MAIKRON

Journal of Artificial Intelligence and Computing Applications

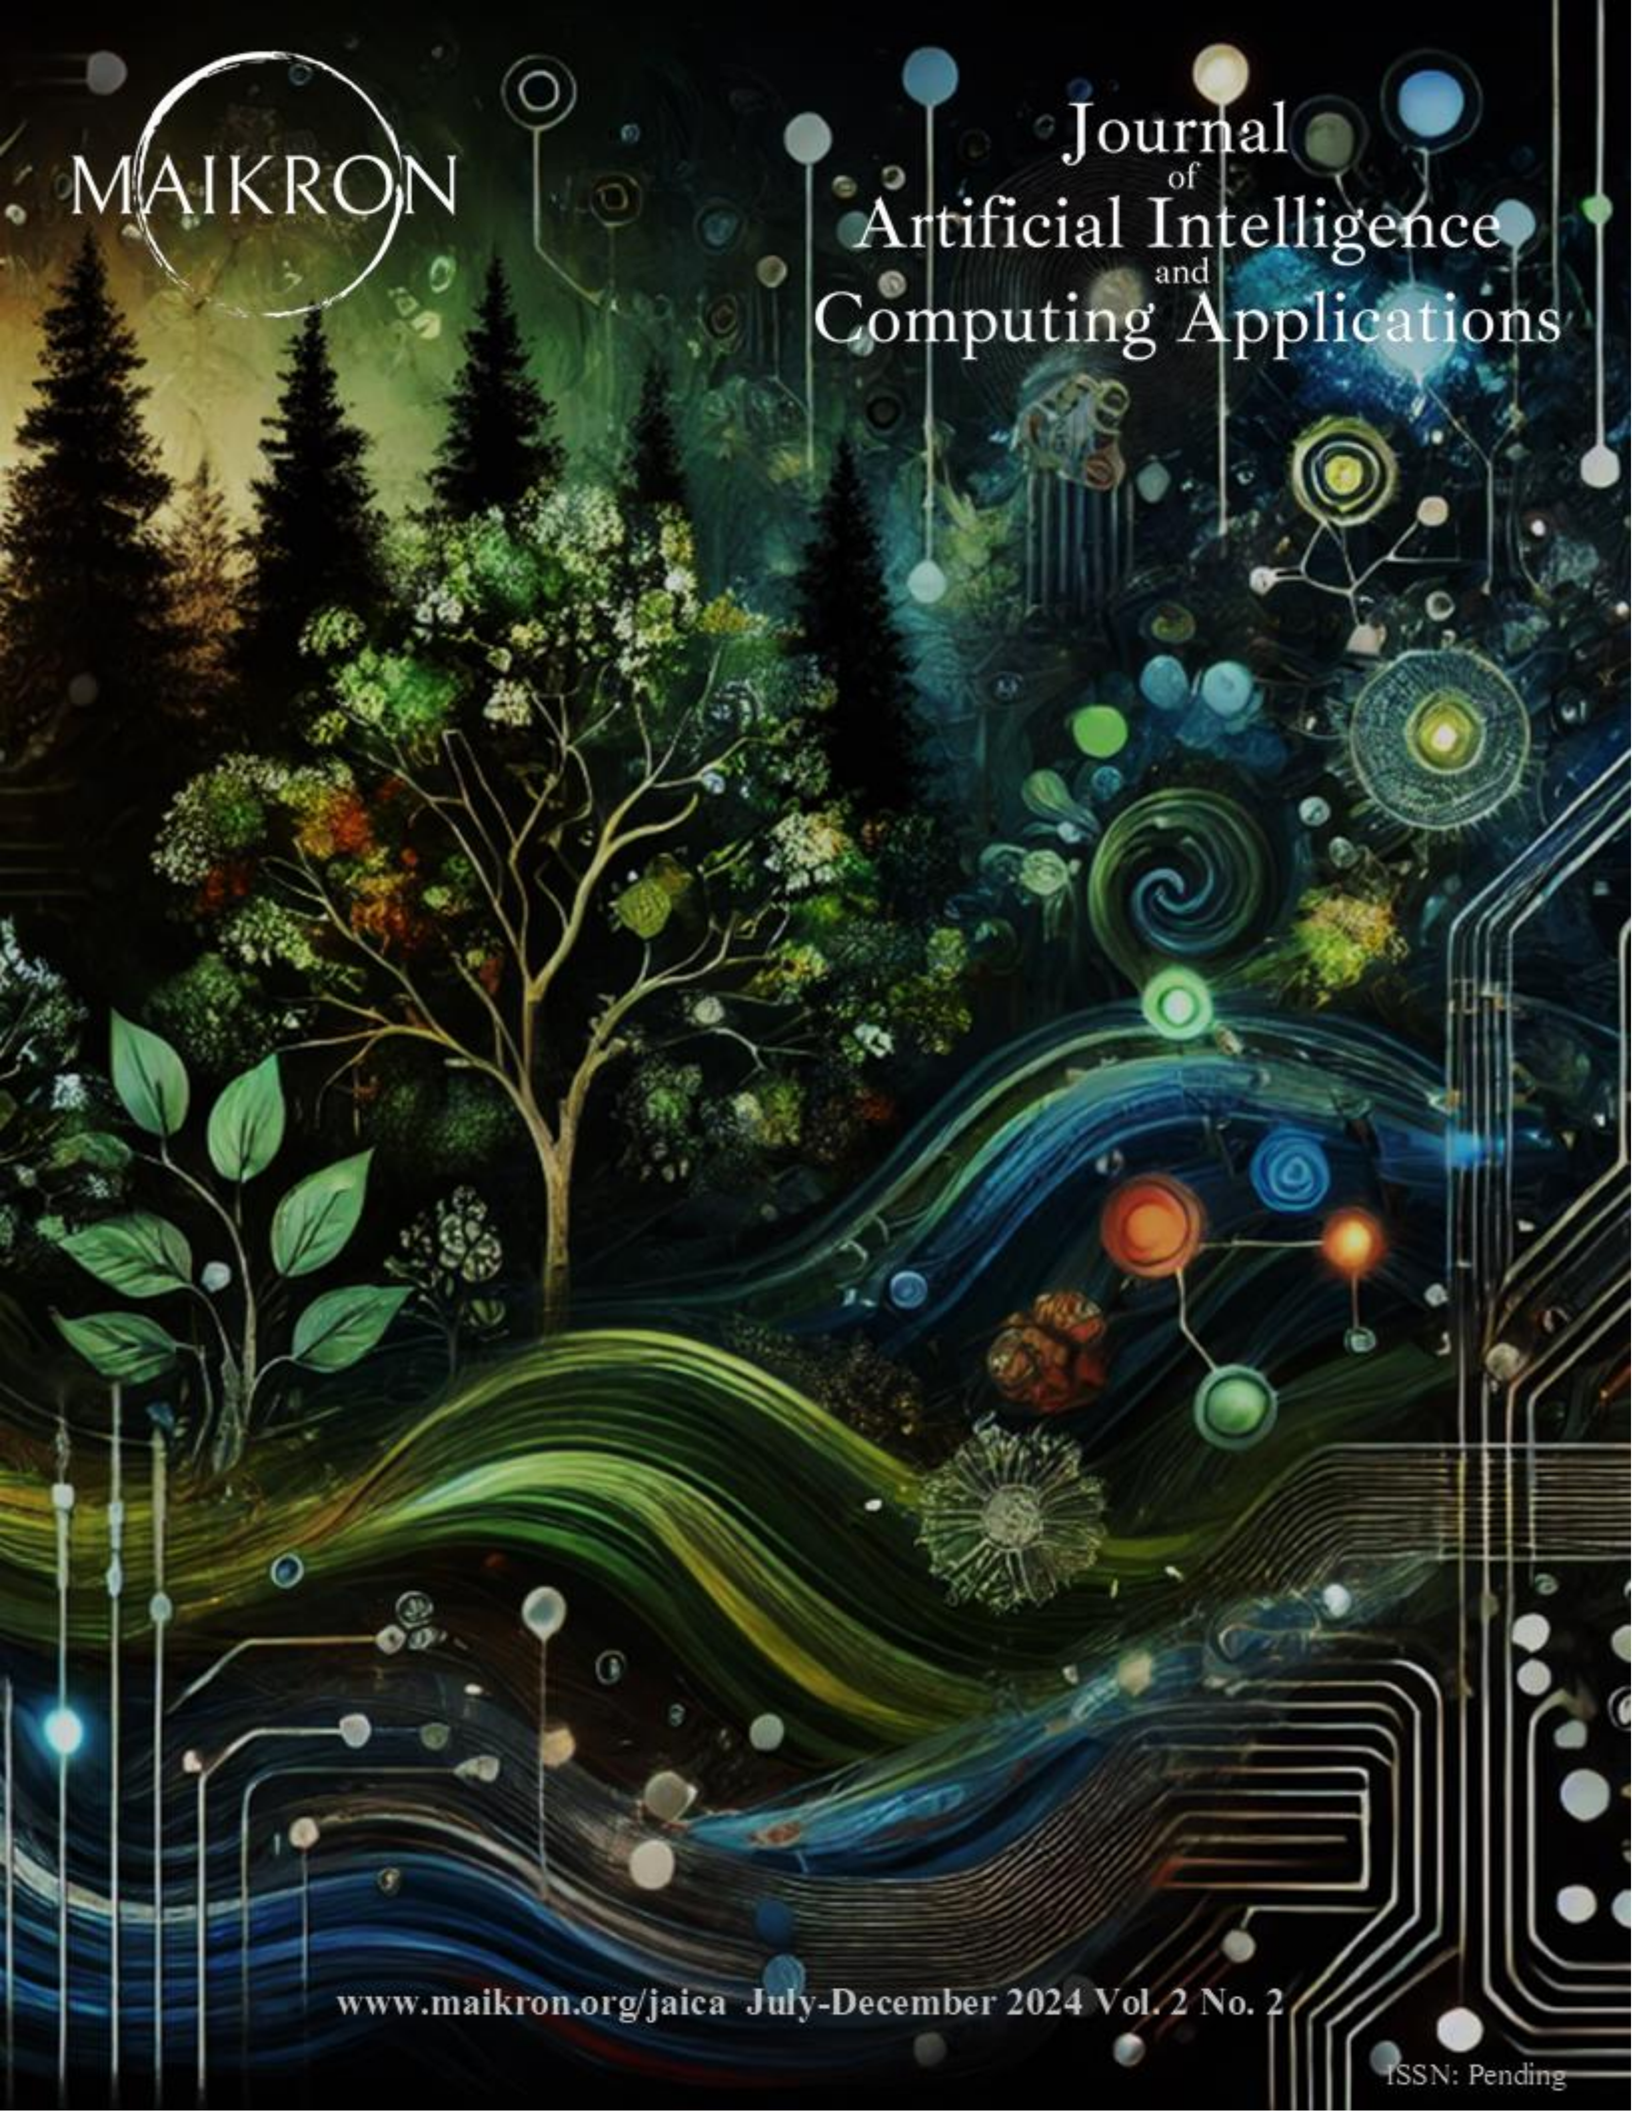

www.maikron.org/jaica July-December 2024 Vol. 2 No. 2

ISSN: Pending

Journal of **Artificial Intelligence** and **Computing Applications**

Journal of Artificial Intelligence and Computing Applications, Vol. 2, Núm. 2, julio – diciembre de 2024, es una publicación semestral editada por la Asociación para el Avance de las Aplicaciones Inteligentes y Tecnologías con Impacto Social, Calle 12 No. 103, Lotes No. 56 y 57, Ucú, Yucatán, C.P. 97357, Tel. (55) 28 55 70 20. URL: https://www.maikron.org/jaica, correo electrónico: contact@maikron.org. Editor responsable: Mauricio Gabriel Orozco del Castillo. Certificado de Reserva de Derechos al uso Exclusivo del Título: 04-2024-032513154000-102, eISSN: en trámite, por el Instituto Nacional del Derecho de Autor (INDAUTOR). Responsable de la última actualización de este número, Mauricio Gabriel Orozco del Castillo, Editor Responsable. Fecha de la última modificación: 31 de diciembre de 2024, Calle 12 No. 103, Lotes No. 56 y 57, C.P. 97357, Ucú, Yucatán. El contenido de los artículos es responsabilidad de los autores y no refleja el punto de vista de los árbitros, del Editor o de la asociación. Se autoriza la reproducción total o parcial de los textos siempre y cuando se cite la fuente completa y la dirección electrónica de la publicación.

Journal of Artificial Intelligence and Computing Applications
Vol. 2, No. 2, July-December 2024

Contents

Original Research Articles

Journal of Artificial Intelligence and Computing Applications

MAIKRON

Foreword

Dear Readers, Contributors, and Colleagues,

It is with immense pride and enthusiasm that I introduce the third issue of the *Journal of Artificial Intelligence and Computing Applications* (JAICA). As we continue our journey into this new year, this issue represents another significant step forward, showcasing the depth and diversity of scholarship that defines our publication.

This issue features a special collection of contributions, including two conference reports, two critical perspectives, one original research article, and two exploration papers. Among these, I would like to highlight an extraordinary inclusion: a critical perspective published in both English and Spanish. This bilingual presentation, an exception to our usual format, underscores the relevance and societal impact of the article titled *"The Role of AI in Combating Mistletoe Infestation in Mexican Forests."* By publishing it in Spanish, we aim to ensure accessibility to the local audience, as the article addresses the mistletoe infestation issue specific to Mexico. At the same time, the English version extends its reach and fosters dialogue across a broader audience, emphasizing the global significance of the intersection between AI and ecological conservation. Each article embodies the spirit of JAICA by exploring real-world challenges through the lens of cutting-edge AI technologies.

Complementing these articles are two conference reports, which offer an in-depth look at the *32nd International Conference on Case-Based Reasoning (ICCBR) 2024* and the *2024 International Conference on Artificial Intelligence for Mental Health (ICAIMH).* These reports provide insights into groundbreaking discussions, emerging trends, and collaborative opportunities in AI-driven research and its applications in mental health and socio-ecological domains.

We are also proud to present an original research article titled *"An Explainable Clustering-Based Approach for Cyber Situational Awareness on Masquerade Attacks Detection."* This study proposes a novel and interpretable methodology that combines optimized clustering algorithms with noise reduction techniques to improve intrusion detection in cybersecurity. The results point to a significant step forward in enhancing Cyber Threat Intelligence by bridging performance and explainability.

In our *Applied AI Exploration Papers (AAIEP)* section, we feature two compelling contributions. The first, *"A Deep Learning Approach for Automated Identification of Triatoma infestans Using YOLOv8,"* demonstrates how AI can support public health by automating the detection of *T. infestans*, the insect vector of Chagas disease. The second, *“An Exploratory Application of Empirical Mode Decomposition and Recurrent Neural Networks for Meteorological Time Series Prediction,”* presents a hybrid forecasting model using EMD and LSTM to better predict weather variables in the tropical region of Mérida, Yucatán. Both papers highlight the versatility of AI in tackling domain-specific challenges and the importance of early-stage prototypes in shaping future innovation.

I extend my heartfelt gratitude to the dedicated editorial team and reviewers, whose expertise ensures the highest standards of quality; and to the sponsors and collaborators, including Maikron, for their unwavering support. Finally, I thank our readers for their continued interest and engagement, which inspires us to push the boundaries of AI and computing applications.

As we look ahead, JAICA remains committed to its mission of fostering a collaborative and inclusive platform for AI research with real-world impact. Together, we will continue to explore the limitless possibilities of artificial intelligence and its transformative potential.

Warmest regards,

Mauricio G. Orozco-del-Castillo

Editor-in-Chief

Journal of Artificial Intelligence and Computing Applications (JAICA)

Journal of Artificial Intelligence and Computing Applications (2024) 2(2): 1–3

Conference Reports

Conference Report on the 2024 International Conference on Case-Based Reasoning (ICCBR 2024)

Juan Recio-Garcia
Universidad Complutense de Madrid, Spain

Mauricio G. Orozco-del-Castillo
Tecnológico Nacional de México / IT de Mérida, México

Derek Bridge
University College Cork, Ireland

The 32nd International Conference on Case-Based Reasoning (ICCBR 2024) was held in Mérida, Mexico, from July 1 to July 4, 2024. As the premier international event dedicated to Case-Based Reasoning (CBR), ICCBR 2024 continued its tradition of fostering cutting-edge research and application development in the field. This edition marked the first time the conference was hosted in Mexico, underscoring a strategic effort to engage local researchers and highlight the significance of CBR in addressing global challenges. Under the theme "AI for socio-ecological welfare," the conference featured a Special Track on Artificial Intelligence, designed to introduce Mexican researchers to the CBR paradigm while emphasizing the potential of AI in promoting social and ecological well-being. The event included keynote addresses from prominent experts, workshops on emerging topics like explainability in AI and the integration of CBR with Large Language Models (LLMs), as well as the Doctoral Consortium aimed at mentoring early-career researchers. Through its diverse program, ICCBR 2024 successfully brought together academics and practitioners to explore innovative solutions and expand the frontiers of CBR research.

The 32nd ICCBR was organized through the collaborative efforts of esteemed institutions and dedicated individuals. As Program Committee Chairs, we—Juan A. Recio-Garcia, Mauricio G. Orozco-del-Castillo, and Derek Bridge (University College Cork, Ireland)—took responsibility for shaping the academic program and maintaining the conference's high standards of quality. We were supported by invaluable contributions from our colleagues, including Lukas Malburg (Trier University, Germany), who excelled in his role as Workshops Chair, and Juan Carlos Valdiviezo-Navarro (CentroGeo, Mexico), who expertly led the AI-Track.

The local organization was expertly managed by a team from the Tecnológico Nacional de México/IT de Mérida, including Local Chairs Nora Cuevas-Cuevas, Carlos Bermejo-Sabbagh, and Pedro Ortiz-Sanchez, as well as Ana Martin-Casado, from Universidad de la Rioja. These efforts were bolstered by the contributions of several volunteers from the AAAIMX Student Chapter, whose dedication was crucial to the event's success.

The conference was supported by various sponsors and collaborating organizations, including Jarkol Technologies, Maikron, and the Instituto Tecnológico de Mérida. Additional sponsorship came from entities like the Honorary Chair BOSCH-UCM on Artificial Intelligence and the Science Foundation Ireland Insight Centre for Data Analytics, which ensured the financial and technical feasibility of the conference.

ICCBR 2024 faced the unique milestone of being the first edition of the conference held in Mexico. This presented both logistical challenges and opportunities, such as navigating local infrastructure while successfully broadening the global reach of the CBR community. By hosting the event in Mérida, the organizers aimed to engage Mexican researchers and foster greater regional participation in the field of CBR.

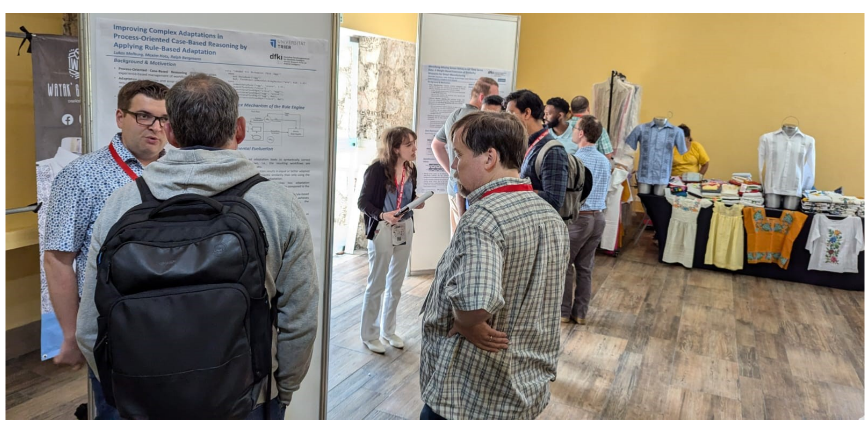

Figure 1. Poster session at ICCBR

https://doi.org/10.5281/zenodo.14700658

Figure 2. Oral presentation at ICCBR

The overarching vision of ICCBR 2024 was to showcase the potential of AI for socio-ecological welfare, emphasizing the intersection of CBR and pressing global challenges. Through its meticulous organization, the conference aimed to inspire collaboration, innovation, and the application of CBR in diverse domains.

ICCBR 2024 featured two distinguished keynote speakers whose talks underscored the event's theme, CBR for Socio-Ecological Welfare, and offered deep insights into the current advancements and applications of CBR. Prof. Enrique Sucar-Sucar, a renowned figure in artificial intelligence, presented an insightful opening talk titled "Causal Models: Representation, Reasoning, and Discovery." In his presentation, Prof. Sucar-Sucar discussed the role of causal models in reasoning and discovery, emphasizing their potential to advance decision-making in complex systems. With a distinguished academic background that includes degrees from ITESM, Stanford University, and Imperial College, and a career adorned with numerous accolades such as the National Science Prize in 2016, Prof. Sucar-Sucar set the tone for the conference by highlighting the potential of CBR and AI in addressing socio-ecological challenges. Prof. Nirmalie Wiratunga delivered the closing keynote, "Intelligent Sharing of Explanation Experience by Users for Users: Case-Based Reasoning for Explanation Strategy Reuse." As a leader in AI and intelligent systems research, Prof. Wiratunga focused on the use of CBR for creating and reusing explanation strategies, with applications in retrieval-augmented Q&A systems and Large Language Models. Her presentation provided thought-provoking perspectives on enhancing user-centric AI systems. As the head of the Artificial Intelligence & Reasoning Research Group at RGU and a frequent contributor to the ICCBR community, Prof. Wiratunga concluded the conference with insights that encouraged innovative approaches and applications in the field.

The conference program featured a series of thematic tracks and engaging sessions. The Retrieval track explored innovative methods for case retrieval, including integrating kNN retrieval with graphical models and visualization techniques for improving CBR comprehension. The Explainable AI track addressed the intersection of XAI and CBR, presenting semi-factual explanations and user-centered evaluation metrics for explainable systems. One of the highlights of the conference was the CBR and LLMs track, showcasing groundbreaking work on integrating CBR with Large Language Models, including retrieval-augmented generation and CBR frameworks for legal question answering.

