

Bio-inspired Synthetic Review Generation via Physicochemical Model for REST-MEX 2026

Estefania Lazcano-Díaz^{1,*}, Miguel Angel Álvarez-Carmona²

¹*Centro de Investigación y de Estudios Avanzados (Cinvestav), Unidad Monterrey*

²*Centro de Investigación en Matemáticas A.C., Sede Monterrey*

Abstract

This paper presents a bio-inspired approach for synthetic tourist review generation submitted to the REST-MEX 2026 Artificial Data Generation Track. The proposed method combines concepts from chemical chain reactions and crystal growth to construct synthetic reviews through the composition and progressive expansion of linguistic units. Unlike prompt-based generation using Large Language Models, the proposed framework follows explicit generation rules, providing a transparent and interpretable mechanism for data augmentation. Experimental results show that the generated synthetic corpus outperformed the Majority Class baseline under the official evaluation metrics, demonstrating the feasibility of bio-inspired text generation as an alternative strategy for data-centric sentiment analysis.

Keywords

Sentiment Analysis, Synthetic Text Generation, Bio-inspired Algorithms, Chemical Reaction Model, Crystal Growth, Tourism Reviews, REST-MEX 2026, REST-MEX 2026

1. Introduction

One of the defining properties of natural and artificial intelligent systems is their capacity to generate complex structures from relatively simple components. This principle is widely observed in self-organizing natural processes, where global organization emerges from local interactions rather than from centralized control. Examples include range from molecular reactions that produce increasingly complex compounds to crystal formation, where ordered structures progressively develop from microscopic nuclei. These natural phenomena have inspired several computational approaches in optimization, machine learning, and artificial intelligence, suggesting that biological and physicochemical processes can provide valuable insights for solving computational problems [1].

Natural Language Processing (NLP) faces a similar challenge. Human language is itself a complex system composed of words, phrases, and semantic relationships that interact to form coherent discourse. Consequently, the generation of synthetic text can also be viewed as an emergent process in which small linguistic units are combined according to structural and semantic constraints. While recent years have been dominated by Large Language Models (LLMs), alternative generation paradigms inspired by natural systems remain comparatively underexplored [2].

Within tourism research, NLP has become a key tool for analyzing the vast volume of textual information generated by travelers. Online reviews, comments, and recommendations are valuable sources for understanding visitor experiences, destination image, service quality, and customer satisfaction. These data have supported applications ranging from sentiment analysis and destination image assessment to recommendation systems, hospitality assistants, and tourism intelligence platforms [3, 4, 5, 6, 7].

Among these applications, sentiment analysis has received particular attention, as it allows for the automatic identification of tourists' opinions at a scale difficult to achieve through manual inspection. Deep learning architectures and transformer-based models have substantially improved classification performance in recent years, leading to increasingly sophisticated solutions for tourism opinion mining

IberLEF 2026, September 2026, León, Spain

*Corresponding author.

✉ estefania.lazcano@cimat.mx (E. Lazcano-Díaz); miguel.alvarez@cimat.mx (M. A. Álvarez-Carmona)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[8, 9]. However, regardless of the classification architecture employed, the quality and representativeness of the training data remain fundamental factors influencing model performance.

Recognizing the growing importance of data-centric artificial intelligence, the REST-MEX assessment campaign has progressively evolved from traditional classification tasks to more complex research problems. Previous editions addressed recommendation systems, sentiment analysis, and multi-task learning on Mexican tourism reviews, establishing benchmark datasets that have become valuable resources for the Spanish NLP community [10, 11, 12, 13]. The 2026 edition introduces a substantially different perspective: instead of asking participants to improve the classifier itself, the Artificial Data Generation Track assesses the ability to generate synthetic tourism reviews that improve the performance of a fixed sentiment analysis model [14, 15].

