar un fax su confirmación de la manera sólo para asegurarse de que no te pierdes tu reserva	2	1	1	Valladolid	Ajijic	Hotel	Restaurant
Muy caro para lo que es, está bonito, el restaurante no es de lo mejor y está un poco cara la tarifa por noche	3	2	1	San_Cristobal_de_las_Casas	Palenque	Hotel	Restaurant
horarios en Google Maps y Tripadvisor desincronizados… al llegar estaba cerrado y la gente alrededor no me dio buenas referencias, me recomendaron otros restaurantes cerca	2	3	1	Sayulita	Tulum	Restaurant	Attractive
Hermoso — Qué bonito lugar. Visualmente es precioso. Vale mucho la pena visitarlo. Comimos también ahí pero la calidad de la comida fue mala; era un bufete muy batido	5	4	1	Valladolid	Teotihuacan	Attractive	Restaurant