In addition to these sessions, ICCBR 2024 included two dynamic workshops. The X-CBR workshop focused on explainability in CBR systems, featuring discussions on counterfactual explanations and hybrid approaches for complex problem-solving. The CBR-LLMs workshop explored the synergy between CBR and Large Language Models, presenting papers on topics such as persistent memory for LLMs and the role of CBR in augmenting LLM capabilities. These workshops provided a platform for deep discussions and the exchange of ideas on emerging topics.

The Doctoral Consortium was another notable feature of the conference, providing early-career researchers with an excellent platform to present their work and receive constructive feedback from experienced mentors. Highlights of the consortium included discussions on using CBR for mental health interventions and innovative indexing techniques for case retrieval. Through this initiative, the conference fostered an environment of learning, collaboration, and mentorship, contributing to the development of the next generation of researchers in the field of CBR.

ICCBR 2024 saw the participation of 50 registered attendees, representing a diverse range of countries and regions. The majority of participants came from Mexico (12), followed by the United States (9), Germany (8), China (7), and Spain (6), with additional representation from the United Kingdom, Ireland, Italy, New Zealand, Norway, and Canada. Geographically, attendees were distributed across the Americas (22), Europe (27), and other regions (1), reflecting the truly international nature of the conference. The submission process demonstrated the growing global interest in CBR. A total of 74 papers were submitted, of which 29 were accepted, yielding an acceptance rate of approximately 39%. The program featured a balanced mix of oral presentations and poster sessions, showcasing a broad spectrum of research that ranged from theoretical advancements to practical applications.

The 32nd ICCBR conference embraced the theme *CBR for Socio-Ecological Welfare*, directing the community's attention to the growing potential of CBR in addressing global challenges, particularly in environmental sustainability, resource management, and social equity. This thematic focus was evident across the tracks and workshops, with presentations highlighting the integration of CBR with modern AI techniques, such as Large Language Models, and its application in diverse fields including mental health, urban planning, and legal systems. Emerging trends included a strong emphasis on explainability, as showcased in multiple sessions on Explainable AI, as well as the expanding role of hybrid approaches combining CBR with deep learning and neural networks.

The conference also featured interactive elements that enriched the experience for participants. Workshops such as X-CBR and CBR-LLMs allowed attendees to engage deeply with specific topics, fostering meaningful dialogue and collaboration. The poster session provided an informal setting for researchers to present their work and exchange ideas, encouraging dynamic interactions among attendees. A key highlight of the conference was the Best Paper Awards ceremony, which recognized outstanding contributions in two categories. The overall Best Paper Award was presented to Nirmalie Wiratunga, Ramitha Abeyratne, Lasal Jayawardena, Kyle Martin, Stewart Massie, Ikechukwu Nkisi-Orji, Ruvan Weerasinghe, Anne Liret, and Bruno Fleisch for their work *"CBR-RAG: Case-Based Reasoning for Retrieval Augmented Generation in LLMs for Legal Question Answering"*. The Best Student Paper Award was given to Mirko Lenz, Lukas Malburg, and Ralph Bergmann for their contribution *"CBRkit: An Intuitive Case-Based Reasoning Toolkit for Python"*. These awards underscored the innovative and impactful research presented at the conference, celebrating both seasoned researchers and emerging talent in the field.

The success of the 32nd ICCBR conference would not have been possible without the dedication and efforts of many individuals and organizations. We extend our deepest gratitude to the organizing committee, whose tireless work ensured the seamless execution of the event. We are deeply grateful to our keynote speakers, Prof. Enrique Sucar-Sucar and Prof. Nirmalie Wiratunga, for sharing their insights and expertise, which greatly enriched the conference program. We also thank the reviewers and program committee members for their diligent evaluations, ensuring the quality of the submissions. The conference was made possible through the generous support of our sponsors and partners. We acknowledge Jarkol Technologies, Maikron, and the Instituto Tecnológico de Mérida for their financial and logistical contributions. We also thank the Honorary Chair BOSCH-UCM on Artificial Intelligence and the Science Foundation Ireland Insight Centre for Data Analytics for their sponsorship and encouragement of academic excellence. Finally, we would like to recognize the efforts of the volunteers from the AAAIMX Student Chapter and ITM-ACM Student Chapter, whose enthusiasm and commitment ensured the smooth running of the event. Their contributions were instrumental in creating a welcoming and vibrant atmosphere for all participants. To everyone who played a role in making ICCBR 2024 a success, we extend our heartfelt thanks.

ICCBR 2024 marked a significant milestone in the evolution of the CBR community. This year's event was impactful and unique, not only for its thematic focus on CBR for Socio-Ecological Welfare but also for its dynamic engagement across diverse topics and innovative intersections with AI technologies like Large Language Models. The location itself added a distinctive flavor, as hosting the conference in Mexico for the first time fostered regional engagement and encouraged new collaborations within the global CBR community.

The major takeaways from the conference highlighted the growing relevance of CBR in addressing complex, real-world challenges. Keynotes by Prof. Enrique Sucar-Sucar and Prof. Nirmalie Wiratunga emphasized the importance of integrating causal reasoning and explanation reuse into practical applications, setting the stage for future research directions. Sessions on topics such as hybrid CBR-AI systems, explainable AI, and retrieval-augmented generation showcased how CBR continues to evolve as a versatile and impactful methodology. The recognition of outstanding contributions through the Best Paper Awards further underscored the innovative work being conducted in the field, celebrating both seasoned researchers and emerging talents.

Looking ahead, the ICCBR community is excited to announce that the 33rd International Conference on Case-Based Reasoning will take place in Biarritz, France, from June 30 to July 3, 2025. With its picturesque setting and a reputation for fostering cutting-edge research, Biarritz promises to provide an inspiring backdrop for the next chapter in the ICCBR journey. We look forward to welcoming researchers, practitioners, and enthusiasts to continue advancing the frontiers of CBR in this vibrant and collaborative forum.

The conference proceedings, published by Springer as part of the collection *Lecture Notes in Artificial Intelligence (LNAI)*, are available at: https://link.springer.com/book/10.1007/978-3-031-63646-2

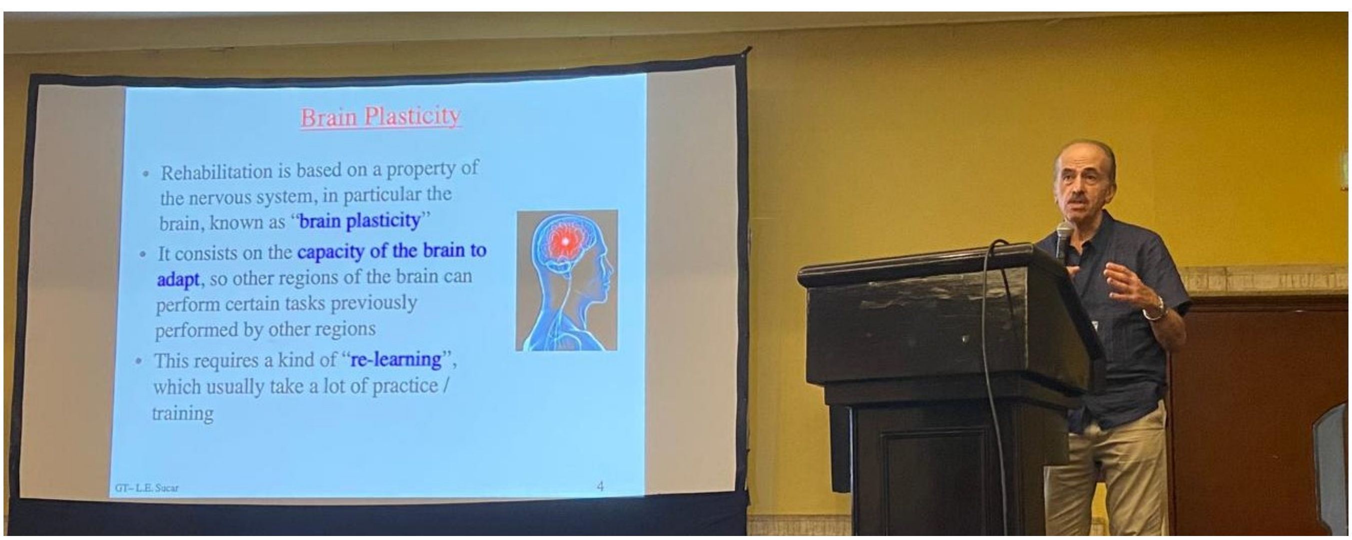

Figure 3. Opening keynote by Prof. Enrique Sucar-Sucar

Conference Reports

Conference Report on the 2024 International Conference on Artificial Intelligence for Mental Health (ICAIMH 2024)

Mauricio G. Orozco-del-Castillo
Tecnológico Nacional de México / IT de Mérida, México

Juan Recio-Garcia
Universidad Complutense de Madrid, Spain

THE International Conference on Artificial Intelligence for Mental Health (ICAIMH) 2024 took place from July 1 to July 4, 2024, at the Holiday Inn in Mérida, Yucatán, Mexico. This second edition of ICAIMH built upon the success of its inaugural conference, which was also held in Mérida, at the Tecnológico Nacional de México/Instituto Tecnológico de Mérida. The central theme of the conference, *How artificial intelligence (AI) can benefit mental health*, brought together experts and practitioners from diverse disciplines, including AI, computing, psychology, psychiatry, and healthcare. By fostering interdisciplinary collaboration, ICAIMH 2024 sought to advance innovative solutions for mental health challenges, bridging the gap between cutting-edge AI technologies and the pressing needs of mental health care systems.

ICAIMH 2024 was organized by the *Association for the Advancement of Intelligent Applications and Technologies with Social Impact* (Maikron), a non-profit association dedicated to leveraging AI for societal benefits. The event was made possible through the support of key partners, including ICCBR 2024, which served as the principal associated conference, alongside the AAAIMX Mexican Student Chapter, the ACM ITM Student Chapter, Universidad Complutense de Madrid, Centro de Investigación y de Estudios Avanzados (CINVESTAV), CentroGeo, Tecnológico Nacional de México/Instituto Tecnológico de Mérida, and Jarkol Technologies.

The program committee played a critical role in shaping the conference's agenda and ensuring its academic rigor. Members of the committee included Esperanza Carolina Orozco-del-Castillo, Juan Carlos Valdiviezo-Navarro, Carlos Bermejo-Sabbagh, Nora Cuevas-Cuevas, Rasikh Tariq, and Pedro Ortiz-Sánchez, whose combined expertise spanned multiple disciplines relevant to AI and mental health. This year's conference also featured a special collaboration with ICCBR 2024, introducing a dedicated topic on Mental Health within the "Special Track on Artificial Intelligence for Socio-Ecological Welfare." Accepted works from ICAIMH 2024 were published as part of the ICCBR 2024 Workshops Proceedings in CEUR-WS (https://ceur-ws.org/Vol-3708/), which is indexed in Scopus, further enhancing the visibility and academic impact of the research presented.

ICAIMH 2024 featured a carefully curated program of keynote lectures, showcasing the expertise of distinguished speakers from renowned institutions. The conference opened on July 2 with three thought-provoking

Figure 1. Official banner of ICAIMH 2024. On the left, the ICAIMH logo depicts a brain encased within a protective technological framework, symbolizing the role of AI as both an enabler and a safeguard for mental health advancements. On the right, the 2024 edition logo showcases a jaguar, a revered figure in Mayan culture often associated with strength, protection, and deep introspection, highlighting the connection between cultural heritage and mental health.

https://doi.org/10.5281/zenodo.14707018

Figure 2. Lecture by Prof. Enrique Súcar-Súcar

talks. Prof. Enrique Sucar-Succar, from the Instituto Nacional de Astrofísica, Óptica y Electrónica, delivered the first keynote, “Gesture Therapy: From Motor Rehabilitation to Cognitive Stimulation,” which explored the potential of AI-driven therapies for cognitive and physical recovery. The second keynote, presented by Dr. Israel Sánchez-Domínguez of the Universidad Nacional Autónoma de México, focused on “IoT as a Measurement and Monitoring Tool,” highlighting the integration of IoT technologies for mental health applications. Dr. José Luis Batún-Cutz, from the Universidad Autónoma de Yucatán, closed the day with an insightful analysis, “Suicides in Yucatán from 2012 to 2022: An Analysis Using MCA and Stochastic Modeling,” emphasizing the role of advanced modeling in understanding mental health trends.

On July 3, the program continued with two captivating talks. Dr. Eduardo Barbará-Morales, from Universidad Anáhuac Mayab, presented “Image Biomarkers and Their Applications in Artificial Intelligence for Characterizing Morphological Changes Associated with Alzheimer’s Disease,” which demonstrated the intersection of neuroimaging and AI in addressing neurodegenerative disorders. The second keynote of the day, delivered by Dr. Luis Alberto Muñoz-Ubando of the Instituto Tecnológico y de Estudios Superiores de Monterrey, titled “Advanced Telerobotics: From EEG to Imaginary Language and Action,” explored the applications of telerobotics in mental and cognitive health.

Unfortunately, due to the imminent arrival of Hurricane Beryl, a Category 2 storm forecasted to impact the Yucatán Peninsula on July 5, 2024, the conference organizers made the decision to cancel the sessions scheduled for July 4. This precautionary measure was taken to ensure the safety of all participants and speakers. The canceled sessions included the talks “Your Emotions Are Trying to Tell You Something, Do You Know What It Is?” by Dr. Esperanza Carolina Orozco-del-Castillo from CINVESTAV, and “Second Victims, Who Cares for Those Who Care?” by Dr. Ana María Martín-Casado from Universidad Internacional de la Rioja and Hospital Universitario de Guadalajara.

ICAIMH 2024 attracted a total of 100 registered attendees, reflecting a diverse and interdisciplinary audience committed to advancing the intersection of AI and mental health. Participants represented a wide range of affiliations, including esteemed academic institutions such as the Instituto Nacional de Astrofísica, Óptica y Electrónica (INAOE), Universidad Nacional Autónoma de México (UNAM), Universidad Autónoma de Yucatán (UADY), TecNM/ITM and Universidad Anáhuac Mayab. Other notable attendees included professionals from healthcare organizations like the Instituto Mexicano del Seguro Social (IMSS) and the Secretaría de Salud de Yucatán, as well as members of public service institutions such as the Ayuntamiento de Mérida and DIF Mérida.

Figure 3. Lecture by Prof. Israel Sánchez-Domínguez

Figure 4. Lecture by Prof. Luis Alberto Muñoz-Ubando

Importantly, ICAIMH 2024 was free of charge, making it accessible to anyone interested in the intersection of AI and mental health. This open format encouraged participation from a wide array of attendees, including students, researchers, practitioners, and members of the general public, fostering a truly inclusive and collaborative environment. The event’s accessibility contributed significantly to the diversity of perspectives represented and the richness of the discussions.

The conference also welcomed participants from private industry, including representatives from MID Data Solutions, SATEC Mexico, and Diagnósticos Biometh, alongside independent practitioners and researchers. This diversity spanned academia, healthcare, public services, and industry, fostering an interdisciplinary environment conducive to meaningful collaboration. The broad participation emphasized the relevance of AI in mental health across multiple sectors and highlighted the conference’s ability to bring together experts from varied backgrounds to address pressing challenges in mental health.

ICAIMH 2024 showcased a rich and interdisciplinary program centered on the theme of leveraging AI to address pressing mental health challenges. The conference explored a wide array of focus areas, including AI-driven prevention, detection, and treatment of mental health conditions, the use of chatbots and intelligent agents for mental health support, IoT applications for biomarker monitoring, and advanced neuroimaging techniques for analyzing morphological changes in neurological disorders.

The program's interdisciplinary nature was evident in the range of speakers and sessions, which bridged the fields of AI, psychology, psychiatry, and public health. Keynotes and discussions highlighted the integration of AI technologies with practical applications in mental health, fostering collaborations between technologists and mental health professionals. This emphasis on cross-disciplinary dialogue created a unique platform for sharing insights and addressing complex mental health issues.

The success of ICAIMH 2024 would not have been possible without the contributions and support of numerous individuals and organizations. We extend our heartfelt gratitude to the *Association for the Advancement of Intelligent Applications and Technologies with Social Impact* (Maikron) for spearheading the organization of the event and ensuring its alignment with the overarching mission of leveraging AI for societal benefit. Special thanks go to the program committee members whose expertise and dedication shaped the program's academic rigor and inclusivity.

We are deeply grateful to our keynote speakers, who shared their valuable insights and expertise, enriching the conference with their thought-provoking presentations. Their contributions highlighted the transformative potential of AI in addressing complex mental health challenges.

We also acknowledge the critical support of key partners, including ICCBR 2024, whose collaboration provided a unique opportunity to integrate ICAIMH 2024 submissions into its Special Track. The additional partners played significant roles in ensuring the conference's success.

Finally, we thank the institutions, organizations, and individuals from academia, healthcare, public services, and industry who participated in the event, fostering an environment of collaboration and interdisciplinary engagement. ICAIMH 2024 exemplified the power of collective effort and the importance of bringing together diverse perspectives to address the pressing challenges at the intersection of AI and mental health.

ICAIMH 2024 highlighted the transformative potential of AI in addressing critical mental health challenges, reinforcing the importance of interdisciplinary collaboration between AI researchers, mental health professionals, and industry practitioners. The conference demonstrated how cutting-edge AI technologies, such

Figure 6. Lecture by Prof. José Luis Batún-Cutz

as intelligent agents, IoT applications, and advanced neuroimaging techniques, can revolutionize mental health care by enhancing prevention, detection, and treatment strategies. The discussions and presentations urged the need for ethical and human-centered approaches to ensure these innovations benefit individuals and communities in meaningful ways.

The event's major takeaway was the recognition that bridging the gap between AI and mental health care requires not only technical advancements but also a sustained effort to foster collaboration across diverse fields. By providing a platform for dialogue and knowledge exchange, ICAIMH 2024 contributed to building a stronger, more integrated community dedicated to tackling mental health issues with innovative and impactful solutions.

Looking ahead, ICAIMH is excited to announce its return to Mérida, Yucatán, for its third edition in July 2025. Building on the success of this year's conference, ICAIMH 2025 will continue to explore the intersection of AI and mental health, fostering interdisciplinary collaboration and driving progress in this vital area of research. For more information about ICAIMH, please visit the official website at www.icaimh.org, or learn more about the upcoming 2025 edition at www.icaimh.org/2025. We look forward to welcoming participants from around the world to join us once again in Mérida for another inspiring and impactful event.

Figure 5. Prof. Eduardo Barbará-Morales receiving the certificate after his keynote lecture

Journal of Artificial Intelligence and Computing Applications (2024) 2(2): 7–17

Critical perspective

The Role of AI in Combating Mistletoe Infestation in Mexican Forests: A Call for Interdisciplinary Action

Juan Carlos Valdiviezo-Navarro
CONAHCYT-CentroGeo, Yucatán, Mexico

Mauricio G. Orozco-del-Castillo
Tecnológico Nacional de México/IT de Mérida, Yucatán, Mexico

Este artículo ha sido publicado tanto en inglés como en español para garantizar la accesibilidad a una audiencia más amplia. Se puede acceder a la versión en español de este artículo en el siguiente DOI: https://doi.org/10.5281/zenodo.14720118.

Keywords: mistletoe detection, artificial intelligence in conservation, ecological management tools

PERSPECTIVE

Mistletoe is a hemiparasitic plant that attaches to host trees, extracting water and nutrients, which weakens the trees, reduces biodiversity, and disrupts critical ecosystem services. While often romanticized in cultural contexts, mistletoe poses a serious ecological threat, particularly in urban and forested ecosystems where infestations can spread unchecked. This article examines the role of artificial intelligence (AI) and remote sensing in addressing the detection and management challenges posed by mistletoe. Through a critical evaluation of methodologies ranging from texture-based machine learning to advanced deep learning models such as ResNet-34, this paper reflects on the successes, limitations, and implications of these approaches. Our interdisciplinary research highlights the transformative potential of combining AI with ecological expertise to develop scalable and efficient tools for conservation. However, we also identify key challenges, including the need for equitable access, ethical considerations, and scalability across diverse ecological contexts. Moreover, we emphasize the importance of engaging the broader community, as misconceptions about mistletoe hinder conservation efforts. By integrating public awareness with technological advancements, we advocate for a balanced and sustainable approach to ecological management. This paper aims to provoke critical dialogue and inspire actionable strategies for leveraging AI in addressing global conservation challenges.

Introduction

Mistletoe infestations represent a localized instance of a broader global issue: the ecological and socioeconomic challenges posed by invasive plant species. These species, characterized by their ability to rapidly establish and proliferate in non-native ecosystems, have significant and far-reaching impacts. Globally, invasive plants contribute to biodiversity loss by outcompeting native species, disrupting food webs, and altering ecosystem services [1, 2]. In this sense, their spread can compromise carbon sequestration, temperature regulation, and habitat support for native wildlife, while simultaneously creating socioeconomic burdens by affecting agriculture, forestry, and public health.

Mistletoe is a hemiparasitic plant that requires a host tree to extract water, minerals, and other nutrients by inserting a specialized root called haustoria. This parasitic relationship weakens host trees over time, reducing their structural stability, decreasing biodiversity, and threatening critical ecosystem services [3, 4]. In Mexico City, three mistletoe species—*Cladocolea loniceroides*, *Phoradendron velutinum*, and *Struthanthus interruptus*—have been identified as significant contributors to ecological disruption and one of the main problems of urban green spaces[5].

Among these species, *S. interruptus* and *P. velutinum* stand out due to their extensive impact on urban ecosystems and forested areas. For example, *S. interruptus* is prevalent in 9 of the 16 boroughs of Mexico City [5], while *P. velutinum* is particularly noted for its ability to blend seamlessly with host vegetation in conservation forests [6]. Both species form dense clusters in advanced infestation stages, dominating tree crowns and disrupting natural tree structures [7]. These infestations compromise forest health, reduce biodiversity, and hinder essential ecosystem services, including carbon sequestration and habitat support for native species [3, 4]. The widespread ecological and economic impacts of these mistletoe species underscore the urgency of developing effective and scalable detection and management strategies.