Many existing approaches to synthetic text generation rely on prompting Large Language Models. Although these models produce highly fluent text, their generation process often remains difficult to interpret, trace, and control. In contrast, we investigate whether principles inspired by natural self-organization can provide an alternative mechanism for constructing coherent synthetic reviews. Rather than generating complete opinions, the proposed approach models textual generation as a gradual construction process driven by two complementary metaphors: chemical chain reactions, in which linguistic fragments are combined according to semantic compatibility reaction rules, and crystal growth, in which complete reviews emerge through the progressive expansion of small textual nuclei.

The resulting generation process does not attempt to replicate the behavior of neural language models. Instead, it exploits compositionality as the main principle for producing domain-consistent tourism reviews. By combining reaction-based composition with crystal-inspired expansion, the proposed method aims to generate synthetic opinions while preserving semantic coherence and explicit control throughout the construction process. The generated corpus is subsequently incorporated into the official REST-MEX 2026 training dataset in order to evaluate its impact on performance of a fixed sentiment classification model.

2. Related Work

The use of computational methods in tourism research has grown considerably during the last decade. Online travel reviews, social media posts, destination images, and digital interaction traces have become valuable sources for studying how tourists perceive, describe and evaluate destinations. Several studies have shown that artificial intelligence can support the analysis of destination image, digital reputation, and visitor experience from both textual and visual perspectives [4, 16, 17, 18]. These contributions demonstrate that tourism-related data are not only useful for descriptive analysis, but also for developing predictive models and decision-support systems.

Natural Language Processing has played a central role in this transformation. Previous research has applied NLP techniques to online travel reviews to detect topics, extract opinions, classify sentiment, and identify factors associated with tourist satisfaction [3, 19, 8]. In the Mexican tourism context, these methods have been used to study attractions, hotels, destination reputation, recommendation systems, and hospitality services [6, 7, 5]. These works highlight the relevance of Spanish-language tourism corpora and the need for specialized resources that capture regional linguistic and cultural and domain-specific characteristics.

The REST-MEX evaluation campaign has contributed directly to this line of research by providing benchmark tasks focused on Mexican tourism texts. Its first editions addressed recommendation and classification problems in tourism reviews [10, 11], whereas later editions focused more specifically on sentiment analysis and the classification of Mexican tourist opinions [12, 13]. Several participant systems explored transformer-based architectures, domain adaptation, random instance selection, balancing strategies, and optimization-based approaches for polarity prediction [20, 21, 22, 23]. These works show that sentiment analysis in tourism remains as challenging task, particularly when dealing with imbalanced datasets, domain-specific vocabulary, regional linguistic variation, and the subjective nature of tourist opinions.

Class imbalance has been one of the most persistent difficulties in REST-MEX and related sentiment analysis tasks. Tourist review corpora usually contain a large number of negative or neutral examples. This imbalance is not only a numerical issue, but also a problem of representativeness, since negative and neutral reviews often capture critical tourist experiences that are less frequent but highly informative. This distribution can bias classifiers toward majority classes and reduce their ability to detect dissatisfied experiences. Previous work has explored balancing strategies and combinations of sentiment analysis systems as possible mechanisms to improve robustness [9, 23]. However, the generation of artificial instances as a direct mechanism for improving sentiment classification remains an open and relevant research problem.

The REST-MEX 2026 Artificial Data Generation Track explicitly addresses this issue by evaluating synthetic review generation rather than model design. In this setting, the objective is to create artificial tourist opinions that improve the behavior of a fixed sentiment classifier when incorporated in to the training corpus [14, 15]. This formulation is aligned with recent data-centric perspectives in machine learning, where improving the quality, diversity, and representativeness of data is considered as important as modifying the learning algorithm itself. Therefore, the track pack provides an appropriate scenario for exploring alternative data generation strategies that are not only effective, but also interpretable and controllable.

Many current synthetic text generation strategies rely on Large Language Models. Prompt-based systems have shown strong capabilities in producing fluent text, and similar approaches have been used in other NLP tasks involving hallucination detection, classification, and controlled generation [24]. Nevertheless, LLM-based generation may provide limited interpretability regarding how each synthetic example is constructed. This limitation motivates the exploration of alternative mechanisms in which the generation process can be described through explicit rules, intermediate structures, and controlled transformations.