Detecting and managing mistletoe infestations presents significant challenges due to the plant's ability to visually blend with host vegeta-

E-mail address: mauricio.orozco@itmerida.edu.mx
https://doi.org/10.5281/zenodo.14699764

tion, particularly in its early stages [6]. As shown in Figure 1, mistletoe clusters closely resemble the healthy foliage of their host trees, making manual identification difficult, especially in dense vegetation. This challenge is further compounded by environmental variability, including changes in lighting, viewing angles, and seasonal foliage density, which reduce the reliability of traditional methods [7]. Manual surveys, conducted by trained experts, remain the main approach for mistletoe detection, but they are labor intensive, time consuming, and resource demanding, rendering them impractical for large-scale implementation [6]. Additionally, the physical inaccessibility of certain forested areas exacerbates these inefficiencies, allowing infestations to spread undetected. Together, these limitations underscore the significant difficulties faced in managing mistletoe infestations and highlight the pressing need for improved methods to enhance the efficiency and scalability of detection efforts.

The urgent need for innovative and scalable solutions to address the mistletoe infestation problem in Mexico catalyzed the formation of an interdisciplinary research initiative. This collaborative effort brought together experts from Centro de Investigación en Ciencias de Información Geoespacial, A.C. (CentroGeo), Tecnológico Nacional de México/Instituto Tecnológico de Mérida (ITM), Instituto Tecnológico de Pabellón de Arteaga (ITPA), and included support from Secretaría de Educación, Ciencia, Tecnología e Innovación de la Ciudad de México (SECTEI), Secretaría del Medio Ambiente de la Ciudad de México (SEDEMA), and the UNAM-Huawei Innovation Space. By combining expertise in artificial intelligence (AI), remote sensing, ecological management, and field operations, the team seeks to tackle the multifaceted challenges posed by mistletoe infestations in both urban and forested ecosystems.

Central to the team strategy is the integration of advanced technologies with ecological knowledge. In particular, remote sensing technologies, such as high-resolution multispectral imagery collected via unmanned aerial vehicles (UAVs), have provided the ability to monitor large areas efficiently and capture detailed information related to vegetation health, phenology, among others. These data, combined with cutting-edge AI techniques, have enabled the development of scalable detection tools. Hence, by leveraging machine learning algorithms and deep learning models, the team aims to create methodologies capable of distinguishing mistletoe clusters from surrounding healthy vegetation with high precision, even under challenging environmental conditions.

The overarching goal of the initiative is to develop innovative, technology-driven methodologies for the automated detection and management of mistletoe infestations. Designed to address the limitations of traditional approaches, these methodologies prioritize scalability, precision, and efficiency while ensuring interpretability and actionable insights for ecological practitioners. By focusing on scalability, the tools aim to support diverse ecological contexts, from urban parks to forestry conservation areas, contributing to biodiversity preservation and ecosystem health.

To address the challenges of mistletoe detection, comprehensive datasets have been collected from diverse study sites, including the *San Bartolo Ameyalco* ecological conservation area, the *Ramón López Velarde Garden*, and Xochimilco pier, in Mexico City. These locations were chosen in conjunction with SEDEMA due to their high prevalence of mistletoe infestations and their diverse environmental conditions, which provide a robust foundation for developing robust detection models.

The datasets were built using a P4 multispectral drone, which provides detailed spatial and spectral information critical for distinguishing mistletoe from host vegetation. This device allowed to capture high-resolution imagery, preserving the fine-grained details of mistletoe clusters, such as their broad, velvety leaves and distinct coloration. These characteristics are essential for accurate identification, particularly in early infestation stages where mistletoe blends seamlessly with its host trees. The original aerial image on the left of Figure 2 highlights this visual similarity, demonstrating the challenge of differentiating mistletoe from healthy vegetation. To address this, preprocessing steps such as multispectral band co-registration and normalization were applied to ensure consistency across the dataset.

The identification and segmentation of mistletoe clusters in each set of images were performed manually by an ecological specialist. These labeled maps served as the foundation for supervised learning, providing accurate ground-truth data to train machine learning models. The left image of Figure 2 shows an example of segmentation map overlaid (in red) to the original corresponding image. This combination of high-resolution imagery and expert segmentation created a robust dataset for advancing automated detection methodologies.

The integration of AI and remote sensing holds promise not only for addressing the immediate challenge of mistletoe infestations, but also for broader ecological applications. These tools enable conservationists to monitor biodiversity, assess the health of ecosystems, and develop targeted management strategies for invasive species. As highlighted in this research, the use of these technologies requires a multidisciplinary approach that combines ecological expertise with computational advancements, ensuring that solutions are both scientifically robust and practically actionable.

The objectives of this paper are threefold. First, to reflect on the insights gained from the ongoing collaborative efforts of our interdisciplinary research team, emphasizing both successes and challenges encountered in applying AI and remote sensing to ecological conservation. Second, to encourage dialogue within the scientific and conservation communities, fostering interdisciplinary collaboration and inspiring innovation. Finally, to outline po-

Figure 1. Visual representation of mistletoe clusters (*S. interruptus*) highlighted in blue, captured in the Ramón López Velarde Garden, Mexico City in November 2024. The hemiparasitic mistletoe blends seamlessly with its host vegetation, complicating manual detection and emphasizing the need for automated identification methods.

tential directions for future research, emphasizing the importance of sustainable and equitable deployment of AI technologies in addressing pressing environmental challenges.

Current Advances

The initial phase of our research explored the potential of evolutionary computating approaches, specifically Genetic Programming (GP), to address the challenge of mistletoe detection. GP was employed to develop a spectral index method capable of distinguishing mistletoe infestations in UAV-acquired multispectral imagery [6]. By simulating natural evolutionary processes, GP iteratively optimized combinations of spectral bands to create a highly effective index for detecting *Phoradendron Velutinum* species during its flowering stage.

This methodology demonstrated remarkable accuracy, exceeding 96% in controlled conditions, and provided a foundational proof of concept for automated detection. The focus on the flowering stage of *P. velutinum* leveraged the unique spectral characteristics of the parasitic plant, enabling precise differentiation from host trees and surrounding vegetation. This achievement evidenced the potential of computational approaches to complement traditional ecological methods, particularly in addressing labor-intensive tasks such as manual field inspections.

The success of GP not only validated the feasibility of using automated spectral analysis for mistletoe detection, but also highlighted the complexity of scaling such methods to broader ecological contexts and the monitoring of different mistletoe species. These insights set the stage for subsequent explorations into more adaptable techniques capable of handling environmental variability and addressing a wider range of species. Hence, our research expanded to explore the use of texture descriptors and machine learning techniques for mistletoe detection.

Recognizing the importance of spatial patterns in differentiating mistletoe from host trees, we employed descriptors such as Gray Level Co-occurrence Matrices (GLCM) and Gabor filters, along with color transformations and vegetation indices, to capture the unique

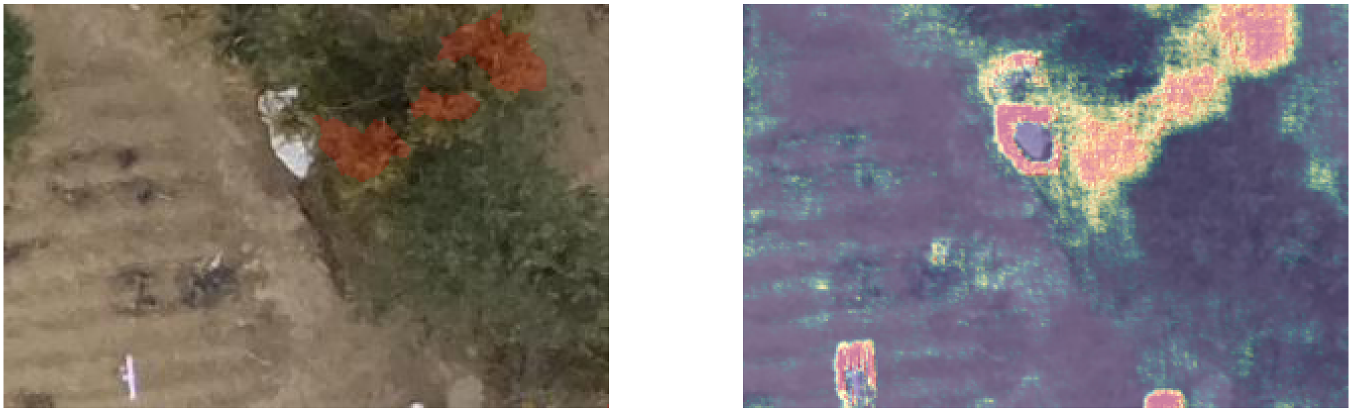

Figure 2. Visualization of the ResNet-34 CNN model's processing for the identification of mistletoe species known as *P. Velutinum*. On the left, the original high-resolution RGB image with an overlaid binary segmentation mask (red) highlighting mistletoe regions (expert-labeled), and on the right, the CNN-generated heatmap emphasizing areas identified as mistletoe.

textural and spectral characteristics of mistletoe clusters [8]. These features were selected with the aim of improving the classification performance of *S. Interruptus* species carried out by Support Vector Machines (SVMs), which achieved an accuracy close to 60% in controlled environments [8]. These results demonstrated the potential of combining texture and spectral descriptors with machine learning models for the current application. However, the performance of these methods was significantly affected by the close similarity between the mistletoe and host tree foliage. Moreover, the lack of other spectral features limited the adaptability of our approach to more complex and heterogeneous settings, such as large forested areas with various mistletoe species.

The results of SVM revealed the potential of using computational techniques to classify mistletoe infestations, but also evidenced their limitations in adapting to environmental variability and capturing the complexity of texture patterns. These challenges highlighted the need for methodologies capable of learning intricate features directly from the data, without relying on predefined handcrafted descriptors. This realization led us to explore deep learning approache, which offered a more dynamic and adaptive framework for handling complex ecological tasks.

Among deep learning techniques, Convolutional Neural Networks (CNNs) emerged as a natural and highly suitable choice. CNNs are specifically designed to extract spatial hierarchies of features from images, making them particularly effective for texture-rich tasks, like mistletoe detection. By learning patterns directly from UAV-acquired imagery, CNNs provided the flexibility needed to account for variability in environmental conditions, such as changes in lighting, foliage density, and tree structure. This adaptability made CNNs an ideal next step in our research, bridging the gap between earlier efforts and the need for robust and scalable solutions.

With the purpose of testing the feasibility of advanced image-based classification techniques, we began by the development of a simpler application oriented to the detection of dead trees in urban parks [9]. This task allowed us to focus on well-defined visual features, such as the absence of foliage or grenness, while making use of UAV-acquired data to establish baseline models. To do so, we implemented ResNet-34, a convolutional neural network (CNN) architecture known for its balance between computational efficiency and representational power [10, 11, 12]. ResNet-34 utilizes residual blocks with shortcut connections, which mitigate the vanishing gradient problem and enable effective training in deep networks [13]. This architecture excels in extracting fine-grained features, making it particularly suitable for ecological monitoring tasks where subtle distinctions, such as identifying invasive species or degraded vegetation, are critical.

This phase of our research not only confirmed the potential of CNNs for ecological and vegetation monitoring, but also laid the groundwork for subsequent efforts to capture more subtle texture patterns associated with mistletoe infestations. By successfully applying CNNs to the detection of dead trees, we were able to refine our methodologies and build confidence in the scalability and adaptability of deep learning models for broader ecological challenges, such as mistletoe species classification.

Our research then progressed to implement ResNet-34 for the more intricate task of mistletoe classification in urban green spaces. To deepen our understanding of the visual features and characteristics associated with such parasitic plant, we conducted experiments that perturbed the input images with changes in brightness, noise, and resolution. These perturbations were designed to explore how variations in environmental conditions affected the model's ability to identify mistletoe clusters. This approach provided valuable insights into the specific visual features that ResNet-

34 prioritized during classification, such as color, texture, and spatial patterns, enhancing our ecological understanding of mistletoe and its distinct visual signatures.

In addition, we integrated explainability tools, such as class activation maps (also known as heatmaps), to further validate the model predictions. These visualizations highlighted the regions of the images that the model relied on for its classifications, allowing us to ensure that its decisions aligned with ecological expectations. This combination of perturbation analysis and explainability not only strengthens the model interpretability, but also intends to provide end-users, such as forestry experts and conservation practitioners, with actionable insights into the visual characteristics of mistletoe, facilitating more informed ecological management.

The implementation of convolutional neural network models significantly improved our ability to identify mistletoe clusters; however, these models also introduced substantial computational demands. Training and deploying CNNs at scale requires considerable processing power, particularly when handling high-resolution UAV imagery. This challenge was addressed through our collaboration with the UNAM-Huawei Innovation Space, which provided access to high-performance computational infrastructure, enabling the optimization and efficient training of our models. Nevertheless, certain experiments required an evaluation on diverse hardware configurations and computational environments to assess performance and adaptability across different scenarios. For instance, the optimal combination of hardware setups, resource allocation strategies, and system configurations to minimize training times remains inadequately understood, highlighting the need for further exploration into hardware-performance trade-offs for ecological deep learning applications.

To address the previous exploration, we conducted a comparative analysis of hardware platforms, evaluating configurations ranging from consumer-grade laptops to high-performance servers. The study revealed critical insights into the trade-offs between computational efficiency, cost, and model performance. High-performance workstations and specialized hardware provided significant reductions in training time and improved the scalability of our models. By optimizing hardware performance, we not only increased the feasibility of deploying deep learning models for mistletoe detection, but also set the stage for broader applications in ecological conservation. This phase of research reinforced the need for interdisciplinary collaborations and resource-sharing initiatives to overcome computational barriers, paving the way for more robust and scalable solutions to complex environmental challenges.

A primary goal of our research has been to develop scalable and practical tools for ecological conservation, transitioning from experimental methodologies to real-world applications. Our detection systems, ranging from infestation maps to automated classification models, are designed with usability in mind, addressing the needs of forestry personnel, urban park managers, and policymakers. These tools have the potential to automate labor-intensive tasks, provide actionable insights for resource allocation, and monitor the effectiveness of management strategies. Additionally, explainability tools, such as class activation maps, ensure interpretability and trust in the models' outputs, bridging the gap between advanced computational methods and practical conservation efforts.

Discussion

The advancements achieved so far in this initiative represent significant progress toward addressing the ecological challenges posed by mistletoe infestations. However, as with any innovative approach, the journey has revealed not only successes but also limitations that warrant critical reflection. Beyond the technical achievements, this work raises broader questions about the role of AI in ecological conservation, the ethical and practical implications of deploying such technologies, and the challenges of ensuring scalability and accessibility. This section aims to provide an analysis of these aspects, exploring the lessons learned, the unresolved challenges, and the opportunities for future innovation. By reflecting on these dimensions, we hope to inspire continued dialogue and interdisciplinary collaboration in the pursuit of sustainable and impactful conservation solutions.

The methodologies explored reflect an evolving response to the complex challenges posed by mistletoe infestations. GP demonstrated significant potential by automating the design of a spectral index that achieved an overall accuracy of 96.6% in detecting mistletoe infestations during the flowering stage of *P. velutinum*. This precision was achieved in controlled conditions with multispectral UAV imagery, underscoring GP's ability to derive tailored solutions for specific ecological challenges. However, its dependency on flowering-stage spectral characteristics highlights limitations in generalizing across variable conditions and phenological stages.

This success was juxtaposed with challenges in scalability and adaptability. While GP excelled in specific scenarios, its reliance on spectral indices limited its robustness in heterogeneous settings. The detailed analysis of GP solutions revealed a strong reliance on visible spectrum bands, particularly the R and B channels, aligning with the characteristic pigmentation of mistletoe. This specificity, while effective in the targeted study area, raises questions about the adaptability of the GP-based algorithms to other environments and mistletoe species.

Subsequent exploration of texture descriptors and machine learning classifiers represented an important step in addressing the spatial complexity of mistletoe infestations. Techniques such as Gray Level Co-occurrence Matrices (GLCM), Gabor filters, and Local Binary Patterns (LBPs) were employed to capture the distinct textural characteristics of mistletoe regions, complemented by vegetation indices derived from UAV-acquired multispectral imagery. These approaches showed the critical role of texture in differentiating mistletoe from healthy vegetation, particularly in controlled scenarios. However, classification experiments conducted with SVMs, revealed notable challenges. These included sensitivity to environmental variability, such as lighting and seasonal changes, as well as the inherent limitations of handcrafted features, which restricted scalability and generalization to heterogeneous ecosystems. While the methodology offered valuable insights, the results highlighted the necessity of more adaptable techniques capable of learning features directly from data, setting the stage for deep learning approaches.

Deep learning techniques, while offering greater adaptability, require large, annotated datasets to fully exploit their potential. Our dataset, however, was of limited size, posing significant constraints for training CNNs to perform tasks like semantic segmentation, which demand pixel-level annotations. To address this limitation, we adopted an image classification approach by dividing larger aerial images into smaller tiles for analysis. This method maximized the utility of the available data and leveraged the robust feature extraction capabilities of CNN architectures such as ResNet-34. Rather than immediately addressing the complex task of distinguishing mistletoe from green vegetation, we opted for a more incremental approach. By focusing initially on the classification of dead trees versus green vegetation, we refined our methodologies, built a solid foundation for model training, and validated the feasibility of our approach. This stepwise progression balanced the technical challenges of deep learning implementation with the practical constraints of limited data availability and ecological complexity.

Our research achieved a significant milestone with the classification of dead trees using UAV-acquired multispectral imagery. The comparative study of ResNet-34 and DenseNet-121 provided insights into the applicability of these architectures for detecting vegetation anomalies indicative of ecological disturbances. Both models demonstrated robust performance, with DenseNet-121 slightly outperforming ResNet-34 in metrics such as accuracy and F1-score, achieving values of approximately 97% [9]. This high level of precision was complemented by the integration of heatmaps, which illuminated the areas within input images most relevant to the classification process. These visual tools enhanced the interpretability of the models' decisions, bridging the gap between computational outputs and actionable ecological insights. While these results validated the feasibility of using CNNs for tree health monitoring, they also set a precedent for tackling more complex classification tasks, such as distinguishing mistletoe infestations from healthy vegetation.

Building on the success of dead tree classification, we transitioned to the more intricate task of identifying mistletoe infestations. Confident in our approach, we extracted tiles in the same manner as before, this time focusing specifically on mistletoe. These tiles, derived from UAV-acquired multispectral imagery, formed the basis for training the model to distinguish mistletoe (*S. interruptus*) from the surrounding vegetation. Figure 3 showcases representative tiles of mistletoe, capturing the distinct visual features necessary for classification.

The ResNet-34 model achieved notable accuracy, with baseline performance metrics consistently exceeding 83% across both hold-out and k-fold validation strategies. Heatmap visualizations effectively highlighted biologically relevant features, aligning with annotated mistletoe regions and validating the ecological relevance of the model's classifications.

However, the study also highlighted specific limitations. The model exhibited sensitivity to environmental perturbations, such as changes in color channels and the introduction of noise, particularly Salt-and-Pepper noise at higher intensities. While this sensitivity underscores areas for further refinement—such as enhancing robustness through noise-tolerant training methodologies or adaptive preprocessing strategies—it also offers a unique opportunity to better understand the visual characteristics of mistletoe. By analyzing the model's responses under varying conditions, we gain deeper insights into the defining features of mistletoe, such as its texture, color, and spatial patterns. This enhanced understanding can inform both manual identification by ecological practitioners and the development of more sophisticated automated detection methods. Additionally, the reliance on tile-based classification, while practical given the dataset constraints, presents challenges in capturing the broader spatial context necessary for large-scale ecological monitoring.

Looking forward, these results highlight the potential for integrating semantic segmentation techniques to overcome the spatial limitations of tile-based classification. By advancing the model's capacity to identify mistletoe at a pixel level, future efforts could significantly enhance detection precision and scalability. However, the computational demands of such refined methodologies remark the importance of efficient hardware solutions. As deep learning models grow in complexity, addressing the trade-offs between computational efficiency and performance becomes essential to ensuring their feasibility for large-scale ecological applications.

Our results provide valuable insights into optimizing the computational demands of deep learning models for mistletoe detection. Multiple configurations were evaluated, including Intel and AMD processors, consumer-grade GPUs, and high-performance workstations, to assess their suitability for training ResNet-34. Notably, AMD processors consistently outperformed their Intel counterparts in training efficiency, with mid-range AMD setups striking a balance between cost and performance. Additionally, disabling hy-

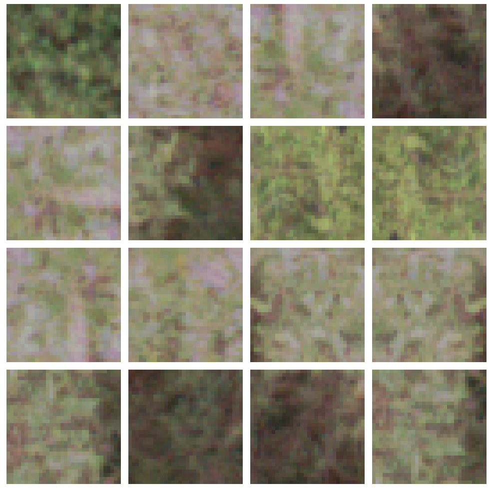

Figure 3. Sample tiles used in the training dataset for mistletoe (*S. interruptus*). Each tile measures 32×32 pixels and corresponds to a segment extracted from the original aerial imagery. These tiles were used to train and validate the ResNet-34 model.

perthreading (HT) on CPUs accelerated training times across most platforms, highlighting the importance of carefully tuning hardware settings for deep learning tasks.

Consumer-grade GPUs demonstrated effectiveness for short-term training tasks, delivering competitive performance relative to high-performance workstations in scenarios with moderate computational demands. However, high-performance servers exhibited unparalleled scalability and reduced training times, making them indispensable for extensive datasets or long-term workloads. These findings highlight the trade-offs between computational efficiency and cost, emphasizing the need to select hardware configurations that align with the scale and scope of ecological applications.

These outcomes reinforce the feasibility of deploying advanced models for ecological monitoring, provided that hardware choices are optimized for specific operational constraints. They naturally transition to broader considerations of deployment, such as addressing accessibility challenges in under-resourced regions and balancing computational performance with environmental sustainability.

While our methodologies have demonstrated significant potential, they are not without limitations, many of which present critical challenges for broader implementation and scalability. A key constraint lies in the reliance on high-resolution UAV-acquired imagery, which, while enabling precise detection, demands substantial resources for data collection and processing. This dependence raises questions about the feasibility of deploying such methods in resource-limited regions or over large-scale ecosystems where UAV coverage and computational capacity may be restricted.

Another limitation is the challenge of generalizing our models to diverse ecological contexts. The high accuracy achieved in controlled or semi-controlled environments does not guarantee equivalent performance in more complex and heterogeneous landscapes. Variations in vegetation types, mistletoe species, and environmental conditions, such as lighting or seasonal changes, can significantly affect model robustness. This highlights the need for extensive field validation and the development of models capable of adapting to the dynamic nature of ecological systems.

Scaling these approaches to larger areas also introduces unresolved challenges, including the computational demands of processing expansive datasets and the need for automation in UAV flight planning and data annotation. Moreover, biases in the training data—stemming from an overrepresentation of specific ecological conditions or geographic areas—may inadvertently limit the models' applicability. Addressing these biases will require careful dataset curation, with a focus on improving representativeness and diversity in training samples.

These gaps invite further exploration and innovation within the scientific and conservation communities. How can we balance the need for high-resolution data with the goal of creating scalable and accessible solutions? Are there hybrid approaches, such as integrating satellite imagery for broader coverage with UAV data for fine-tuning, that could bridge this gap? Additionally, adaptive learning techniques, where models are continuously updated with new data from diverse environments, could offer a path toward greater generalizability and resilience.

By acknowledging these limitations, we aim to stimulate debate and collaboration on how to address them. While the challenges are significant, they also present opportunities to refine and expand the applicability of AI-driven conservation tools, paving the way for more inclusive and impactful ecological solutions.

The methodologies and insights developed in this research hold significant potential for advancing AI applications beyond the detection of mistletoe infestations. The integration of UAV-acquired imagery, machine learning, and deep learning techniques establishes a scalable framework that can be adapted for a variety of ecological challenges. For example, these approaches could be utilized to monitor biodiversity in threatened habitats, assess the spread of other invasive species, or evaluate the health of forests affected by climate change. The ability of AI models to process and analyze large volumes of high-resolution imagery positions them as transformative tools in ecological conservation.

However, the broader application of these technologies raises important ethical and practical questions. Data ownership and privacy concerns are particularly relevant when UAVs are deployed in areas with human activity, potentially capturing unintended information. Establishing clear protocols for data governance and ensuring compliance with privacy regulations will be essential to maintain public trust and ethical integrity. Similarly, the accessibility of advanced AI tools in under-resourced areas remains a pressing challenge. High-performance computing and UAV infrastructure are often limited to well-funded institutions, creating disparities in who can benefit from these innovations. Addressing these inequities will require collaborative efforts to design cost-effective solutions and share resources more equitably.

Environmental sustainability is another critical consideration. UAV flights, particularly those involving multiple missions, contribute to carbon emissions, while the energy-intensive nature of high-performance computing adds to the environmental footprint of AI applications. These factors remark the importance of optimizing resource use and exploring greener alternatives, such as edge computing or more energy-efficient hardware, to align technological progress with conservation principles.

Looking forward, this research invites critical reflection on the sustainability and inclusivity of AI in ecological applications. How can we ensure that these tools are accessible to under-resourced regions and adaptable to diverse ecological contexts? Are there ways to integrate traditional ecological knowledge with AI systems to create more holistic and inclusive conservation strategies? Furthermore, what frameworks are needed to minimize the environmental impact of deploying these technologies at scale?

By raising these questions, we aim to provoke thoughtful discussion within the scientific and conservation communities. While AI holds immense promise for addressing complex ecological challenges, its responsible application requires a balance between innovation and sustainability. Through interdisciplinary collaboration and ethical foresight, these tools can evolve to become more inclusive, equitable, and aligned with the long-term goals of conservation.

The advancements achieved in this research lay a solid foundation for exploring new avenues in ecological monitoring and conservation. Expanding the scope of testing beyond mistletoe detection to other ecological challenges offers promising opportunities to further validate and refine these methodologies. For instance, the same UAV-acquired imagery and machine learning frameworks could be adapted to detect other invasive species, such as bark beetles or parasitic vines, which similarly threaten forest health. Additionally, the texture and spectral analysis techniques employed here could be applied to assess broader forest health indicators, such as identifying early signs of disease or monitoring vegetation recovery after disturbances.

Emerging technologies present exciting possibilities for enhancing the scalability and precision of these efforts. Vision Transformers, a state-of-the-art architecture for image analysis, could be employed to capture complex spatial relationships and improve detection accuracy in highly heterogeneous environments. By leveraging their ability to process global context within imagery, Vision Transformers may provide a more comprehensive understanding of ecological patterns.

A critical next step is the piloting of these methodologies in real-world deployments. Collaborations with local conservation practitioners and forestry agencies will be essential to adapt these tools to the specific needs and constraints of field applications. Pilot programs could serve as testing grounds for refining workflows, addressing usability challenges, and building trust among end-users. Additionally, partnerships with community-based organizations could facilitate the integration of local ecological knowledge, enriching the data and increasing the relevance of AI-driven insights.

Interdisciplinary collaboration was a cornerstone of this project, driving both its successes and its evolution in addressing the ecological challenge of mistletoe infestations. By integrating expertise from CentroGeo, ITM, and ITPA, as well as leveraging partnerships with SECTEI, SEDEMA, CORENA, and the UNAM-Huawei Innovation Space, the team successfully bridged technical innovation with ecological applications. This collaborative framework facilitated the development of advanced methodologies, from the design of spectral indices to the implementation of deep learning models, ensuring that each phase of the project was informed by a diverse range of perspectives and skills. Similarly, access to high-performance computing resources through institutional partnerships allowed the team to overcome significant computational challenges, accelerating the refinement of deep learning models and enabling detailed performance evaluations.

However, this collaboration also revealed inherent challenges that underscore the complexity of interdisciplinary research. Differences in disciplinary priorities and methodologies occasionally led to misalignments in project goals and expectations. For example, while technical partners prioritized computational efficiency and accuracy, ecological collaborators often emphasized the interpretability and field applicability of the tools being developed. Effective communication and iterative feedback loops are essential in resolving these tensions, but they also highlighted the need for clearer frameworks to align multidisciplinary objectives from the outset.

Reflecting on these experiences, several lessons emerge that could inform future interdisciplinary projects. First, establishing clear, shared goals at the beginning of a project can help align priorities and reduce friction between disciplines. Second, creating structured communication channels—such as regular workshops or collaborative platforms—can facilitate ongoing dialogue and ensure that all voices are heard throughout the research process. Finally, fostering a culture of mutual respect and openness to diverse perspectives is critical for maximizing the potential of interdisciplinary teams.

A key priority for future research and implementation is fostering dialogue that not only refines AI applications for mistletoe detection but also raises awareness of the ecological threat it poses. This dialogue must extend beyond scientists and technical experts to include policymakers, conservation practitioners, and the general public. While mistletoe is often romanticized, it is essential to communicate its grave impact on forests, where unchecked infestations weaken trees, reduce biodiversity, and disrupt critical ecosystem services. Engaging communities in this conversation can help shift perceptions and build support for necessary conservation measures.

Involving local communities is particularly important, as their traditional ecological knowledge can enhance AI-driven approaches and ground them in practical realities. Collaborative frameworks are needed to ensure that detection tools and management strategies are accessible and actionable, particularly in under-resourced areas where mistletoe infestations could be most severe. Furthermore, public education initiatives can help bridge the gap between the cultural symbolism of mistletoe and the ecological urgency of its control, fostering a shared understanding of the problem and the solutions required.

As we advance these efforts, it is crucial to ensure that AI applications remain equitable, transparent, and aligned with ecological stewardship. The future of mistletoe management depends not only on technological innovation but also on our ability to engage and empower all stakeholders, from conservation scientists to the broader community. By combining advanced tools with widespread awareness and collaborative action, we can protect forests from the devastating impact of mistletoe infestations while fostering sustainable conservation practices.

Conclusions

This work critically examines the role of AI and remote sensing in addressing the ecological challenges posed by mistletoe infestations, offering insights into the methodologies and approaches that have shaped this emerging field. By leveraging GP, texture descriptors, and CNNs, our research highlights both the potential and limitations of these technologies in ecological conservation. GP demonstrated its effectiveness in developing spectral indices tailored to specific ecological challenges, providing a starting point for automation in mistletoe detection. Similarly, texture-based machine learning methods, while valuable for capturing spatial patterns, evidenced the constraints of handcrafted features in complex and heterogeneous environments.

The shift toward deep learning, particularly through CNNs, reflects an evolution in methodology that not only addressed earlier limitations but also revealed the scalability and adaptability of AI-driven solutions for conservation. These advancements serve as a critical reflection on the application of cutting-edge technologies in ecological monitoring, illustrating how AI can transform labor-intensive tasks into scalable, efficient processes. However, they also raise important questions about the practicality and accessibility of these approaches, particularly in resource-limited settings.

While the methodologies and insights presented in this work demonstrate significant potential for addressing ecological challenges, they are accompanied by critical limitations that warrant further exploration. Environmental variability complicates the generalization of these models. The performance of machine learning and deep learning approaches, optimized for specific ecological contexts, may degrade in heterogeneous environments characterized by diverse vegetation types, seasonal changes, or fluctuating light conditions. Addressing this variability will require adaptive methodologies capable of learning from diverse datasets, coupled with extensive field validation to ensure robustness across ecological settings.

Computational demands also present a significant challenge, particularly for resource-intensive deep learning models. The energy consumption and infrastructure required to train and deploy these systems raise important questions about their environmental and practical sustainability. Collaborative efforts, such as leveraging edge computing for real-time analysis or optimizing hardware configurations for efficiency, represent critical avenues for mitigating these constraints.

The integration of AI into ecological conservation demands sustained interdisciplinary collaboration, grounded in both ethical foresight and practical deployment strategies. The complexities of mistletoe detection and broader ecological challenges require expertise from diverse fields, including computational sciences, ecology, and social sciences, to ensure that technological advancements remain contextually relevant and inclusive. Future efforts must prioritize partnerships that bridge the gap between cutting-edge research and actionable conservation practices, fostering collaboration among researchers, policymakers, and local communities.

The practical deployment of AI tools also presents critical opportunities for innovation. Real-world implementation, supported by field trials and user-centric design, will be key to refining these technologies and ensuring their usability by non-specialist stakeholders, such as park managers and conservation workers. Integrating emerging technologies like Vision Transformers and edge computing could further enhance the adaptability and efficiency of AI-driven solutions, making them more accessible and sustainable.

Looking forward, the role of AI in conservation is poised to expand far beyond mistletoe detection. These tools offer immense potential to tackle diverse challenges, from monitoring biodiversity to mitigating the impacts of climate change. However, their impact will ultimately depend on our ability to align technological progress with the principles of sustainability, equity, and transparency.

By fostering interdisciplinary dialogue and embracing a forward-thinking approach, the scientific community can advance AI applications that not only address immediate ecological threats but also contribute to the broader goal of sustainable development. The path ahead is both challenging and filled with opportunity, requiring collective action, ethical reflection, and a commitment to collaboration that transcends disciplinary boundaries. Through these efforts, AI can become a transformative force in conservation, driving innovation while safeguarding the planet for future generations.

CRediT authorship contribution statement

Juan Carlos Valdiviezo-Navarro: Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing. **Mauricio G. Orozco-del-Castillo:** Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT in order to improve readability. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

We express our gratitude to Secretaría del Medio Ambiente (SEDEMA) for the support during field work. We also recognize the efforts of our team to develop this project: Paola Andrea Mejía Zuluaga, León Felipe Dozal García, Alejandra Aurelia López Caloca, Alejandro Téllez Quiñones, Adriana Lizzette Luna Nieves, Amilcar Morales Gamas, and Mario Ledesma Arriola (CentroGeo); Nivia Iracemi Escalante García, Eduardo Flores Gallegos, and José Ernesto Olvera González (TecNM/ITPA); Nora L. Cuevas Cuevas, Carlos Bermejo-Sabbagh, William Orozco-González, and Pedro A.G. Ortiz-Sánchez (TecNM/ITM). Part of this work has been financed by Secretaria de Educación, Ciencia, Tecnología e Innovación (SECTEI) of Mexico City through Grant Number 3581c24.

References

[1] P. K. Rai and J. Singh, "Invasive alien plant species: Their impact on environment, ecosystem services and human health," *Ecological Indicators*, vol. 111, p. 106020, 2020.

[2] B. Gallardo, M. Clavero, M. I. Sánchez, and M. Vilà, "Global ecological impacts of invasive species in aquatic ecosystems," *Global Change Biology*, vol. 22, 2016.

[3] M. Pompa-García and J. A. Rodríguez, "Variation of carbon uptake from forest species in mexico: a review," *Madera Y Bosques*, vol. 23, pp. 225–235, 2017.

[4] N. Mesa-Sierra, M. de la Peña-Domene, J. Campo, and C. Giardina, "Restoring mexican tropical dry forests: A national review," *Sustainability*, 2022.

[5] D. Cibrián Tovar, "Diagnóstico fitosanitario de áreas verdes infestadas por muérdago en la Ciudad de México," 2016.

[6] P. A. Mejia-Zuluaga, L. Dozal, and J. Valdiviezo-N., "Genetic programming approach for the detection of mistletoe based on uav multispectral imagery in the conservation area of mexico city," *Remote Sensing*, vol. 14, p. 801, 2022.

[7] T. Paap, T. Burgess, and M. Wingfield, "Urban trees: bridge-heads for forest pest invasions and sentinels for early detection," *Biological Invasions*, vol. 19, pp. 3515–3526, 2017.

[8] P. A. Mejia-Zuluaga, J. C. Valdiviezo-N., and L. Dozal, "Texture descriptors and machine learning algorithms for mistletoe detection in urban forests using multispectral imagery," in *Remote Sensing*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:264299698

[9] W. Orozco-González, J. C. Valdiviezo-Navarro, M. G. O. del Castillo, P. A. Mejia-Zuluaga, and L. Dozal, "Classification of dead trees in urban parks using aerial imagery and convolutional neural networks," in *ICCBR Workshops*, 2024. [Online]. Available: https://api.semanticscholar.org/CorpusID:270701684

[10] A. Abrol, A. Sharma, K. Karnic, and R. Ranjan, "Human activity recognition using resnet-34 model," *International Journal of Recent Technology and Engineering*, vol. 10, no. 5, pp. 224–228, 2021.

[11] X. Li, X. Xu, X. He, X. Wei, and H. Yang, "Intelligent crack detection method based on gm-resnet," *Sensors (Basel, Switzerland)*, vol. 23, 2023.

[12] W. Liu, L. Jiangtao, G. Zhao, L. Sun, H. Wang, W. Li, and B. Sun, "Improvement of cifar-10 image classification based on modified resnet-34," in *Advances in Natural Computation, Fuzzy Systems and Knowledge Discovery*, 2020.

[13] S. Salman and X. Liu, "Sparsity as the implicit gating mechanism for residual blocks," in *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2019, pp. 1–6.

Journal of Artificial Intelligence and Computing Applications (2024) 2(2): 18–29

Perspectiva crítica

El papel de la IA en el combate a la infestación de muérdago en los bosques mexicanos: un llamado a la acción interdisciplinaria

Juan Carlos Valdiviezo-Navarro
CONAHCYT-CentroGeo, Yucatán, Mexico

Mauricio G. Orozco-del-Castillo
Tecnológico Nacional de México/IT de Mérida, Yucatán, Mexico

This article has been published in both English and Spanish to ensure accessibility to a broader audience. The English version of this article can be accessed at the following DOI: https://doi.org/10.5281/zenodo.14699764.

Palabras clave: detección de muérdago, inteligencia artificial en conservación, herramientas de gestión ecológica

PERSPECTIVA

El muérdago es una planta hemiparásita que se adhiere a los árboles huéspedes, extrayendo agua y nutrientes, lo que debilita los árboles, reduce la biodiversidad y altera los servicios ecosistémicos críticos. Si bien a menudo se romantiza en contextos culturales, el muérdago plantea una grave amenaza ecológica, particularmente en los ecosistemas urbanos y boscosos donde las infestaciones pueden propagarse sin control. Este artículo examina el papel de la inteligencia artificial (IA) y la teledetección para abordar los desafíos de detección y gestión que plantea el muérdago. A través de una evaluación crítica de metodologías que van desde el aprendizaje automático basado en texturas hasta modelos avanzados de aprendizaje profundo como ResNet-34, este artículo reflexiona sobre los éxitos, las limitaciones y las implicaciones de estos enfoques. Nuestra investigación interdisciplinaria destaca el potencial transformador de combinar la IA con la experiencia ecológica para desarrollar herramientas escalables y eficientes para la conservación. Sin embargo, también identificamos desafíos clave, incluida la necesidad de acceso equitativo, consideraciones éticas y escalabilidad en diversos contextos ecológicos. Además, enfatizamos la importancia de involucrar a la comunidad en general, ya que los conceptos erróneos sobre el muérdago obstaculizan los esfuerzos de conservación. Al integrar la conciencia pública con los avances tecnológicos, abogamos por un enfoque equilibrado y sostenible para la gestión ecológica. Este documento tiene como objetivo provocar un diálogo crítico e inspirar estrategias viables para aprovechar la IA para abordar los desafíos de conservación global.

Introducción

Las infestaciones de muérdago representan un ejemplo localizado de un problema global más amplio: los desafíos ecológicos y socioeconómicos que plantean las especies de plantas invasoras. Estas especies, caracterizadas por su capacidad para establecerse y proliferar rápidamente en ecosistemas no nativos, tienen impactos significativos y de gran alcance. A nivel mundial, las plantas invasoras contribuyen a la pérdida de biodiversidad al superar a las especies nativas, alterar las redes alimentarias y alterar los servicios ecosistémicos [1, 2]. En este sentido, su propagación puede comprometer el secuestro de carbono, la regulación de la temperatura y el mantenimiento del hábitat para la vida silvestre nativa, al tiempo que crea cargas socioeconómicas al afectar la agricultura, la silvicultura y la salud pública.

El muérdago es una planta hemiparásita que requiere un árbol huésped para extraer agua, minerales y otros nutrientes mediante la inserción de una raíz especializada llamada *haustoria*. Esta relación parasitaria debilita los árboles hospedantes con el tiempo, reduciendo su estabilidad estructural, disminuyendo la biodiversidad y amenazando servicios ecosistémicos críticos [3, 4]. En la Ciudad de México, tres especies de muérdago—*Cladocolea loniceroides*, *Phoradendron velutinum* y *Struthanthus interruptus*—han sido identificadas como contribuyentes importantes a la alteración ecológica y uno de los principales problemas de los espacios verdes urbanos [5].

La detección y el manejo de infestaciones de muérdago presentan desafíos importantes debido a la capacidad de la planta para mezclarse visualmente con la vegetación huésped, particularmente en sus primeras etapas [6]. Como se muestra en la Figura 1, los grupos de muérdago se parecen mucho al follaje sano de sus árboles hospedantes, lo que di-

Correo electrónico: mauricio.orozco@itmerida.edu.mx
https://doi.org/10.5281/zenodo.14720118

ficulta la identificación manual, especialmente en vegetación densa. Este desafío se ve agravado aún más por la variabilidad ambiental, incluidos los cambios en la iluminación, los ángulos de visión y la densidad del follaje estacional, que reducen la confiabilidad de los métodos tradicionales [7]. Las encuestas manuales, realizadas por expertos capacitados, siguen siendo el enfoque principal para la detección del muérdago, pero requieren mucho desarrollo manual, tiempo y recursos, lo que las hace poco prácticas para una implementación a gran escala [6]. Además, la inaccesibilidad física de ciertas áreas boscosas exacerba estas ineficiencias, permitiendo que las infestaciones se propaguen sin ser detectadas. En conjunto, estas limitaciones destacan las importantes dificultades que enfrentamos en el manejo de las infestaciones de muérdago y resaltan la necesidad apremiante de mejores métodos para mejorar la eficiencia y escalabilidad de los esfuerzos de detección.

Entre estas especies, *S. interrumpus* y *P. velutinum* se destacan por su amplio impacto en los ecosistemas urbanos y áreas boscosas. Por ejemplo *S. interruptus* prevalece en 9 de las 16 alcaldías de la Ciudad de México [5], mientras que *P. velutinum* se destaca particularmente por su capacidad de combinarse perfectamente con la vegetación huésped en los bosques de conservación [6]. Ambas especies forman grupos densos en etapas avanzadas de infestación, dominando las copas de los árboles y alterando las estructuras naturales de los árboles [7]. Estas infestaciones comprometen la salud de los bosques, reducen la biodiversidad y obstaculizan los servicios ecosistémicos esenciales, incluido el secuestro de carbono y el apoyo al hábitat de especies nativas [3, 4]. Los impactos ecológicos y económicos generalizados de estas especies de muérdago enfatizan la urgencia de desarrollar estrategias de detección y gestión eficaces y escalables.