Bioinspired and physicochemical metaphors offer a useful conceptual framework for this purpose. In particular, chemical reaction models allow textual fragments to be represented as units that combine into larger semantic structures, whereas crystallization models describe the gradual expansion of coherent forms from initial nuclei. Although these mechanisms have been more commonly studied in optimization and soft computing [1], they can also be adapted to text generation by treating phrases as compositional elements. In this work, we explore this direction by combining reaction-based composition and crystal-inspired expansion to generate synthetic tourist reviews for REST-MEX 2026.

3. Methodology

The proposed method generates synthetic tourist reviews through a symbolic bio-inspired process based on two complementary mechanisms: chemical chain reactions and crystal growth. Instead of producing complete reviews directly, the system constructs opinions by combining small linguistic units and progressively expanding them into coherent texts.

Figure 1 provides an overview of the proposed framework. The generation process consists of four main stages: linguistic molecule extraction, crystal seed selection, reaction-based composition, and crystal growth. The resulting synthetic reviews are then incorporated into the original REST-MEX 2026 training corpus and used to fine-tune the baseline sentiment classification model.

3.1. Dataset

The system was developed using the official REST-MEX 2026 training corpus, which contains 208,051 tourist reviews from Mexican *Pueblos Mágicos*. Each review is annotated with a polarity label in the range (1)–(5). The proposed approach focuses on generating synthetic reviews that preserve the vocabulary, structure, and thematic patterns observed in the original corpus.

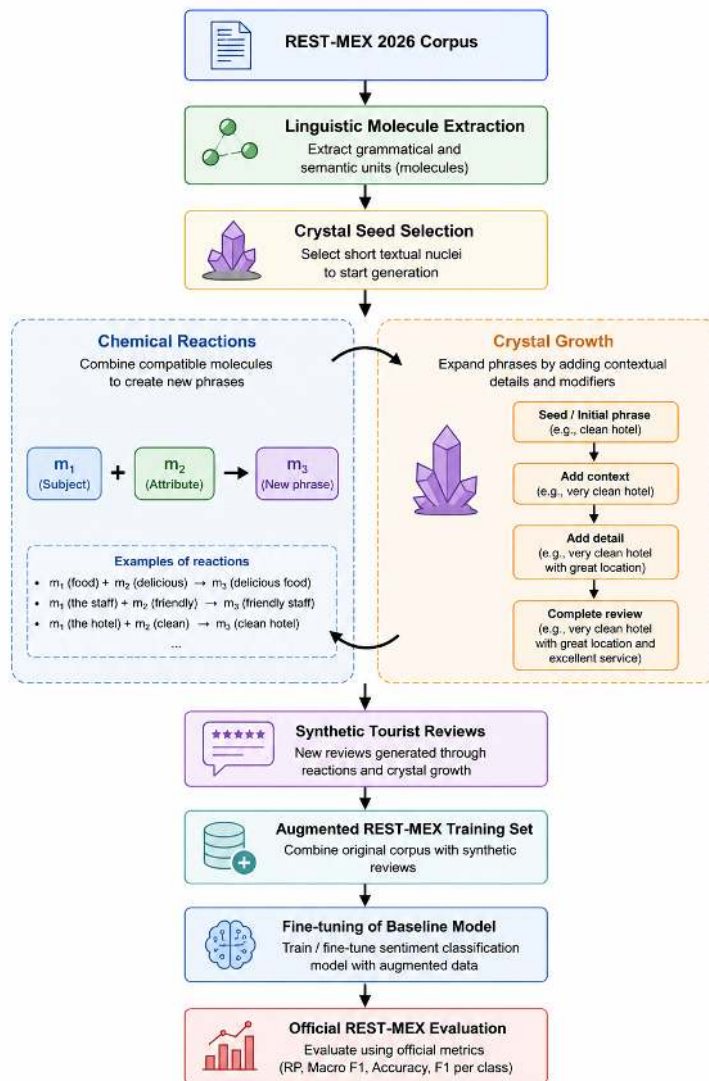


Figure 1: Overview of the proposed bio-inspired synthetic review generation framework.