La urgente necesidad de soluciones innovadoras y escalables para abordar el problema de la infestación de muérdago en México catalizó la formación de una iniciativa de investigación interdisciplinaria. Este esfuerzo colaborativo reunió a expertos del Centro de Investigación en Ciencias de Información Geoespacial, A.C. (CentroGeo), el Tecnológico Nacional de México/Instituto Tecnológico de Mérida (ITM), el Instituto Tecnológico de Pabellón de Arteaga (ITPA), e incluyó el apoyo de la Secretaría de Educación, Ciencia, Tecnología e Innovación de la Ciudad de México (SECTEI), Secretaría del Medio Ambiente de la Ciudad de México (SEDEMA) y el Espacio de Innovación UNAM-Huawei. Combinando experiencia en inteligencia artificial (IA), teledetección, gestión ecológica y operaciones de campo, el equipo busca abordar los desafíos multifacéticos que plantean las infestaciones de muérdago en ecosistemas tanto urbanos como boscosos.

Un elemento central de la estrategia del equipo es la integración de tecnologías avanzadas con conocimientos ecológicos. En particular, las tecnologías de detección remota, como las imágenes multiespectrales de alta resolución recopiladas mediante vehículos aéreos no tripulados (UAV, de *Unmanned Aerial Vehicles*), han brindado la capacidad de monitorear grandes áreas de manera eficiente y capturar información detallada relacionada con la salud de la vegetación y la fenología, entre otros. Estos datos, combinados con técnicas de IA de vanguardia, han permitido el desarrollo de herramientas de detección escalables. Por lo tanto, aprovechando los algoritmos de aprendizaje automático y los modelos de aprendizaje profundo, el equipo pretende crear metodologías capaces de distinguir los grupos de muérdago de la vegetación saludable circundante con alta precisión, incluso en condiciones ambientales desafiantes.

El objetivo general de la iniciativa es desarrollar metodologías innovadoras impulsadas por la tecnología para la detección y gestión automatizadas de las infestaciones de muérdago. Diseñadas para abordar las limitaciones de los enfoques tradicionales, estas metodologías priorizan la escalabilidad, la precisión y la eficiencia, al tiempo que garantizan interpretabilidad y conocimientos prácticos para los profesionales de la ecología. Al centrarse en la escalabilidad, las herramientas apuntan a apoyar diversos contextos ecológicos, desde parques urbanos hasta áreas de conservación forestal, contribuyendo a la preservación de la biodiversidad y la salud de los ecosistemas.

Para abordar los desafíos de la detección del muérdago, se han recopilado conjuntos de datos completos de diversos sitios de estudio, incluido el área de conservación ecológica *San Bartolo Ameyalco*, el *Jardín Ramón López Velarde* y el muelle de Xochimilco, en la Ciudad de México. Estas ubicaciones fueron elegidas en conjunto con SEDEMA debido a su alta prevalencia de infestaciones de muérdago y sus diversas condiciones ambientales, que brindan una base sólida para desarrollar modelos de detección sólidos.

Los conjuntos de datos se construyeron utilizando un dron multiespectral P4, que proporciona información espacial y espectral detallada, fundamental para distinguir el muérdago de la vegetación huésped. Este dispositivo permitió capturar imágenes de alta resolución, preservando los detalles de grano fino de los racimos de muérdago, como sus hojas anchas y aterciopeladas y su coloración distintiva. Estas características son esenciales para una identificación precisa, particularmente en las primeras etapas de infestación, donde el muérdago se mezcla perfectamente con sus árboles huéspedes. La imagen aérea original a la izquierda de la Figura 2 resalta esta similitud visual, lo que demuestra el desafío de diferenciar el muérdago de la vegetación saludable. Para abordar esto, se aplicaron pasos de preprocesamiento, como el co-registro y la normalización de bandas multiespectrales, para garantizar la coherencia en todo el conjunto de datos.

La identificación y segmentación de los grupos de muérdago en cada conjunto de imágenes fue realizada manualmente por un especialista en ecología. Estos mapas etiquetados sirvieron como base para el aprendizaje supervisado, proporcionando datos precisos y reales para entrenar modelos de aprendizaje automático. La imagen izquierda de la Figura 2

Figura 1. Representación visual de grupos de muérdago (*S. interruptus*) resaltados en azul, capturados en el Jardín Ramón López Velarde, Ciudad de México, en noviembre de 2024. El muérdago hemiparásito se mezcla perfectamente con la vegetación huésped, lo que complica la detección manual y enfatiza la necesidad de métodos de identificación automatizados.

muestra un ejemplo de un mapa de segmentación (en rojo) sobrepuesto a la imagen original correspondiente. Esta combinación de imágenes de alta resolución y segmentación experta creó un conjunto de datos sólido para avanzar en las metodologías de detección automatizadas.

La integración de la IA y la teledetección es prometedora no sólo para abordar el desafío inmediato de las infestaciones de muérdago, sino también para aplicaciones ecológicas más amplias. Estas herramientas permiten a los conservacionistas monitorear la biodiversidad, evaluar la salud de los ecosistemas y desarrollar estrategias de manejo específicas para especies invasoras. Como se destaca en esta investigación, el uso de estas tecnologías requiere un enfoque multidisciplinario que combine experiencia ecológica con avances computacionales, asegurando que las soluciones sean científicamente sólidas y prácticas.

Los objetivos de este trabajo son triples. Primero, reflexionar sobre los conocimientos adquiridos a partir de los esfuerzos de colaboración continuos de nuestro equipo de investigación interdisciplinario, enfatizando tanto los éxitos como los desafíos encontrados en la aplicación de la IA y la teledetección a la conservación ecológica. En segundo lugar, fomentar el diálogo dentro de las comunidades científica y conservacionista, fomentando la colaboración interdisciplinaria e inspirando la innovación. Finalmente, delinear posibles direcciones para investigaciones futuras, enfatizando la importancia del despliegue sostenible y equitativo de tecnologías de IA para abordar los desafíos ambientales apremiantes.

Avances actuales

La fase inicial de nuestra investigación exploró el potencial de los enfoques de computación evolutiva, específicamente la Programación Genética (GP, de *Genetic Programming*), para abordar el desafío de la detección del muérdago. Se empleó GP para desarrollar un método de índice espectral capaz de distinguir infestaciones de muérdago en imágenes multiespectrales adquiridas por UAV [6]. Al simular procesos evolutivos naturales, GP optimizó iterativamente combinaciones de bandas espectrales para crear un índice altamente efectivo para detectar especies de *Phoradendron Velutinum* durante

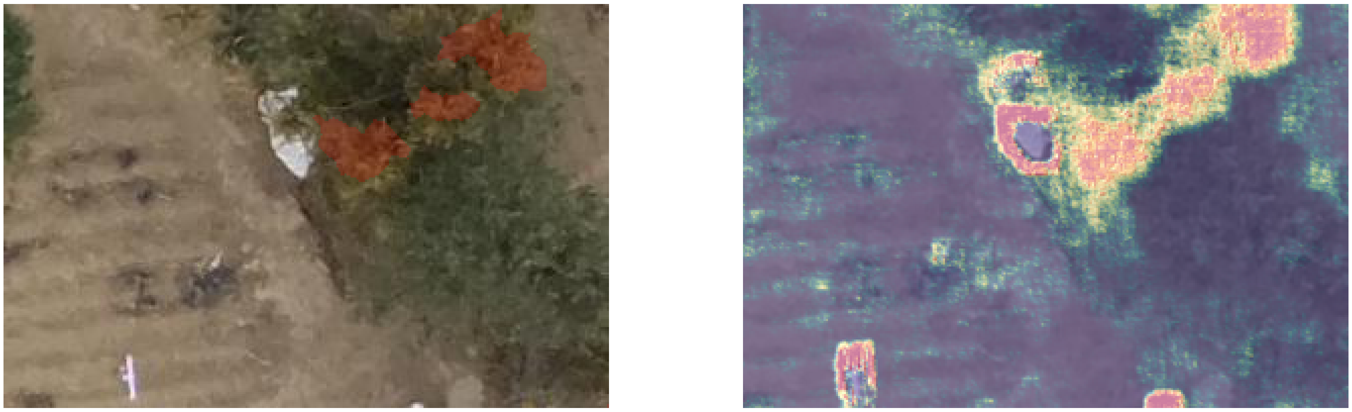

Figura 2. Visualización del procesamiento del modelo CNN ResNet-34 para la identificación de especies de muérdago conocidas como *P. Velutinum*. A la izquierda, la imagen RGB de alta resolución original con una máscara de segmentación binaria superpuesta (roja) que resalta las regiones de muérdago (etiquetadas por expertos) y a la derecha, el mapa de calor generado por CNN que enfatiza las áreas identificadas como muérdago.

su etapa de floración.

Esta metodología demostró una precisión notable, superior al 96 % en condiciones controladas, y proporcionó una prueba de concepto fundamental para la detección automatizada. El enfoque en la etapa de floración de *P. velutinum* aprovechó las características espectrales únicas de la planta parásita, permitiendo una diferenciación precisa de los árboles huéspedes y la vegetación circundante. Este logro puso de manifiesto el potencial de los enfoques computacionales para complementar los métodos ecológicos tradicionales, particularmente al abordar tareas que requieren mucho desarrollo manual, como las inspecciones manuales de campo.

El éxito de GP no solo validó la viabilidad de utilizar análisis espectral automatizado para la detección de muérdago, sino que también destacó la complejidad de ampliar dichos métodos a contextos ecológicos más amplios y al seguimiento de diferentes especies de muérdago. Estos conocimientos sientan las bases para exploraciones posteriores hacia técnicas más adaptables capaces de manejar la variabilidad ambiental y abordar una gama más amplia de especies. Por lo tanto, nuestra investigación se amplió para explorar el uso de descriptores de textura y técnicas de aprendizaje automático para la detección de muérdago.

Reconociendo la importancia de los patrones espaciales para diferenciar el muérdago de los árboles hospedantes, empleamos descriptores como matrices de coocurrencia de nivel de grises (GLCM) y filtros Gabor, junto con transformaciones de color e índices de vegetación, para capturar las características texturales y espectrales únicas de los grupos de muérdago [8]. Estas características fueron seleccionadas con el objetivo de mejorar el rendimiento de clasificación de *S. Interruptus* realizada mediante Máquinas de Vectores de Soporte (SVM, de *Support Vector Machines*), que alcanzó una precisión cercana al 60 % en entornos controlados [8]. Estos resultados demostraron el potencial de combinar descriptores espectrales y de textura con modelos de aprendizaje automático para la aplicación actual. Sin embargo, el rendimiento de estos métodos se vio significativamente afectado por la gran similitud entre el follaje del muérdago y del árbol huésped. Además, la falta de otras características espectrales limitó la adaptabilidad de nuestro enfoque a entornos más complejos y heterogéneos, como grandes áreas boscosas con varias especies de muérdago.

Los resultados de SVM revelaron el potencial del uso de técnicas computacionales para clasificar las infestaciones de muérdago, pero también evidenciaron sus limitaciones para adaptarse a la variabilidad ambiental y capturar la complejidad de los patrones de textura. Estos desafíos resaltaron la necesidad de metodologías capaces de aprender características complejas directamente a partir de los datos, sin depender de descriptores artesanales predefinidos. Esta comprensión nos llevó a explorar un enfoque de aprendizaje profundo, que ofrecía un marco más dinámico y adaptable para manejar tareas ecológicas complejas.

Entre las técnicas de aprendizaje profundo, las redes neuronales convolucionales (CNN, de *Convolutional Neural Networks*) surgieron como una opción natural y muy adecuada. Las CNN están diseñadas específicamente para extraer jerarquías espaciales de características de imágenes, lo que las hace particularmente efectivas para tareas ricas en texturas, como la detección de muérdago. Al aprender patrones directamente de las imágenes adquiridas por los UAV, las CNN proporcionaron la flexibilidad necesaria para tener en cuenta la variabilidad de las condiciones ambientales, como los cambios en la iluminación, la densidad del follaje y la estructura de los árboles. Esta adaptabilidad convirtió a las CNN en un siguiente paso ideal en nuestra investigación, cerrando la brecha entre los esfuerzos anteriores y la necesidad de soluciones sólidas y escalables.

Con el propósito de probar la viabilidad de técnicas avanzadas de clasificación basadas en imágenes, comenzamos por el desarrollo de una aplicación más sencilla orientada a la detección de árboles muertos en parques urbanos [9]. Esta tarea nos permitió centrarnos en características visuales bien definidas, como la ausencia de follaje o áreas verdes, mientras utilizamos datos adquiridos por UAV para establecer modelos de referencia. Para hacerlo, implementamos ResNet-34, una arquitectura de CNN conocida por su equilibrio entre eficiencia computacional y poder de representación [10, 11, 12]. ResNet-34 utiliza bloques residuales con conexiones de acceso directo, que mitigan el problema del desvanecimiento del gradiente y permiten un entrenamiento eficaz en redes profundas [13]. Esta arquitectura sobresale en la extracción de características de grano fino, lo que la hace particularmente adecuada para tareas de monitoreo ecológico donde las distinciones sutiles, como la identificación de especies invasoras o vegetación degradada, son críticas.

Esta fase de nuestra investigación no sólo confirmó el potencial de las CNN para el monitoreo ecológico y de la vegetación, sino que también sentó las bases para esfuerzos posteriores para capturar patrones de textura más sutiles asociados con las infestaciones de muérdago. Al aplicar con éxito las CNN a la detección de árboles muertos, pudimos perfeccionar nuestras metodologías y generar confianza en la escalabilidad y adaptabilidad de los modelos de aprendizaje profundo para desafíos ecológicos más amplios, como la clasificación de especies de muérdago.

Luego, nuestra investigación avanzó para implementar ResNet-34 para la tarea más compleja de clasificación del muérdago en espacios verdes urbanos. Para profundizar nuestra comprensión de los rasgos visuales y las características asociadas con dicha planta parásita, realizamos experimentos que perturbaron las imágenes de entrada con cambios en el brillo, el ruido y la resolución. Estas perturbaciones fueron diseñadas para explorar cómo las variaciones en las condiciones ambientales afectaban la capacidad del modelo para identificar grupos de muérdago. Este enfoque proporcionó información valiosa sobre las características visuales específicas que ResNet-34 priorizó durante la clasificación, como el color, la textura y los patrones espaciales, mejorando nuestra comprensión ecológica del muérdago y sus distintas firmas visuales.

Además, integramos herramientas de explicabilidad, como mapas de activación de clases (también conocidos como mapas de calor), para validar aún más las predicciones del modelo. Estas visualizaciones resaltaron las regiones de las imágenes en las que se basó el modelo para sus clasificaciones, lo que nos permitió asegurar que sus decisiones se alinearan con las expectativas ecológicas. Esta combinación de análisis de perturbaciones y explicabilidad no solo fortalece la interpretabilidad del modelo, sino que también pretende proporcionar a los usuarios finales, como expertos forestales y profesionales de la conservación, conocimientos prácticos sobre las características visuales del muérdago, facilitando una gestión ecológica más informada.

La implementación de modelos de CNN mejoró significativamente nuestra capacidad para identificar grupos de muérdago; sin embargo, estos modelos también introdujeron demandas computacionales sustanciales. Entrenar e implementar modelos de CNN a escala requiere una potencia de procesamiento considerable, particularmente cuando se manejan imágenes de vehículos aéreos no tripulados de alta resolución. Este desafío se abordó a través de nuestra colaboración con el Espacio de Innovación UNAM-Huawei, que brindó acceso a infraestructura computacional de alto rendimiento, permitiendo la optimización y entrenamiento eficiente de nuestros modelos. Sin embargo, ciertos experimentos requirieron una evaluación de diversas configuraciones de hardware y entornos computacionales para evaluar el rendimiento y la adaptabilidad en diferentes escenarios. Por ejemplo, la combinación óptima de configuraciones de hardware, estrategias de asignación de recursos y configuraciones de sistemas para minimizar los tiempos de entrenamiento aún no se comprende adecuadamente, lo que destaca la necesidad de explorar más a fondo las compensaciones entre hardware y rendimiento para aplicaciones ecológicas de aprendizaje profundo.

Para abordar la exploración anterior, realizamos un análisis comparativo de plataformas de hardware, evaluando configuraciones que van desde computadoras portátiles de consumo hasta servidores de alto rendimiento. El estudio reveló conocimientos críticos sobre las compensaciones entre eficiencia computacional, costo y rendimiento del modelo. Las estaciones de trabajo de alto rendimiento y el hardware especializado proporcionaron reducciones significativas en el tiempo de entrenamiento y mejoraron la escalabilidad de nuestros modelos. Al optimizar el rendimiento del hardware, no solo aumentamos la viabilidad de implementar modelos de aprendizaje profundo para la detección de muérdago, sino que también sentamos las bases para aplicaciones más amplias en la conservación ecológica. Esta fase de investigación reforzó la necesidad de colaboraciones interdisciplinarias e iniciativas de intercambio de recursos para superar las barreras computacionales, allanando el camino para soluciones más sólidas y escalables a desafíos ambientales complejos.

Un objetivo principal de nuestra investigación ha sido desarrollar herramientas prácticas y escalables para la conservación ecológica, pasando de metodologías experimentales a aplicaciones del mundo real. Nuestros sistemas de detección, que van desde mapas de infestación hasta modelos de clasificación automatizados, están diseñados teniendo en cuenta la usabilidad y abordando las necesidades del personal forestal, los administradores de parques urbanos y los formuladores de políticas. Estas herramientas tienen el potencial de automatizar tareas que requieren mucho desarrollo manual, proporcionar información útil para la asignación de recursos y monitorear la efectividad de las estrategias de gestión. Además, las herramientas de explicabilidad, como los mapas de activación de clases, garantizan la interpretabilidad y la confianza en los resultados de los modelos, cerrando la brecha entre los métodos computacionales avanzados y los esfuerzos prácticos de conservación.

Discusión

Los avances logrados hasta ahora en esta iniciativa representan un progreso significativo para abordar los desafíos ecológicos que plantean las infestaciones de muérdago. Sin embargo, como ocurre con cualquier enfoque innovador, el viaje ha revelado no sólo éxitos sino también limitaciones que justifican una reflexión crítica. Más allá de los logros técnicos, este trabajo plantea preguntas más amplias sobre el papel de la IA en la conservación ecológica, las implicaciones éticas y prácticas del despliegue de tales tecnologías y los desafíos de garantizar la escalabilidad y la accesibilidad. Esta sección tiene como objetivo proporcionar un análisis de estos aspectos, explorando las lecciones aprendidas, los desafíos no resueltos y las oportunidades para la innovación futura. Al reflexionar sobre estas dimensiones, esperamos inspirar un diálogo continuo y una colaboración interdisciplinaria en la búsqueda de soluciones de conservación sostenibles e impactantes.

Las metodologías exploradas reflejan una respuesta en evolución a los complejos desafíos que plantean las infestaciones de muérdago. GP demostró un potencial significativo al automatizar el diseño de un índice espectral que logró una precisión general del 96,6 % en la detección de infestaciones de muérdago durante la etapa de floración de *P. velutinum.* Esta precisión se logró en condiciones controladas con imágenes multiespectrales de vehículos aéreos no tripulados, lo que destaca la capacidad de GP para derivar soluciones personalizadas para desafíos ecológicos específicos. Sin embargo, su dependencia de las características espectrales de la etapa de floración resalta las limitaciones a la hora de generalizar entre condiciones variables y etapas fenológicas.

Este éxito se yuxtapuso a desafíos en materia de escalabilidad y adaptabilidad. Si bien GP destacó en escenarios específicos, su dependencia de índices espectrales limitó su solidez en entornos heterogéneos. El análisis detallado de las soluciones GP reveló una fuerte dependencia de las bandas del espectro visible, particularmente los canales R y B, alineándose con la pigmentación característica del muérdago. Esta especificidad, si bien es eficaz en el área de estudio objetivo, plantea dudas sobre la adaptabilidad de los algoritmos basados en GP a otros entornos y especies de muérdago.

La exploración posterior de descriptores de textura y clasificadores de aprendizaje automático representó un paso importante para abordar la complejidad espacial de las infestaciones de muérdago. Se emplearon técnicas como matrices de coocurrencia de nivel de grises (GLCM, de *Gray Level Co-occurrence Matrix*), filtros de Gabor y patrones binarios locales (LBP, de *Local Binary Patterns*) para capturar las distintas características de textura de las regiones de muérdago, complementadas con índices de vegetación derivados de imágenes multiespectrales adquiridas por UAV. Estos enfoques mostraron el papel fundamental de la textura para diferenciar el muérdago de la vegetación sana, particularmente en escenarios controlados. Sin embargo, los experimentos de clasificación realizados con SVM revelaron desafíos notables. Estos incluían la sensibilidad a la variabilidad ambiental, como la iluminación y los cambios estacionales, así como las limitaciones inherentes de las características artesanales, que restringían la escalabilidad y la generalización a ecosistemas heterogéneos. Si bien la metodología ofreció información valiosa, los resultados resaltaron la necesidad de técnicas más adaptables capaces de aprender características directamente de los datos, sentando las bases para enfoques de aprendizaje profundo.

Las técnicas de aprendizaje profundo, si bien ofrecen una mayor adaptabilidad, requieren grandes conjuntos de datos anotados para explotar plenamente su potencial. Sin embargo, nuestro conjunto de datos era de tamaño limitado, lo que planteaba limitaciones importantes para entrenar las CNN para realizar tareas como la segmentación semántica, que exigen anotaciones a nivel de píxel. Para abordar esta limitación, adoptamos un enfoque de clasificación de imágenes dividiendo imágenes aéreas más grandes en mosaicos más pequeños para su análisis. Este método maximizó la utilidad de los datos disponibles y aprovechó las sólidas capacidades de extracción de características de las arquitecturas CNN como ResNet-34. En lugar de abordar de inmediato la compleja tarea de distinguir el muérdago de la vegetación verde, optamos por un enfoque más gradual. Al centrarnos inicialmente en la clasificación de árboles muertos versus vegetación verde, refinamos nuestras metodologías, construimos una base sólida para el entrenamiento de modelos y validamos la viabilidad de nuestro enfoque. Esta progresión gradual equilibró los desafíos técnicos de la implementación del aprendizaje profundo con las limitaciones prácticas de la disponibilidad limitada de datos y la complejidad ecológica.

Nuestra investigación logró un hito importante con la clasificación de árboles muertos utilizando imágenes multiespectrales adquiridas por UAV. El estudio comparativo de ResNet-34 y DenseNet-121 proporcionó información sobre la aplicabilidad de estas arquitecturas para detectar anomalías de la vegetación indicativas de perturbaciones ecológicas. Ambos modelos demostraron un rendimiento sólido, con DenseNet-121 superando ligeramente a ResNet-34 en métricas como precisión y puntuación F1, alcanzando valores de aproximadamente el 97 % [9]. Este alto nivel de precisión se complementó con la integración de mapas de calor, que iluminaron las áreas dentro de las imágenes de entrada más relevantes para el proceso de clasificación. Estas herramientas visuales mejoraron la interpretabilidad de las decisiones de los modelos, cerrando la brecha entre los resultados computacionales y los conocimientos ecológicos procesables. Si bien estos resultados validaron la viabilidad de utilizar CNN para monitorear la salud de los árboles, también sentaron un precedente para abordar tareas de clasificación más complejas, como distinguir las infestaciones de muérdago de la vegetación sana.

Aprovechando el éxito de la clasificación de los árboles muertos, pasamos a la tarea más compleja de identificar las infestaciones de muérdago. Confiados en nuestro enfoque, extrajimos los mosaicos de la misma manera que antes, esta vez centrándonos específicamente en el muérdago. Estos mosaicos, derivados de imágenes multiespectrales adquiridas por vehículos aéreos no tripulados, formaron la base para entrenar el modelo para distinguir el muérdago (*S. interruptus*) de la vegetación circundante. La Figura 3 muestra mosaicos representativos de muérdago, capturando las distintas características visuales necesarias para la clasificación.

El modelo ResNet-34 logró una precisión notable, con métricas de rendimiento de referencia que superaron constantemente el 83 % en las estrategias de validación de reserva y k-fold. Las visualizaciones de mapas de calor resaltaron de manera efectiva características biológicamente relevantes, alineándose con las regiones de muérdago anotadas y validando la relevancia ecológica de las clasificaciones del modelo.

Sin embargo, el estudio también destacó limitaciones específicas. El modelo mostró sensibilidad a las perturbaciones ambientales, como cambios en los canales de color y la introducción de ruido, particularmente ruido de sal y pimienta a intensidades más altas. Si bien esta sensibilidad destaca áreas que requieren un mayor refinamiento, como mejorar la robustez a través de metodologías de entrenamiento tolerantes al ruido o estrategias de preprocesamiento adaptativo, también ofrece una oportunidad única para comprender mejor las características visuales del muérdago. Al analizar las respuestas del modelo en diferentes condiciones, obtenemos conocimientos más profundos sobre las características definitorias del muérdago, como su textura, color y patrones espaciales. Esta mejor comprensión puede contribuir tanto a la identificación manual por parte de los profesionales ecológicos como al desarrollo de métodos de detección automatizados más sofisticados. Además, la dependencia de la clasificación basada en mosaicos, si bien es práctica dadas las limitaciones del conjunto de datos, presenta desafíos a la hora de capturar el contexto espacial más amplio necesario para el monitoreo ecológico a gran escala.

De cara al futuro, estos resultados resaltan el potencial de integrar técnicas de segmentación semántica para superar las limitaciones espaciales de la clasificación basada en mosaicos. Al mejorar la capacidad del modelo para identificar el muérdago a nivel de píxel, los esfuerzos futuros podrían mejorar significativamente la precisión y la escalabilidad de la detección. Sin embargo, las demandas computacionales de metodologías tan refinadas resaltan la importancia de soluciones de hardware eficientes. A medida que los modelos de aprendizaje profundo crecen en complejidad, abordar las compensaciones entre la eficiencia computacional y el rendimiento se vuelve esencial para garantizar su viabilidad para aplicaciones ecológicas a gran escala.

Nuestros resultados proporcionan información valiosa para optimizar las demandas computacionales de los modelos de aprendizaje profundo para la detección de muérdago. Se evaluaron múltiples configuraciones, incluidos procesadores Intel y AMD, GPU de consumo y estaciones de trabajo de alto rendimiento, para evaluar su idoneidad para el entrenamiento de ResNet-34. En particular, los procesadores AMD superaron consistentemente a sus homólogos Intel en eficiencia de entrenamiento, y las configuraciones AMD de rango medio lograron un equilibrio entre costo y rendimiento. Además, la desactivación del hyperthreading (HT) en las CPU aceleró los tiempos de entrenamiento en la mayoría de las plataformas, lo que resalta la importancia de ajustar cuidadosamente la configuración del hardware para las tareas de aprendizaje profundo.

Las GPU de consumo demostraron eficacia para tareas de entrenamiento a corto plazo, brindando un rendimiento competitivo en relación con estaciones de trabajo de alto rendimiento en escenarios con demandas computacionales moderadas. Sin embargo, los servidores de alto rendimiento exhibieron una escalabilidad incomparable y tiempos de entrenamiento reducidos, lo que los hizo indispensables para conjuntos de datos extensos o cargas de trabajo a largo plazo. Estos hallazgos resaltan las compensaciones entre la eficiencia computacional y el costo, enfatizando la necesidad de seleccionar configuraciones de hardware que se alineen con la escala y el alcance de las aplicaciones ecológicas.

Estos resultados refuerzan la viabilidad de implementar modelos avanzados para el monitoreo ecológico, siempre que las opciones de hardware estén optimizadas para restricciones operativas específicas. Naturalmente, pasan a consideraciones más amplias de implementación, como abordar los desafíos de accesibilidad en regiones de escasos recursos y equilibrar el rendimiento computacional con la sostenibilidad ambiental.

Si bien nuestras metodologías han demostrado un potencial significativo, no están exentas de limitaciones, muchas de las cuales presentan desafíos críticos para una implementación y escalabilidad más amplias. Una limitación clave reside en la dependencia de imágenes de alta resolución adquiridas por vehículos aéreos no tripulados, que, si bien permiten una detección precisa, exigen recursos sustanciales para la recopilación y el procesamiento de datos. Esta dependencia plantea dudas sobre la viabilidad de implementar tales métodos en regiones con recursos limitados o en ecosistemas de gran escala donde la cobertura de los UAV y la capacidad computacional pueden estar restringidas.

Otra limitación es el desafío de generalizar nuestros modelos a diversos contextos ecológicos. La alta precisión lograda en entornos controla-

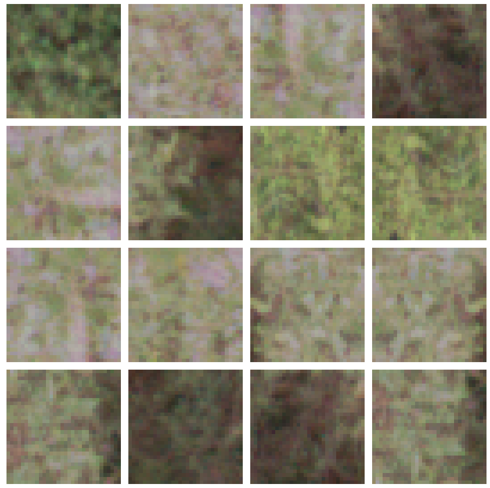

Figura 3. Mosaicos de muestra utilizados en el conjunto de datos de entrenamiento para muérdago (*S. interruptus*). Cada mosaico mide 32×32 píxeles y corresponde a un segmento extraído de las imágenes aéreas originales. Estos mosaicos se utilizaron para entrenar y validar el modelo ResNet-34.

dos o semicontrolados no garantiza un rendimiento equivalente en paisajes más complejos y heterogéneos. Las variaciones en los tipos de vegetación, las especies de muérdago y las condiciones ambientales, como la iluminación o los cambios estacionales, pueden afectar significativamente la solidez del modelo. Esto destaca la necesidad de una extensa validación de campo y el desarrollo de modelos capaces de adaptarse a la naturaleza dinámica de los sistemas ecológicos.

Ampliar estos enfoques a áreas más grandes también presenta desafíos no resueltos, incluidas las demandas computacionales del procesamiento de conjuntos de datos expansivos y la necesidad de automatización en la planificación de vuelos de UAV y la anotación de datos. Además, los sesgos en los datos de entrenamiento (derivados de una sobrerrepresentación de condiciones ecológicas o áreas geográficas específicas) pueden limitar inadvertidamente la aplicabilidad de los modelos. Abordar estos sesgos requerirá una cuidadosa conservación de los conjuntos de datos, centrándose en mejorar la representatividad y la diversidad de las muestras de entrenamiento.

Estas brechas invitan a una mayor exploración e innovación dentro de las comunidades científica y conservacionista. ¿Cómo podemos equilibrar la necesidad de datos de alta resolución con el objetivo de crear soluciones escalables y accesibles? ¿Existen enfoques híbridos, como la integración de imágenes satelitales para una cobertura más amplia con datos de UAV para realizar ajustes, que puedan cerrar esta brecha? Además, las técnicas de aprendizaje adaptativo, en las que los modelos se actualizan continuamente con nuevos datos de diversos entornos, podrían ofrecer un camino hacia una mayor generalización y resiliencia.

Al reconocer estas limitaciones, pretendemos estimular el debate y la colaboración sobre cómo abordarlas. Si bien los desafíos son importantes, también presentan oportunidades para perfeccionar y ampliar la aplicabilidad de las herramientas de conservación impulsadas por la IA, allanando el camino para soluciones ecológicas más inclusivas e impactantes.

Las metodologías y los conocimientos desarrollados en esta investigación tienen un potencial significativo para hacer avanzar las aplicaciones de IA más allá de la detección de infestaciones de muérdago. La integración de imágenes adquiridas por UAV, aprendizaje automático y técnicas de aprendizaje profundo establece un marco escalable que puede adaptarse a una variedad de desafíos ecológicos. Por ejemplo, estos enfoques podrían utilizarse para monitorear la biodiversidad en hábitats amenazados, evaluar la propagación de otras especies invasoras o evaluar la salud de los bosques afectados por el cambio climático. La capacidad de los modelos de IA para procesar y analizar grandes volúmenes de imágenes de alta resolución los posiciona como herramientas transformadoras en la conservación ecológica.

Sin embargo, la aplicación más amplia de estas tecnologías plantea importantes cuestiones éticas y prácticas. Las preocupaciones sobre la propiedad de los datos y la privacidad son particularmente relevantes cuando los UAV se despliegan en áreas con actividad humana, capturando potencialmente información no deseada. Establecer protocolos claros para la gobernanza de datos y garantizar el cumplimiento de las normas de privacidad será esencial para mantener la confianza pública y la integridad ética. De manera similar, la accesibilidad a herramientas avanzadas de IA en áreas de escasos recursos sigue siendo un desafío apremiante. La informática de alto rendimiento y la infraestructura de UAV a menudo se limitan a instituciones bien financiadas, lo que genera disparidades en cuanto a quién puede beneficiarse de estas innovaciones. Abordar estas desigualdades requerirá esfuerzos de colaboración para diseñar soluciones rentables y compartir recursos de manera más equitativa.

La sostenibilidad ambiental es otra consideración crítica. Los vuelos de UAV, en particular los que implican misiones múltiples, contribuyen a las emisiones de carbono, mientras que la naturaleza de uso intensivo de energía de la informática de alto rendimiento aumenta la huella ambiental de las aplicaciones de IA. Estos factores resaltan la importancia de optimizar el uso de recursos y explorar alternativas más ecológicas, como la informática de punta o hardware más eficiente desde el punto de vista energético, para alinear el progreso tecnológico con los principios de conservación.

De cara al futuro, esta investigación invita a una reflexión crítica sobre la sostenibilidad y la inclusión de la IA en aplicaciones ecológicas. ¿Cómo podemos garantizar que estas herramientas sean accesibles para las regiones de escasos recursos y adaptables a diversos contextos ecológicos? ¿Existen formas de integrar el conocimiento ecológico tradicional con los sistemas de IA para crear estrategias de conservación más holísticas e inclusivas? Además, ¿qué marcos se necesitan para minimizar el impacto ambiental del despliegue de estas tecnologías a escala?

Al plantear estas preguntas, pretendemos provocar un debate reflexivo dentro de las comunidades científica y conservacionista. Si bien la IA es inmensamente prometedora para abordar desafíos ecológicos complejos, su aplicación responsable requiere un equilibrio entre innovación y sostenibilidad. A través de la colaboración interdisciplinaria y la previsión ética, estas herramientas pueden evolucionar para volverse más inclusivas, equitativas y alineadas con los objetivos de conservación a largo plazo.

Los avances logrados en esta investigación sientan una base sólida para explorar nuevas vías en el monitoreo y la conservación ecológicos. Ampliar el alcance de las pruebas más allá de la detección de muérdago a otros desafíos ecológicos ofrece oportunidades prometedoras para validar y perfeccionar aún más estas metodologías. Por ejemplo, las mismas imágenes adquiridas por UAV y los mismos marcos de aprendizaje automático podrían adaptarse para detectar otras especies invasoras, como los escarabajos de la corteza o las enredaderas parásitas, que amenazan de manera similar la salud de los bosques. Además, las técnicas de análisis espectral y de textura empleadas aquí podrían aplicarse para evaluar indicadores más amplios de salud forestal, como la identificación de signos tempranos de enfermedades o el seguimiento de la recuperación de la vegetación después de perturbaciones.

Las tecnologías emergentes presentan posibilidades interesantes para mejorar la escalabilidad y precisión de estos esfuerzos. *Vision Transformers* (VT), una arquitectura de última generación para el análisis de imágenes, podría emplearse para capturar relaciones espaciales complejas y mejorar la precisión de la detección en entornos altamente heterogéneos. Al aprovechar su capacidad para procesar el contexto global dentro de imágenes, VT puede proporcionar una comprensión más completa de los patrones ecológicos.

Un siguiente paso fundamental es la puesta a prueba de estas metodologías en implementaciones del mundo real. Las colaboraciones con profesionales de la conservación locales y agencias forestales serán esenciales para adaptar estas herramientas a las necesidades y limitaciones específicas de las aplicaciones de campo. Los programas piloto podrían servir como campo de pruebas para perfeccionar los flujos de trabajo, abordar los desafíos de usabilidad y generar confianza entre los usuarios finales. Además, las asociaciones con organizaciones comunitarias podrían facilitar la integración del conocimiento ecológico local, enriqueciendo los datos y aumentando la relevancia de los conocimientos impulsados por la IA.

La colaboración interdisciplinaria fue la piedra angular de este proyecto, impulsando tanto sus éxitos como su evolución para abordar el desafío ecológico de las infestaciones de muérdago. Al integrar la experiencia de CentroGeo, ITM e ITPA, así como aprovechar las asociaciones con SECTEI, SEDEMA, CORENA y el Espacio de Innovación UNAM-Huawei, el equipo unió con éxito la innovación técnica con aplicaciones ecológicas. Este marco de colaboración facilitó el desarrollo de metodologías avanzadas, desde el diseño de índices espectrales hasta la implementación de modelos de aprendizaje profundo, asegurando que cada fase del proyecto estuviera informada por una amplia gama de perspectivas y habilidades. De manera similar, el acceso a recursos informáticos de alto rendimiento a través de asociaciones institucionales permitió al equipo superar importantes desafíos computacionales, acelerando el refinamiento de los modelos de aprendizaje profundo y permitiendo evaluaciones detalladas del desempeño.

Sin embargo, esta colaboración también reveló desafíos inherentes que destacan la complejidad de la investigación interdisciplinaria. Las diferencias en las prioridades y metodologías disciplinarias en ocasiones llevaron a desajustes en los objetivos y expectativas del proyecto. Por ejemplo, mientras los socios técnicos priorizaban la eficiencia y precisión computacional, los colaboradores ecológicos a menudo enfatizaban la interpretabilidad y aplicabilidad de campo de las herramientas que se estaban desarrollando. La comunicación efectiva y los ciclos de retroalimentación iterativos son esenciales para resolver estas tensiones, pero también resaltaron la necesidad de marcos más claros para alinear los objetivos multidisciplinarios desde el principio.

Al reflexionar sobre estas experiencias, surgen varias lecciones que podrían informar futuros proyectos interdisciplinarios. En primer lugar, establecer objetivos claros y compartidos al comienzo de un proyecto puede ayudar a alinear las prioridades y reducir la fricción entre disciplinas. En segundo lugar, la creación de canales de comunicación estructurados (como talleres periódicos o plataformas colaborativas) puede facilitar el diálogo continuo y garantizar que se escuchen todas las voces durante todo el proceso de investigación. Finalmente, fomentar una cultura de respeto mutuo y apertura a perspectivas diversas es fundamental para maximizar el potencial de los equipos interdisciplinarios.

Una prioridad clave para futuras investigaciones e implementación es fomentar un diálogo que no solo perfeccione las aplicaciones de IA para la detección de muérdago, sino que también genere conciencia sobre la amenaza ecológica que representa. Este diálogo debe extenderse más allá de los científicos y expertos técnicos para incluir a los formuladores de políticas, los profesionales de la conservación y el público en general. Si bien a menudo se romantiza el muérdago, es esencial comunicar su grave impacto en los bosques, donde las infestaciones no controladas debilitan los árboles, reducen la biodiversidad y alteran servicios ecosistémicos críticos. Involucrar a las comunidades en esta conversación puede ayudar a cambiar las percepciones y generar apoyo para las medidas de conservación necesarias.

Involucrar a las comunidades locales es particularmente importante, ya que sus conocimientos ecológicos tradicionales pueden mejorar los enfoques impulsados por la IA y fundamentarlos en realidades prácticas. Se necesitan marcos de colaboración para garantizar que las herramientas de detección y las estrategias de gestión sean accesibles y viables, especialmente en zonas de escasos recursos donde las infestaciones de muérdago podrían ser más graves. Además, las iniciativas de educación pública pueden ayudar a cerrar la brecha entre el simbolismo cultural del muérdago y la urgencia ecológica de su control, fomentando una comprensión compartida del problema y las soluciones requeridas.

A medida que avanzamos en estos esfuerzos, es crucial garantizar que las aplicaciones de IA sigan siendo equitativas, transparentes y alineadas con la gestión ecológica. El futuro del manejo del muérdago depende no sólo de la innovación tecnológica sino también de nuestra capacidad para involucrar y empoderar a todas las partes interesadas, desde los científicos conservacionistas hasta la comunidad en general. Al combinar herramientas avanzadas con una conciencia generalizada y una acción colaborativa, podemos proteger los bosques del impacto devastador de las infestaciones de muérdago y al mismo tiempo fomentar prácticas de conservación sostenibles.

Conclusiones

Este trabajo examina críticamente el papel de la IA y la teledetección para abordar los desafíos ecológicos que plantean las infestaciones de muérdago, ofreciendo información sobre las metodologías y enfoques que han dado forma a este campo emergente. Al aprovechar GP, descriptores de textura y CNN, nuestra investigación destaca tanto el potencial como las limitaciones de estas tecnologías en la conservación ecológica. GP demostró su eficacia en el desarrollo de índices espectrales adaptados a desafíos ecológicos específicos, proporcionando un punto de partida para la automatización en la detección de muérdago. De manera similar, los métodos de aprendizaje automático basados en texturas, si bien son valiosos para capturar patrones espaciales, evidenciaron las limitaciones de las características hechas a mano en entornos complejos y heterogéneos.

El cambio hacia el aprendizaje profundo, particularmente a través de las CNN, refleja una evolución en la metodología que no solo abordó las limitaciones anteriores sino que también reveló la escalabilidad y adaptabilidad de las soluciones para la conservación impulsadas por la IA. Estos avances sirven como una reflexión crítica sobre la aplicación de tecnologías de vanguardia en el monitoreo ecológico, ilustrando cómo la IA puede transformar tareas que requieren mucho desarrollo manual en procesos escalables y eficientes. Sin embargo, también plantean preguntas importantes sobre la practicidad y accesibilidad de estos enfoques, particularmente en entornos con recursos limitados.

Si bien las metodologías y los conocimientos presentados en este trabajo demuestran un potencial significativo para abordar los desafíos ecológicos, van acompañados de limitaciones críticas que justifican una mayor exploración. La variabilidad ambiental complica la generalización de estos modelos. El rendimiento de los enfoques de aprendizaje automático y aprendizaje profundo, optimizados para contextos ecológicos específicos, puede degradarse en entornos heterogéneos caracterizados por diversos tipos de vegetación, cambios estacionales o condiciones de luz fluctuantes. Abordar esta variabilidad requerirá metodologías adaptativas capaces de aprender de diversos conjuntos de datos, junto con una extensa validación de campo para garantizar la solidez en todos los entornos ecológicos.

Las demandas computacionales también presentan un desafío importante, particularmente para los modelos de aprendizaje profundo que requieren muchos recursos. El consumo de energía y la infraestructura necesaria para entrenar e implementar estos sistemas plantean cuestiones importantes sobre su sostenibilidad ambiental y práctica. Los esfuerzos de colaboración, como aprovechar la informática de punta para análisis en tiempo real u optimizar las configuraciones de hardware para lograr eficiencia, representan vías críticas para mitigar estas limitaciones.

La integración de la IA en la conservación ecológica exige una colaboración interdisciplinaria sostenida, basada tanto en la previsión ética como en estrategias de implementación prácticas. Las complejidades de la detección del muérdago y los desafíos ecológicos más amplios requieren experiencia de diversos campos, incluidas las ciencias computacionales, la ecología y las ciencias sociales, para garantizar que los avances tecnológicos sigan siendo contextualmente relevantes e inclusivos. Los esfuerzos futuros deben priorizar asociaciones que cierren la brecha entre la investigación de vanguardia y las prácticas de conservación viables, fomentando la colaboración entre investigadores, formuladores de políticas y comunidades locales.

El despliegue práctico de herramientas de IA también presenta oportunidades críticas para la innovación. La implementación en el mundo real, respaldada por pruebas de campo y un diseño centrado en el usuario, será clave para perfeccionar estas tecnologías y garantizar su usabilidad por parte de partes interesadas no especializadas, como administradores de parques y trabajadores de conservación. La integración de tecnologías emergentes como VT y la informática de punta podría mejorar aún más la adaptabilidad y eficiencia de las soluciones impulsadas por la IA, haciéndolas más accesibles y sostenibles.