3.2. Linguistic Molecules

Let $(D = (x_i, y_i)_{i=1}^N)$ be the original training corpus, where (x_i) is a tourist review and $(y_i \in 1, 2, 3, 4, 5)$ is its polarity label. Each review (x_i) is segmented into a set of short textual fragments:

$$x_i = m_{i1}, m_{i2}, \dots, m_{ik},$$

where each fragment (m_{ij}) is interpreted as a *linguistic molecule*. A linguistic molecule represents a minimal semantic unit, such as an opinion about food, service, price, cleanliness, location, comfort, or tourist attractions.

The complete molecular pool is defined as:

$$\mathcal{M} = \bigcup_{i=1}^N m_{i1}, m_{i2}, \dots, m_{ik}.$$

Each molecule is associated with a semantic type $(s(m))$, a polarity orientation $(p(m))$, and a lexical realization $(l(m))$. Thus, each molecule can be represented as:

$$m = (s(m), p(m), l(m)).$$

3.3. Crystal Seed Selection

A subset of molecules is selected as initial textual nuclei. These nuclei, referred to as *Crystal Seeds*, correspond to short fragments that are semantically complete domain-consistent, and stylistically similar to real tourist reviews. Formally, the set of seeds is defined as:

$$\mathcal{S} = \{m \in \mathcal{M} \mid \phi(m) \geq \tau\},$$

where $(\phi(m))$ is a heuristic quality function that estimates whether a fragment is suitable as an initial nucleus, and (τ) is a minimum quality threshold.

The function $(\phi(m))$ considers three basic criteria: grammatical completeness, semantic clarity, and compatibility with the target tourism domain. Examples of Crystal Seeds include:

“The view was beautiful.”

“The staff was very rude.”

“The food was excellent.”

“The room was uncomfortable.”

3.4. Chemical Reaction Operator

Synthetic reviews are generated by applying reaction rules over compatible linguistic molecules. A chemical reaction is modeled as a binary operator:

$$R : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}^*,$$

where $(\mathcal{M}^*)$ denotes the space of composed textual structures. Given two molecules (m_a) and (m_b) , the reaction operator produces a new compound molecule:

$$m_c = R(m_a, m_b).$$

A reaction is allowed only if the two molecules satisfy a compatibility condition:

$$\kappa(m_a, m_b) = 1,$$

where (κ) evaluates whether the semantic types and polarity orientations of the molecules can be combined without contradiction. For instance, molecules related to food and service may react to form a restaurant experience, whereas molecules related to room quality and location may react to form an accommodation experience:

$$\text{Food} + \text{Service} \rightarrow \text{Restaurant Experience},$$

$$\text{Room} + \text{Location} \rightarrow \text{Accommodation Experience}.$$

The polarity of the resulting compound is determined by the orientation of the participating molecules:

$$p(m_c) = g(p(m_a), p(m_b)),$$

where (g) is an aggregation function that preserves the intended sentiment label of the generated review.

3.5. Crystal Growth Operator

After one or more reactions have produced a stable compound molecule, the text is expanded through a crystallization process. Let (z_0) be an initial seed or compound molecule. The growth process is defined recursively as:

$$z_{t+1} = C(z_t, a_t),$$

where (C) is the crystal growth operator and (a_t) is an additional compatible detail selected from the molecular pool. Each iteration adds information related to aspects such as place, service, atmosphere, cleanliness, waiting time, price, or overall experience.

For example, the following sequence illustrates the crystallization process:

The view was beautiful.

↓

The view was beautiful from the room.

↓

The view was beautiful from the room and the balcony.

↓

The view was beautiful from the room and the balcony, especially at sunset.

The process stops when a termination criterion is satisfied:

$$T(z_t) = 1,$$

where (T) verifies whether the generated review has reached an appropriate length, grammatical coherence, and sentiment consistency.