De cara al futuro, el papel de la IA en la conservación está a punto de expandirse mucho más allá de la detección del muérdago. Estas herramientas ofrecen un inmenso potencial para abordar diversos desafíos, desde monitorear la biodiversidad hasta mitigar los impactos del cambio climático. Sin embargo, su impacto dependerá en última instancia de nuestra capacidad para alinear el progreso tecnológico con los principios de sostenibilidad, equidad y transparencia.

Al fomentar el diálogo interdisciplinario y adoptar un enfoque con visión de futuro, la comunidad científica puede promover aplicaciones de IA que no solo aborden las amenazas ecológicas inmediatas sino que también contribuyan al objetivo más amplio del desarrollo sostenible. El camino que tenemos por delante es desafiante y está lleno de oportunidades, y requiere acción colectiva, reflexión ética y un compromiso de colaboración que trascienda las fronteras disciplinarias. A través de estos esfuerzos, la IA puede convertirse en una fuerza transformadora en la conservación, impulsando la innovación y al mismo tiempo salvaguardando el planeta para las generaciones futuras.

Declaración de contribución de autoría de CRediT

Juan Carlos Valdiviezo-Navarro: Conceptualización, Metodología, Análisis formal, Redacción – borrador original, Redacción – revisión y edición. **Mauricio G. Orozco-del-Castillo:** Conceptualización, Metodología, Análisis formal, Redacción – borrador original, Redacción – revisión y edición.

Declaración de IA Generativa y tecnologías asistidas por IA en el proceso de redacción

Durante la preparación de este trabajo, los autores utilizaron ChatGPT para mejorar la legibilidad. Después de utilizar esta herramienta, los autores revisaron y editaron el contenido según fuera necesario y asumieron total responsabilidad por el contenido de la publicación.

Declaración de interés en competencia

Los autores declaran que no tienen intereses financieros en competencia ni relaciones personales conocidas que pudieran haber influido en el trabajo presentado en este artículo.

Reconocimiento

Expresamos nuestro agradecimiento a la Secretaría del Medio Ambiente (SEDEMA) por el apoyo durante el trabajo de campo. También reconocemos el esfuerzo de nuestro equipo para desarrollar este proyecto: Paola Andrea Mejía Zuluaga, León Felipe Dozal García, Alejandra Aurelia López Caloca, Alejandro Téllez Quiñones, Adriana Lizzette Luna Nieves, Amilcar Morales Gamas y Mario Ledesma Arriola (CentroGeo); Nivia Iracemi Escalante García, Eduardo Flores Gallegos y José Ernesto Olvera González (TecNM/ITPA); Nora L. Cuevas Cuevas, Carlos Bermejo-Sabbagh, William Orozco-González y Pedro A.G. Ortiz-Sánchez (TecNM/ITM). Parte de este trabajo ha sido financiado por la Secretaría de Educación, Ciencia, Tecnología e Innovación (SECTEI) de la Ciudad de México a través del Subsidio Número 3581c24.

Referencias

[1] P. K. Rai and J. Singh, "Invasive alien plant species: Their impact on environment, ecosystem services and human health," *Ecological Indicators*, vol. 111, p. 106020, 2020.

[2] B. Gallardo, M. Clavero, M. I. Sánchez, and M. Vilà, "Global ecological impacts of invasive species in aquatic ecosystems," *Global Change Biology*, vol. 22, 2016.

[3] M. Pompa-García and J. A. Rodríguez, "Variation of carbon uptake from forest species in mexico: a review," *Madera Y Bosques*, vol. 23, pp. 225–235, 2017.

[4] N. Mesa-Sierra, M. de la Peña-Domene, J. Campo, and C. Giardina, "Restoring mexican tropical dry forests: A national review," *Sustainability*, 2022.

[5] D. Cibrián Tovar, "Diagnóstico fitosanitario de áreas verdes infestadas por muérdago en la Ciudad de México," 2016.

[6] P. A. Mejia-Zuluaga, L. Dozal, and J. Valdiviezo-N., "Genetic programming approach for the detection of mistletoe based on uav multispectral imagery in the conservation area of mexico city," *Remote Sensing*, vol. 14, p. 801, 2022.

[7] T. Paap, T. Burgess, and M. Wingfield, "Urban trees: bridge-heads for forest pest invasions and sentinels for early detection," *Biological Invasions*, vol. 19, pp. 3515–3526, 2017.

[8] P. A. Mejia-Zuluaga, J. C. Valdiviezo-N., and L. Dozal, "Texture descriptors and machine learning algorithms for mistletoe detection in urban forests using multispectral imagery," in *Remote Sensing*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:264299698

[9] W. Orozco-González, J. C. Valdiviezo-Navarro, M. G. O. del Castillo, P. A. Mejia-Zuluaga, and L. Dozal, "Classification of dead trees in urban parks using aerial imagery and convolutional neural networks," in *ICCBR Workshops*, 2024. [Online]. Available: https://api.semanticscholar.org/CorpusID:270701684

[10] A. Abrol, A. Sharma, K. Karnic, and R. Ranjan, "Human activity recognition using resnet-34 model," *International Journal of Recent Technology and Engineering*, vol. 10, no. 5, pp. 224–228, 2021.

[11] X. Li, X. Xu, X. He, X. Wei, and H. Yang, "Intelligent crack detection method based on gm-resnet," *Sensors (Basel, Switzerland)*, vol. 23, 2023.

[12] W. Liu, L. Jiangtao, G. Zhao, L. Sun, H. Wang, W. Li, and B. Sun, "Improvement of cifar-10 image classification based on modified resnet-34," in *Advances in Natural Computation, Fuzzy Systems and Knowledge Discovery*, 2020.

[13] S. Salman and X. Liu, "Sparsity as the implicit gating mechanism for residual blocks," in *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2019, pp. 1–6.

Journal of Artificial Intelligence and Computing Applications (2024) 2(2): 30–38

Journal
of
Artificial Intelligence
and
Computing Applications

Exploration Paper

A Deep Learning Approach for Automated Identification of *Triatoma infestans* Using YOLOv8

Alejandro D. Espinosa-Chim[1,*], **Abigail Cetina-Romero**[1], **Hector Ivan Chumba-Poot**[1], **Manuel Armando Lara-Puc**[1], **Daniel Leobardo Vallejos-Aguilar**[1], **Miguel Angel Borges-Borges**[1], **and Brian A. Cumi-Guzman**[1]

[1]Tecnológico Nacional de México/IT de Mérida, Yucatán, Mexico

ABSTRACT

Triatoma infestans, a primary vector of Chagas disease, poses a significant public health risk in Latin America. Rapid and accurate identification of this insect is essential for both vector surveillance programs and individual-level decision-making after potential exposure. Traditional identification methods rely on manual inspection, which is time-consuming, error-prone, and dependent on expert knowledge. This study explores the feasibility of an AI-driven detection system based on the medium YOLOv8 model (YOLOv8m) to automate the identification of *T. infestans* from images. The model was trained on a dataset of 91 manually labeled images, with built-in data augmentation techniques dynamically generating 9,100 augmented images over 100 training epochs. The model achieved high accuracy, with a mean average precision at an Intersection over Union threshold of 50% (mAP@50) of 0.9588 and a fitness score of 0.6844, demonstrating its effectiveness under controlled conditions. To assess its reliability, detection examples were analyzed in varied lighting conditions and backgrounds, as well as in scenarios where *T. infestans* appeared alongside visually similar insects. Results show that the model can consistently detect *T. infestans* while avoiding false positives for other insect species, highlighting its potential for real-world deployment. This work provides a proof of concept for the integration of AI in entomological identification tasks. Future improvements include expanding the dataset, fine-tuning model hyperparameters, and adapting the system for mobile or embedded deployment to facilitate field usability. By automating *T. infestans* detection, this study contributes to enhanced vector surveillance efforts and better-informed responses to potential Chagas disease exposure.

Keywords: object detection, vector surveillance, artificial intelligence in entomology

1. Introduction

Triatoma infestans, commonly known as the *"Chinche Besucona"* in some parts of Mexico, is a primary vector of Chagas disease, a life-threatening illness caused by *Trypanosoma cruzi* [1]. This disease affects millions of people, particularly in Latin America, where *T. infestans* is widespread [2]. Early and accurate identification of this insect is crucial for vector control programs, as prompt intervention can significantly reduce transmission risk [3]. Recognizing *T. infestans* requires attention to its distinctive morphological features, including an elongated, flattened body, a dark exoskeleton with reddish-orange markings along the edges of its abdomen, and a prominent, cone-shaped head with an extended proboscis [4] (Figure 1). These characteristics are essential for distinguishing it from other hemipteran insects, some of which may resemble *T. infestans* but

E-mail address: le22080725@merida.tecnm.mx
https://doi.org/10.5281/zenodo.14976415

do not pose the same epidemiological risk.

Failure to correctly identify *T. infestans* can lead to serious public health consequences [4]. If a person is bitten by this insect, it is recommended to capture it rather than kill it and take it to a nearby health center for examination [2]. Experts can determine whether the insect carries *T. cruzi*, helping assess the risk of Chagas disease transmission [5]. Misidentification or failure to detect T. infestans could prevent this crucial step, delaying medical attention [1]. Conversely, if the insect is not T. infestans, the person should pursue a different course of action, making accurate identification critical for public health responses [3]. Despite its importance, manual identification of T. infestans remains a significant challenge [6]. The process is labor-intensive, requiring trained entomologists and public health workers to visually inspect and classify insects [7]. Moreover, this approach is time-consuming and prone to human error and inconsistencies, particularly when differentiating T. infestans from visually similar species [8]. The lack of scalable and automated detection methods limits the efficiency of epidemiological surveillance efforts, emphasizing the need for innovative technological solutions [9].

Recent advancements in artificial intelligence (AI), particularly in computer vision and deep learning, have revolutionized the automation of insect detection [10]. Object detection models, such as YOLO (You Only Look Once), have demonstrated remarkable success in real-time applications across various domains, including medical imaging, agriculture, and environmental monitoring [11]. These models offer a fast and accurate means of identifying objects within an image, making them ideal for tasks requiring rapid decision-making [12].

In the context of vector control, AI-powered detection systems can significantly enhance epidemiological surveillance by providing an efficient, scalable, and reliable solution for identifying *T. infestans* [13]. By automating the identification process, AI can bridge the gap between expert entomologists and the general public, enabling individuals to make informed decisions after an insect bite [14]. Public health authorities can deploy automated detection systems capable of quickly analyzing large volumes of images, thereby improving response times and aiding in more effective disease prevention strategies [15].

The primary objective of this study is to develop a detection model based on the medium variant of YOLOv8 (YOLOv8m) to accurately identify *T. infestans* in real-world conditions. The model was trained on a dataset consisting of images sourced from both online repositories [16, 17] and field photographs, ensuring a diverse and representative training set. Manual annotation was performed to precisely label the insects, improving the quality of the training data and reducing misclassification risks. By creating a robust detection system, this research aims to contribute to the automation of insect vector identification, ultimately assisting in the early detection and control of *T. infestans* populations to mitigate the spread of Chagas disease. Additionally, this model has the potential to allow non-experts to verify whether an insect requires further examination.

This study follows a structured methodology to ensure the development of an effective detection system. The YOLOv8m model was trained on a dataset of 91 manually labeled images, with additional variations introduced dynamically through YOLOv8's built-in data augmentation techniques to enhance model generalization. Over 100 training epochs, these transformations generated approximately 9,100 augmented images, significantly increasing dataset diversity without requiring additional manual labeling. The majority of the original images were sourced from public repositories, ensuring a diverse dataset while maintaining consistent visual characteristics. These included sharp focus, clear lighting, and varied angles, capturing *T. infestans* from top, side, and perspective views (Figure 1). The dataset also ensured that the insect was represented in natural resting positions, often against contrasting backgrounds, such as walls, textiles, and outdoor surfaces. Performance evaluation was conducted by measuring precision, recall, and mean average precision (mAP) to assess the model's effectiveness. Additionally, a real-time detection system was implemented using OpenCV, enabling live insect identification through video input. The study presents preliminary results, highlighting the strengths of the YOLOv8m model while also identifying areas for future improvement, such as fine-tuning hyperparameters, expanding the dataset with more diverse samples, and optimizing the model for better detection accuracy in challenging real-world conditions.

The remainder of this paper is structured as follows. Section 2 provides a detailed account of the AI prototype, including design considerations and innovative features. Section 3 describes the implementation process and presents preliminary results, highlighting unique challenges encountered during development. Section 4 explores the broader implications of the project, its contributions to advancing AI applications, and possible future enhancements. Finally, Section 5 summarizes the study's objectives, key findings, and the next steps for further research and development.

2. Project Description

The proposed system is designed to detect *T. infestans* in real-world conditions using YOLOv8m. The model is trained on a labeled dataset consisting of images from both online sources and field-collected photographs. By leveraging deep learning techniques, the system is optimized for real-time detection, ensuring fast and accurate identification of the insect in various environments. The ability to automate *T. infestans* identification plays a crucial role in vector surveillance, allowing public health organizations to respond promptly and efficiently to potential outbreaks. This approach reduces the dependency on manual identification, minimizing errors and improving the scalability of vector control programs.

Figure 1. Images of *T. infestans* showcasing its distinctive morphological patterns, including its elongated body shape, dark exoskeleton with reddish-orange markings, and prominent proboscis. These images are also examples of samples used in the dataset.

The dataset used for training the detection model consisted of images obtained from both freely available online sources, such as Google Images, and original field photographs taken specifically for this study. This approach ensured a diverse dataset, incorporating variations in lighting, angles, and backgrounds to improve model robustness. To ensure accurate training, all images were manually annotated, precisely labeling instances of *T. infestans* using bounding boxes. A manual filtering process was also performed to remove illustrative images, low-quality samples, and images where the insect was not clearly visible, ensuring that only high-quality images contributed to training. The dataset was then split into 91 images for training and 27 images for validation, maintaining a balanced distribution for effective model evaluation. During training, YOLOv8's built-in augmentation functionalities were applied dynamically, generating a variation of each image per epoch. Since the model was trained for 100 epochs, this resulted in an approximate 9,100 augmented training images and a similar augmentation process for the validation set. This dynamic augmentation strategy allowed the model to learn from a significantly expanded dataset, improving generalization to real-world conditions.

YOLOv8m was selected for this task as a balance between speed and accuracy, making it particularly well-suited for automated insect identification [18, 19]. Among the available YOLOv8 variants, the medium (m) version was chosen to ensure a lightweight yet powerful model capable of handling real-time detection efficiently [20]. Unlike traditional object detection models that require multiple stages for region proposal and classification, YOLO performs detection in a single pass, significantly reducing computational overhead while maintaining high precision [21]. This efficiency is crucial for real-time applications, such as field surveillance and mobile deployment, where rapid detection is necessary for timely interventions [11].

Other possible alternatives, such as Faster R-CNN and EfficientDet, were considered but ultimately not chosen due to their trade-offs [22, 23]. Faster R-CNN is known for its high accuracy but suffers from slower inference times due to its two-stage detection pipeline, making it less suitable for real-time processing [24]. EfficientDet, on the other hand, offers a balance between accuracy and speed, but it requires extensive hyperparameter tuning and more computational resources for optimal performance [25]. YOLOv8 incorporates improved feature extraction, better anchor-free detection, and an optimized model architecture, making it the most effective option for this study [20]. Its ability to generalize well across different environments while maintaining low latency and high precision positioned it as the ideal choice for detecting *T. infestans* in real-world conditions [18].

To improve the model's robustness and generalization across diverse real-world conditions, data augmentation was applied using YOLOv8's built-in augmentation functionalities. These techniques introduced controlled variations in the dataset, enhancing the model's ability to handle changes in lighting, orientation, and scale. The applied transformations included random rotations (up to 10 degrees), translations (shifting up to 10% of the image), scaling (50% zoom), shear transformations (10-degree distortion), and flipping (both vertical and horizontal with a 50% probability each). By leveraging YOLOv8's native augmentation methods, the model was exposed to a wider range of visual scenarios, reducing the risk of overfitting to specific conditions and improving its detection accuracy in real-world environments.

By diversifying the training data, these techniques significantly improved the model's generalization capability, reducing its sensitivity to minor variations in insect appearance and background clutter. This was particularly important given the diversity of real-world settings where *T. infestans* might be encountered, including different lighting conditions, backgrounds, and camera angles. The augmented dataset helped the model develop a more robust feature representation, ultimately leading to improved detection accuracy when applied to previously unseen images in real-world applications.

The training process for the YOLOv8m model was conducted on Google Colab, leveraging its GPU acceleration to efficiently handle the computational demands of deep learning. The training environment was configured with a Tesla T4 GPU, allowing for faster processing and reduced training times compared to CPU-based setups. The dataset, preprocessed with augmentation techniques, was fed into the YOLOv8m training pipeline using the default settings provided by the model's framework.

For this initial version of the model, no hyperparameter tuning was performed, meaning the training followed YOLOv8m's default configurations, including predefined learning rate, batch size, and anchor box settings. While these default parameters provided strong baseline performance, future iterations of the model may benefit from fine-tuning key hyperparameters to further optimize detection accuracy. Adjustments such as learning rate scheduling, batch size optimization, and anchor size adjustments could enhance the model's ability to detect *T. infestans* with greater precision, particularly in challenging real-world conditions. Moving forward, experimenting with transfer learning and adaptive training strategies may also contribute to improved detection performance, ensuring the model remains both accurate and efficient in operational settings.

After training, the YOLOv8m model was deployed for real-time detection using OpenCV, enabling live identification of *T. infestans* through a webcam feed. The system processes incoming video frames in real time, passing each frame through the trained YOLOv8m model to detect the presence of the insect. If a detection is made, the model generates bounding boxes around the identified insect, displaying them on the screen with confidence scores. To ensure high reliability while prioritizing the minimization of false negatives over false positives, a confidence threshold of 80% was applied. This means that only detections with a probability of 80% or higher are considered valid, reducing the likelihood of missing actual *T. infestans* specimens.

The pipeline follows an efficient loop where each video frame is captured, processed, and displayed in quick succession, maintaining smooth real-time performance. This setup allows for instant feedback, making it suitable for field applications where rapid identification of *T. infestans* is crucial for vector control programs. The use of OpenCV ensures the system remains lightweight and deployable on various hardware configurations, from personal computers to potential mobile and embedded applications. Future enhancements may include integrating additional post-processing techniques to further refine detection accuracy and reduce computational overhead.

This project stands out due to its custom dataset creation, real-time deployment capabilities, and scalability for future applications. Unlike generic object detection models, this model was specifically designed for detecting *T. infestans*, incorporating field-collected and manually labeled images to ensure high accuracy in real-world conditions. Its real-time deployment capability, powered by OpenCV and YOLOv8m, allows for immediate identification using a webcam with minimal latency, making it suitable for both laboratory monitoring and field applications. By applying a 90% confidence threshold, the system prioritizes high-accuracy detections, reducing false positives and improving reliability. Additionally, the project is scalable for mobile integration, with future plans to convert the trained model into lightweight formats such as TensorFlow Lite or ONNX for deployment on smartphones or embedded devices. This ensures accessibility for non-experts, enabling real-time vector surveillance in remote locations.

3. Implementation and Results

The complete implementation, including dataset preprocessing, model training scripts, and real-time detection code, is publicly available in the project's GitHub repository: `https://github.com/aaaimx/T_infestans_detection_YOLOv8m`.

After training, the YOLOv8m model was evaluated based on several key performance metrics to assess its accuracy and generalization ability. The model achieved a precision of 0.9606 and a recall of 0.9041, demonstrating a strong ability to correctly identify *T. infestans* while maintaining a relatively low rate of false negatives. Figure 2 illustrates the training performance, showing the loss curve convergence and mAP progression over epochs, which confirm the model's steady improvement.

The mean average precision at an Intersection over Union (IoU) threshold of 50% (mAP@50) reached 0.9588, meaning that the model correctly identified *T. infestans* with high accuracy when the predicted bounding boxes overlapped at least 50% with the ground truth annotations, demonstrating high detection accuracy under standard intersection-over-union thresholds. However, the mAP@50-95, which evaluates performance across multiple IoU thresholds, was 0.6539, suggesting that detection confidence varies depending on object overlap. The fitness score, calculated as 0.6844,

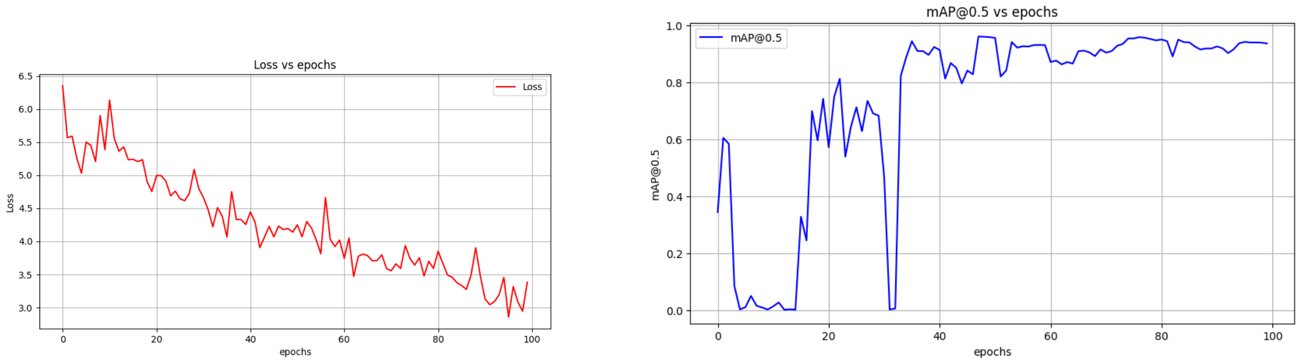

(a) **(b)**

Figure 2. Training Performance Graphs: (a) Loss curve over training epochs and (b) mAP progression during training, illustrating performance trends.

reflects the model's overall precision, recall, and IoU consistency. These results indicate that the model generalizes well to the validation set, achieving high accuracy in controlled conditions but may require further refinements to improve detection under more challenging real-world scenarios, such as variations in lighting, backgrounds, and occlusions.