3.6. Synthetic Corpus Construction

The algorithm was executed repeatedly to generate the synthetic corpus submitted to the REST-MEX 2026 Artificial Data Generation Track. Each generated review was assigned a polarity label and stored together with the corresponding source identifier. Finally, the synthetic instances were merged with the official training corpus and used during the fine-tuning stage of the baseline sentiment classifier under the official evaluation protocol.

4. Results and Discussion

The proposed bio-inspired generation strategy was evaluated using the official metrics defined for the REST-MEX 2026 Artificial Data Generation Track. Since the training corpus presents a severe class imbalance, the competition uses the *Ranking Performance* (RP) score as its primary evaluation metric. Unlike conventional measures such as accuracy, RP assigns greater importance to the minority sentiment classes, providing a more suitable assessment of the usefulness of synthetic data in imbalanced classification settings. Additionally, the organizers report Macro F1, overall Accuracy, and the F1-score for each individual polarity class. Macro F1 evaluates the average classification performance across all sentiment categories, Accuracy measures the proportion of correctly classified instances, and the per-class F1 scores provide a detailed analysis of the classifier’s behavior for each sentiment level.

Table 1 presents the official results obtained by the proposed method together with the Majority Class baseline. The latter simply predicts the most frequent sentiment class for every review, providing a useful lower bound for comparison.

As shown in Table 1, the proposed bio-inspired generation strategy outperformed the Majority Class baseline in terms of RP and Macro F1. In particular, the proposed method achieved an RP score of 0.0146

Table 1

Comparison between the proposed method and the Majority Class baseline.

Method	RP Score	Macro F1	Accuracy	F1(1)	F1(2)	F1(3)	F1(4)	F1(5)
Majority Class Baseline	0.0000	0.0000	0.6564	0.0000	0.0000	0.0000	0.0000	0.7925
Proposed Method	0.0146	0.0156	3.5457	0.0000	0.0500	0.0000	0.0000	0.0290

and a Macro F1 score of 0.0156, whereas the Majority Class baseline obtained zero in both metrics. This result suggests that the generated synthetic reviews provided some useful information for training the classifier, enabling it to identify sentiment categories classifier beyond the trivial majority-class prediction. However, the overall performance remains limited, indicating that the generated corpus only partially improved the classifier’s ability to distinguish among sentiment classes.

The per-class analysis provides additional insight into the behavior of the model. The highest performance was obtained for polarity class 2, with an F1-score of 0.0500, while class 5 reached an F1-score of 0.0290. In contrast, the remaining sentiment categories obtained values close to zero, suggesting that the generated corpus only partially captured the linguistic variability required to discriminate the complete polarity spectrum. These results indicate that the reaction and crystallization mechanisms introduced useful synthetic instances for some sentiment categories; however, their diversity was still insufficient to substantially improve the learning process across all classes.

Although the proposed approach improved over the Majority Class baseline in terms of RP an Macro F1, it did not outperform the original baseline model provided by the organizers. This observation suggests that structural coherence alone is not sufficient for effective synthetic data generation. Sentiment classifiers benefit not only from grammatically correct reviews, but also from a broad lexical, semantic, and contextual diversity that reflects the complexity of real tourist opinions.

From a methodological perspective, the proposed algorithm emphasizes interpretability and traceability over purely statistical generation. Unlike LLM-based generation, whose internal construction process is not always directly observable, each synthetic review produced by our framework can be traced back to a sequence of explicit reaction rules and crystal growth operations. This transparency provides greater control over the generation process and facilitates the analysis of how individual linguistic components contribute to the final synthetic review.

The relatively modest quantitative performance also reveals several opportunities for future improvements. The current reaction operator relies on manually defined semantic compatibility rules, which naturally limit the diversity of the generated reviews. Likewise, the crystallization process performs deterministic expansions from a finite set of textual nuclei, potentially producing reviews that remain too similar to one another. Introducing probabilistic reaction operators, richer semantic representations, adaptive molecule selection, or hybrid generation strategies combining bio-inspired rules with neural language models could substantially increase both lexical diversity and sentiment coverage.

Overall, the experiments demonstrate that the proposed bio-inspired framework constitutes a feasible alternative for synthetic review generation. While its current implementation remains limited in quantitative performance, it improves over a trivial baseline in RP Score and Macro F1 while offering a transparent, interpretable, and fully controllable generation process. These findings suggest that bio-inspired text generation deserves further investigation as a complementary direction for data-centric Natural Language Processing, particularly in applications where explainability and controllability are as important as predictive performance.

5. Conclusions

This paper presented a bio-inspired approach for synthetic tourist review generation in the REST-MEX 2026 Artificial Data Generation Track. In contrast to many current approaches based on Large Language Models, the proposed methodology models text generation as a self-organizing process inspired by two natural phenomena: chemical chain reactions and crystal growth. By representing textual fragments as linguistic molecules and progressively expanding them from crystal seeds, the method generates

coherent reviews through explicit and interpretable generation rules.

The experimental evaluation showed that the proposed approach outperformed the Majority Class baseline according to the official RP and Macro F1 metrics, suggesting that the generated synthetic reviews contributed useful information during the fine-tuning process. However, the generated corpus was not sufficient to surpass the original baseline model provided by the organizers, suggesting that structural coherence alone is not enough to maximize the effectiveness of synthetic data augmentation. Increasing lexical diversity, semantic richness, and contextual variability remains a key challenge for rule-based generation methods.

Despite these limitations, the proposed framework introduces an alternative perspective for synthetic text generation that differs from the dominant paradigm based on neural language models. One of its main advantages is the complete transparency of the generation process, allowing every synthetic review to be explained as the result of successive reaction and crystallization operations. This level of interpretability may be particularly valuable in applications where the origin and construction of synthetic data must be understood and controlled.

Future work will focus on extending the reaction model through probabilistic and adaptive operators, incorporating semantic similarity measures to guide molecular selection, and exploring hybrid architectures that combine the interpretability of bio-inspired generation with the expressive capabilities of Large Language Models. More broadly, this work suggests that concepts borrowed from chemistry and materials science provide an interesting foundation for developing novel data-centric approaches to Natural Language Processing, opening new research opportunities beyond conventional neural text generation.

Acknowledgments

This research was supported by the Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI), Mexico, through the Humanistic Research Project H-2025-G-283, entitled “*Procesamiento de lenguaje natural adaptativo con enfoque humanístico para apoyo comunicacional y diagnóstico en pacientes con enfermedades neurodegenerativas*”. The authors gratefully acknowledge the financial support provided by this project, which made this research possible.

Declaration on Generative AI

Generative AI tools were used solely to improve the language, grammar, and readability of the manuscript. The authors reviewed and edited all AI-assisted content and take full responsibility for the final version of the manuscript.

References

- [1] A. Y. Rodríguez-González, M. Á. Álvarez-Carmona, Machine learning and soft computing: Current trends and applications, 2026.
- [2] A. Díaz-Pacheco, A. B. García-Gutiérrez, A. Y. Rodríguez-González, R. Aranda, M. Á. Álvarez-Carmona, Balancing sentiment analysis datasets through representative-word-guided synthetic review generation: A case study on mexican spanish tourism reviews, *Applied Sciences* 16 (2026) 7398.
- [3] M. Á. 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, Á. 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.
- [4] A. Diaz-Pacheco, M. Á. Álvarez-Carmona, R. Guerrero-Rodríguez, L. A. C. Chávez, A. Y. Rodríguez-González, J. P. Ramírez-Silva, R. Aranda, Artificial intelligence methods to support the research of

destination image in tourism. a systematic review, *Journal of Experimental & Theoretical Artificial Intelligence* 36 (2024) 1415–1445.

- [5] R. Guerrero-Rodríguez, M. A. Álvarez-Carmona, R. Aranda, Á. Díaz-Pacheco, Ai-driven analysis of hotel digital reputation: insights from a mexican destination, *International Journal of Tourism Cities* 12 (2026) 19–36.
- [6] S. Arce-Cardenas, D. Fajardo-Delgado, M. Á. Álvarez-Carmona, J. P. Ramírez-Silva, A tourist recommendation system: a study case in mexico, in: *Mexican international conference on artificial intelligence*, Springer, 2021, pp. 184–195.
- [7] E. P. Ramirez-Villaseñor, H. Pérez-Espinosa, M. A. Álvarez-Carmona, R. Aranda, Design, development, and evaluation of a chatbot for hospitality services assistance in spanish, *Acta universitaria* 33 (2023).
- [8] 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.
- [9] M. Á. Álvarez-Carmona, R. Aranda, R. Guerrero-Rodríguez, A. Y. Rodríguez-González, A. P. López-Monroy, A combination of sentiment analysis systems for the study of online travel reviews: Many heads are better than one, *Computación y Sistemas* 26 (2022) 977–987.
- [10] 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. Y. Rodríguez-González, Overview of rest-mex at iberlef 2021: Recommendation system for text mexican tourism (2021).
- [11] 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) 289–299.
- [12] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, V. Muñoz-Sánchez, A. P. López-Monroy, F. Sánchez-Vega, L. Bustio-Martínez, Overview of rest-mex at iberlef 2023: Research on sentiment analysis task for mexican tourist texts, *Procesamiento del Lenguaje Natural* 71 (2023) 425–436.
- [13] 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, *Procesamiento del Lenguaje Natural* 75 (2025) 499–511.
- [14] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, M. Reyes-Sierra, A. Y. Rodríguez-González, L. Bustio-Martínez, V. Herrera-Semenets, Overview of rest-mex at iberlef 2026: Research on synthetic text modeling for mexican magical towns, *Procesamiento del Lenguaje Natural* 77 (2026).
- [15] A. Bonet-Jover, J. Á. González-Barba, L. Chiruzzo, Overview of IberLEF 2026: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2026)*, co-located with the 42nd Conference of the Spanish Society for Natural Language Processing (SEPLN 2026), CEUR-WS. org, 2026.
- [16] Á. Díaz-Pacheco, R. Guerrero-Rodríguez, M. Á. Álvarez-Carmona, A. Y. Rodríguez-González, R. Aranda, Quantifying differences between ugc and dmo's image content on instagram using deep learning, *Information Technology & Tourism* 26 (2024) 293–329.
- [17] A. Diaz-Pacheco, M. A. Álvarez-Carmona, A. Y. Rodríguez-González, H. Carlos, R. Aranda, Measuring the difference between pictures from controlled and uncontrolled sources to promote a destination. a deep learning approach (2023).
- [18] 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 (2024) 639–661.
- [19] R. Guerrero-Rodríguez, M. Á. 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.
- [20] V. G. Morales-Murillo, H. Gómez-Adorno, D. Pinto, I. A. Cortés-Miranda, P. Delice, Lke-iimas

team at rest-mex 2023: Sentiment analysis on mexican tourism reviews using transformer-based domain adaptation (2023).

- [21] O. G. Toledano-López, M. Á. Álvarez-Carmona, J. Madera, A. Simón-Cuevas, Y. A. López-Rodríguez, H. González Díez, Polarity prediction in tourism cuban reviews using transformer with estimation of distribution algorithms, in: International Workshop on Artificial Intelligence and Pattern Recognition, Springer, 2023, pp. 335–346.
- [22] J. D. Jurado-Buch, S. Minayo-Díaz, J. Tello, K. Chaucanes, L. Salazar, M. Oquendo-Coral, M. Á. Álvarez-Carmona, A single model based on beto to classify spanish tourist opinions through the random instances selection, 2023.
- [23] A. B. García-Gutiérrez, P. E. López-Ávila, P. A. Gallegos-Ávila, R. Aranda, M. Á. Álvarez-Carmona, Balancing of tourist opinions for sentiment analysis task., in: IberLEF@ SEPLN, 2023.
- [24] J. Stack-Sánchez, M. Alvarez-Carmona, A. P. L. Monroy, Nlp_cimat at semeval-2025 task 3: Just ask gpt or look inside. a prompt and neural networks approach to hallucination detection, in: Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025), 2025, pp. 1683–1689.