To visually assess the model's performance, Figure 3 presents examples of *T. infestans* detections. The images illustrate successful identifications where the insect was accurately enclosed within bounding boxes, demonstrating the model's ability to detect the target species across different environments.

Additionally, to evaluate the specificity of the detection system, Figure 4 shows instances where *T. infestans* was detected while other insects, such as spiders and bed bugs, were not misclassified as *T. infestans*. These results highlight the model's ability to differentiate *T. infestans* from other visually similar insects, reducing the likelihood of false positives.

4. Discussion and Potential Impact

Vector-borne diseases pose a significant public health challenge, requiring efficient and accurate surveillance systems to mitigate their spread [26]. Traditional vector monitoring methods rely on manual insect identification, which can be time-consuming, labor-intensive, and prone to human error [27]. The integration of AI in vector surveillance presents a transformative approach, allowing for automated, scalable, and real-time detection of disease-carrying insects [28]. This study demonstrates the feasibility of using deep learning-based object detection, specifically YOLOv8m, to identify *T. infestans* with high accuracy and efficiency, ensuring that individuals—particularly those who may have been bitten—can quickly recognize the insect and take appropriate action, such as capturing it for evaluation by health professionals [29].

Beyond its role in vector control programs, automated identification of *T. infestans* is crucial for individual-level decision-making [30]. When a person is bitten by this insect, capturing it and bringing it to a health center for examination is essential, as experts can determine whether it carries *T. cruzi*, the parasite responsible for Chagas disease [31]. Misidentification or failure to detect *T. infestans* could result in missed medical evaluations, potentially leaving individuals unaware of their exposure risk [32]. Conversely, if the insect is not *T. infestans*, the person should pursue a different course of action, highlighting the importance of rapid and accurate identification [33]. By automating insect detection, this AI-driven system reduces dependence on expert entomologists, accelerates response times in both public health initiatives and personal health decisions, and enhances disease prevention efforts [34]. The ability to deploy AI-based detection models in remote or high-risk areas further strengthens surveillance capabilities, making it easier for public health organizations and individuals to monitor and respond effectively to potential threats [2].

The proposed YOLOv8m-based detection system offers a highly effective solution for vector surveillance, disease prevention, and individual health decision-making. With high detection accuracy, demonstrated by strong precision, recall, and mean average precision (mAP) scores, the model ensures reliable identification of *T. infestans* [35]. Its real-time processing capabilities, powered by OpenCV-based deployment, allow for instantaneous detection using a webcam or other imaging devices, making it suitable for both field applications and personal use [36]. Trained on a diverse dataset with data augmentation techniques, the model generalizes well across various lighting conditions, insect orientations, and backgrounds, making it adaptable to different environments. This model has the potential to impact individual health decisions, as it enables people to determine whether an insect in their home is *T. infestans* and requires medical attention [37]. By providing a scalable and cost-effective solution for early detection and monitoring, this AI-powered system may strengthen vector control efforts, empowering individu-

Figure 3. Examples of *T. infestans* detected by YOLOv8m. The bounding boxes highlight successful identifications under various lighting conditions and orientations.

als with fast, AI-assisted identification, and ultimately aiding in the prevention of Chagas disease transmission [38].

Despite its strong performance, the proposed detection system has certain limitations that need to be addressed for improved real-world applicability. One key challenge is the occurrence of false positives in cluttered backgrounds, where the model occasionally misidentifies objects with similar textures or shapes as *T. infestans*. Additionally, poor lighting conditions—such as dim environments or strong shadows—can reduce detection accuracy, leading to missed identifications (false negatives). Another limitation stems from the dataset constraints; since the model was trained on a small set of 91 original images, its generalization to highly varied or unseen environments may still be limited. Expanding the dataset with more diverse images, including additional field-collected samples under different environmental conditions, could help improve robustness. Moreover, fine-tuning hyperparameters, such as confidence thresholds, learning rates, and anchor sizes, may further enhance detection performance. Finally, testing the model in real-world settings, such as deploying it in vector surveillance programs, would provide valuable feedback to refine detection accuracy and adaptability in operational scenarios. Addressing these limitations will be crucial for optimizing the system for large-scale deployment and real-world usability.

While this study focuses on the detection of *T. infestans*, the underlying AI-based detection framework has the potential to be adapted for a wide range of applications beyond Chagas disease vector surveillance. Similar deep learning techniques could be applied to the identification of other disease-carrying insects that are often misidentified, such as other species of Triatoma that also transmit Chagas disease or similar-looking insects that are harmless. This could help improve public awareness and reduce unnecessary concern, while ensuring that true vectors are identified and handled correctly.

Additionally, this AI-driven approach could prove valuable in agriculture, where automated detection of crop-damaging pests (e.g., locusts, beetles, or caterpillars) could help farmers implement targeted pest control strategies, reducing pesticide overuse and improving crop yields. Beyond public health and agriculture, AI-powered biodiversity monitoring could benefit from such models by enabling the automated classification of insect species, contributing to ecological research and conservation efforts. By refining and expanding this detection framework, deep learning models could support real-time monitoring systems for vector control programs, environmental protection initiatives, and field-based entomological studies, demonstrating the far-reaching impact of AI in entomology, epidemiology, and ecological science.

To maximize the real-world impact of this detection system, several key advancements are planned for future deployment. One major focus is adapting the model for mobile and embedded systems, enabling lightweight deployment on smartphones, drones, or IoT devices for real-time vector surveillance in remote areas. Efforts are underway to convert the trained YOLOv8m model into optimized formats such as TensorFlow Lite or ONNX, reducing computational requirements while maintaining detection accuracy. Additionally, field testing in real-world environments is crucial to validating the model's performance under varied lighting conditions, diverse insect orientations, and natural backgrounds. Collaboration with entomologists and public health experts will provide critical feedback for refining the system, ensuring its practical usability in vector control programs.

Beyond mobile deployment, integrating edge computing and federated learning could further enhance offline detection capabilities, allowing the model to function in areas with limited or no internet connectivity. Edge AI devices could process detections locally, sending only essential metadata to centralized databases for broader epidemiological monitoring. These advancements would significantly improve accessibility, scalability, and efficiency in disease vector surveillance, agricultural pest management, and ecological monitoring, making AI-powered detection a viable tool for real-world applications.

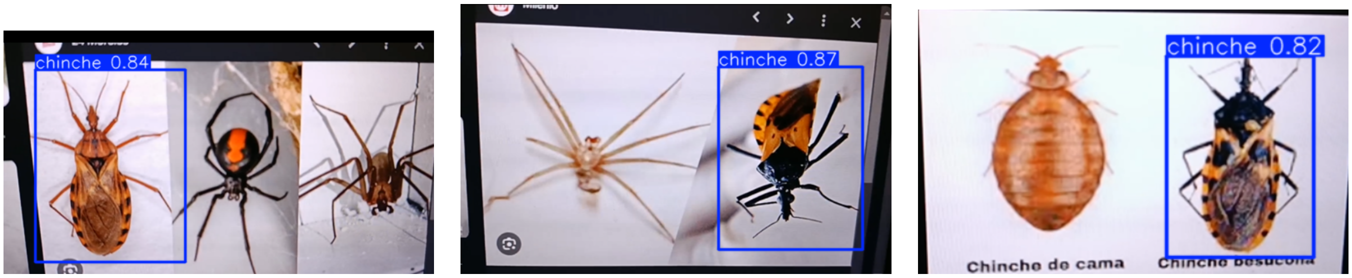

Figure 4. Examples of *T. infestans* detection alongside other insects. The model correctly detects *T. infestans* while avoiding false positives for other species such as spiders and bed bugs.

5. Conclusion

This study developed a YOLOv8m-based detection model to accurately identify *T. infestans*, offering a scalable tool for vector surveillance and individual health decision-making. The model demonstrated high detection accuracy, supported by strong precision, recall, and mean average precision (mAP) scores, ensuring reliable identification across diverse environments. Additionally, its real-time processing capabilities, enabled by OpenCV, allow for instantaneous detection, making it practical for both field applications and personal use. By automating the detection process, this system not only aids public health professionals but also helps non-experts recognize *T. infestans*, ensuring individuals take appropriate actions after a bite—either capturing the insect for examination or avoiding unnecessary concern in cases of misidentification. These advancements contribute to faster response times, improved disease monitoring, and enhanced public health strategies, ultimately reinforcing efforts to mitigate the spread of Chagas disease.

The proposed AI-powered system presents a scalable solution with the potential for widespread use in public health programs, particularly in resource-limited regions where expert identification is not readily available. By reducing dependence on trained entomologists, this model enhances accessibility, enabling quick and accurate recognition of *T. infestans* in real-world conditions. This capability ensures that individuals take the correct post-bite actions, such as submitting the insect for analysis or seeking alternative medical guidance if the insect is not a vector of Chagas disease. By minimizing misidentification and delays in medical intervention, this AI-driven system could play a role in reducing Chagas disease transmission risks, underscoring the potential of AI-based solutions in improving vector surveillance and public health outcomes.

Despite its strengths, the detection system faces challenges that must be addressed to enhance its real-world performance. False positives can occur in cluttered backgrounds, where objects with similar textures or shapes may be misidentified as *T. infestans*. Additionally, poor lighting conditions, such as dim environments or harsh shadows, can lead to false negatives, affecting detection accuracy. Another key limitation is the dataset size, as training on 9,100 images (included augmented images) may not provide sufficient diversity for robust generalization. Expanding the dataset with more field-collected samples and similar-looking non-vector insects could improve model reliability. Moreover, fine-tuning hyperparameters, such as confidence thresholds and learning rates, may further refine detection performance. Future work will focus on mobile deployment, real-world field testing, and collaborations with public health organizations to integrate this tool into vector control programs, ensuring practical usability in endemic regions.

References

[1] R. Castillo-Neyra, C. Barbu, R. Salazar, K. Borrini, C. Náquira, and M. Levy, "Host-seeking behavior and dispersal of triatoma infestans, a vector of chagas disease, under semi-field conditions," *PLoS Neglected Tropical Diseases*, vol. 9, 2015.

[2] J. González, M. Contreras, H. Schenone, H. Adaos, and R. Cabezas, "Chagas disease: impact of triatoma infestans control program in alto del carmen, huasco province, iii region atacama, chile," *Boletin chileno de parasitologia*, vol. 51, pp. 28–30, 1996.

[3] L. Diotaiuti, O. R. de Paula, P. Falcão, and J. Dias, "Evaluation of the chagas' disease vector control program in minas gerais, brazil, with special reference to triatoma sordida," *Bulletin of the Pan American Health Organization*, vol. 28, pp. 211–219, 1994.

[4] A. Loza, A. Talaga, G. Herbas, R. J. Canaviri, T. Cahuasiri, L. Luck, A. Guibarra, R. Gonçalves, J. A. Pereira, S. A. Gomez, A. Picado, L. Messenger, C. Bern, and O. Courtenay, "Systemic insecticide treatment of the canine reservoir of trypanosoma cruzi induces high levels of lethality in triatoma infestans, a principal vector of chagas disease," *Parasites Vectors*, vol. 10, 2017.

[5] V. JorgeGonzález, M. D. C. Contreras, H. Schenone, H. Adaos, and R. Cabezas, "Enfermedad de chagas: impacto del programa de control del triatoma infestans en la comuna de alto del carmen, provincia de huasco, iii región atacama, chile," *Boletin chileno de parasitologia*, vol. 51, pp. 28–30, 1996.

[6] R. Gürtler, G. M. Prokopec, L. Ceballos, C. Petersen, and O. Salomón, "Comparison between two artificial shelter units and timed manual collections for detecting peridomestic triatoma infestans (hemiptera: Reduviidae) in rural northwestern argentina," *Journal of Medical Entomology*, vol. 38, pp. 429–436, 2001.

[7] M. Cecere, R. Gürtler, D. Canale, R. Chuit, and J. E. Cohen, "The role of the peridomiciliary area in the elimination of triatoma infestans from rural argentine communities," *Revista panamericana de salud publica = Pan American journal of public health*, vol. 1, pp. 273–279, 1997.

[8] Y. Provecho, M. S. Gaspe, M. Fernández, and R. Gürtler, "House reinfestation with triatoma infestans (hemiptera: Reduviidae) after community-wide spraying with insecticides in the argentine chaco: A multifactorial process," *Journal of Medical Entomology*, vol. 54, pp. 646–657, 2017.

[9] J. Gurevitz, M. S. Gaspe, G. F. Enriquez, Y. Provecho, U. Kitron, and R. Gürtler, "Intensified surveillance and insecticide-based control of the chagas disease vector triatoma infestans in the argentinean chaco," *PLoS Neglected Tropical Diseases*, vol. 7, 2013.

[10] A. Teixeira, J. Ribeiro, R. Morais, J. Sousa, and A. Cunha, "A systematic review on automatic insect detection using deep learning," *Agriculture*, 2023.

[11] T. D. C. Júnior and R. Rieder, "Automatic identification of insects from digital images: A survey," *Comput. Electron. Agric.*, vol. 178, p. 105784, 2020.

[12] L. Souifi, A. Mdhaffar, I. Rodriguez, M. Jmaiel, and B. Freisleben, "A systematic literature review on insect detection in images," *Signal, Image Processing and Embedded Systems Trends*, 2022.

[13] D. Zeng, Z. Cao, and D. Neill, "Artificial intelligence–enabled public health surveillance—from local detection to global epidemic monitoring and control," *Artificial Intelligence in Medicine*, pp. 437–453, 2020.

[14] S. Saran and P. Singh, "Systematic review on citizen science and artificial intelligence for vector-borne diseases," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2024.

[15] J. Abdulkarim, A. A. Musa, Y. Abdullahi, and U. H. Yamman, "Artificial intelligence may help in the containment of cholera in nigeria," *OIRT Journal of Information Technology*, 2022.

[16] iNaturalist Mexico, "Triatoma dimidiata - photo gallery," 2024, accessed: March 3, 2024. [Online]. Available: https://mexico.inaturalist.org/taxa/308558-Triatoma-dimidiata/browse_photos

[17] Pexels, "Insecto de 6 patas naranja-negro colgando de una planta de hoja verde," 2024, accessed: March 3, 2024. [Online]. Available: https://www.pexels.com/es-es/foto/insecto-de-6-patas-naranja-negro-colgando-de-una-planta-de-hoja-verde-68513/

[18] A. Widayani, A. M. Putra, A. R. Maghriebi, D. Z. C. Adi, and M. H. F. Ridho, "Review of application yolov8 in medical imaging," *Indonesian Applied Physics Letters*, 2024.

[19] C. M. Badgujar, A. Poulose, and H. Gan, "Agricultural object detection with you look only once (yolo) algorithm: A bibliometric and systematic literature review," *ArXiv*, vol. abs/2401.10379, 2024.

[20] M. G. Ragab, S. J. Abdulkadir, A. Muneer, A. Alqushaibi, E. H. H. Sumiea, R. Qureshi, S. M. Al-Selwi, and H. Alhussian, "A comprehensive systematic review of yolo for medical object detection (2018 to 2023)," *IEEE Access*, vol. 12, pp. 57 815–57 836, 2024.

[21] Z. Piao, M. Meng, H. Yang, T. Xue, Z. Jia, and W. Liu, "Distinguishing between aldosterone-producing adenomas and non-functional adrenocortical adenomas using the yolov5 network," *Acta Radiologica*, 2024.

[22] S. J. Yang, Y. Lu, X. Zheng, Y. J. Zhang, F. Xin, P. Sun, Y. Li, S. Liu, S. Li, Y. T. Guo, and S. Liu, "Establishment and clinical testing of pancreatic cancer faster r-cnn ai system based on fast regional convolutional neural network," *Zhonghua Wai Ke Za Zhi [Chinese Journal of Surgery]*, vol. 58, pp. 520–524, 2020.

[23] Q. Ge, T.-Y. Lin, M. Maire, and P. Dollár, "Efficientdet: Scalable and efficient object detection," *arXiv preprint arXiv:1911.09070*, 2019. [Online]. Available: https://arxiv.org/abs/1911.09070

[24] L. Ding, G. Liu, B.-C. Zhao, Y. Zhou, S. Li, Z. Zhang, Y. Guo, A. qin Li, Y. Lu, H. Yao, W. Yuan, G. Wang, D. Zhang, and L. Wang, "Artificial intelligence system of faster region-based convolutional neural network surpassing senior radiologists in evaluation of metastatic lymph nodes of rectal cancer," *Chinese Medical Journal*, vol. 132, pp. 379 – 387, 2019.

[25] Q. V. L. Mingxing Tan, Ruoming Pang, "Efficientdet: Scalable and efficient object detection," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10 781–10 790, 2020. [Online]. Available: https://arxiv.org/abs/1911.09070

[26] Y. Wu, J. Wang, Q. Liu, T. Li, M. Luo, and Z. Gong, "Practice of integrated vector surveillance of arthropod vectors, pathogens and reservoir hosts to monitor the occurrence of tropical vector-borne diseases in 2020 in zhejiang province, china," *Frontiers in Veterinary Science*, 2022.

[27] J. E. Nava-Doctor, C. Sandoval-Ruiz, and A. Fernández-Crispín, "Knowledge, attitudes, and practices regarding vector-borne diseases in central mexico," *Journal of Ethnobiology and Ethnomedicine*, 2021.

[28] S. Saran and P. Singh, "Systematic review on citizen science and artificial intelligence for vector-borne diseases," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2024.

[29] M. Miranda, C. Barceló, D. Arnoldi, X. Augsten, K. Bakran-Lebl, G. Balatsos, M. Bengoa, P. Bindler, K. Boršová, M. Bourquia, D. Bravo-Barriga, V. Čabanová, B. Caputo, M. Christou, S. Delacour, R. Eritja, O. Fassi-Fihri, M. Ferraguti, E. Flacio, E. Frontera, H. Fuehrer, A. García-Pérez, P. Georgiades, S. Gewehr, F. Goiri, M. González, M. Gschwind, R. Gutiérrez-López, C. Horváth, A. Ibáñez-Justicia, V. Jani, P. Kadriaj, K. Kalan, M. Kavran, A. Klobučar, K. Kurucz, J. Lucientes, R. Lühken, S. Magallanes, G. Marini, A. Martinou, A. Michelutti, A. Mihalca, T. Montalvo, F. Montarsi, S. Mourelatos, N. Muja-Bajraktari, P. Müller, G. Notarides, H. Osório, J. Oteo, K. Oter, I. Pajović, J. Palmer, S. Petrinić, C. Răileanu, C. Ries, E. Rogozi, I. Ruíz-Arrondo, I. Sanpera-Calbet, N. Sekulić, K. Sevim, K. Sherifi, C. Silaghi, M. Silva, N. Sokolovska, Z. Soltész, T. Șulescu, J. Šušnjar, S. Teekema, A. Valsecchi, M. Vasquez, E. Velo, A. Michaelakis, W. Wint, D. Petrić, F. Schaffner, and A. della Torre, "Aimsurv: First pan-european harmonized surveillance of aedes invasive mosquito species of relevance for human vector-borne diseases," *GigaByte*, 2022.

[30] O. F. Leite, M. J. Alves, S. S. Souza, R. C. Mayo, V. Andrade, C. Souza, O. Rangel, S. S. Oliveira, V. L. Lima, V. L. Rodrigues, M. E. Carvalho, C. Casanova, and D. M. Wanderley, "Triatoma infestans in area under entomological surveillance for chagas' disease in são paulo state, brazil," *Revista da Sociedade Brasileira de Medicina Tropical*, vol. 34 5, pp. 437–43, 2001.

[31] G. Marti, M. G. Echeverría, M. L. Susevich, J. Becnel, S. Pelizza, and J. García, "Prevalence and distribution of parasites and pathogens of triatominae from argentina, with emphasis on triatoma infestans and triatoma virus trv," *Journal of Invertebrate Pathology*, vol. 102 3, pp. 233–237, 2009.

[32] M. Rojas-Cortez, M. Pinazo, J. Gascón, E. Gamarra, R. M. Grágeda, R. Fernandez, E. Rueda, J. Pinto, H. M. Anzoleaga, Y. E. Caballero, G. S. Urioste, J. Sánchez, R. Saravia, M. Castellon, W. García, L. O. Daza, I. Mur, D. Lozano, F. Carvalho-Costa, F. A. Monteiro, and F. Torrico, "Community-based entomological surveillance in three chagas disease-endemic regions in sub-andean bolivia," *Transactions of the Royal Society of Tropical Medicine and Hygiene*, 2021.

[33] R. Gurgel-Gonçalves, F. Abad-Franch, M. R. de Almeida, M. Obara, R. D. C. M. D. Souza, J. A. de Sena Batista, and D. Rocha, "Triatodex, an electronic identification key to the triatominae (hemiptera: Reduviidae), vectors of chagas disease: Development, description, and performance," *PLoS ONE*, vol. 16, 2021.

[34] M. E. de Carvalho, R. A. da Silva, V. L. Rodrigues, and C. D. de Oliveira, "The chagas disease control program of the são paulo state: the contribution of serology to the epidemiological investigation of triatomine-infested domiciliary units during the 1990s," *Cadernos de Saúde Pública*, vol. 18 6, pp. 1695–703, 2002.

[35] N. R. C. de Lima, "Distribuição, infecção natural e fontes alimentares de triatomíneos coletados em municípios do estado do rio grande do sul e ações educativas na prevenção e vigilância da doença de chagas," 2017.

[36] R. Piccinali, M. Gaunt, and R. Gürtler, "A microsatellite-based analysis of house infestation with triatoma infestans (hemiptera: Reduviidae) after insecticide spraying in the argentine chaco," *Journal of medical entomology*, vol. 55 3, pp. 609–619, 2018.

[37] A. R. de Arias, F. Abad-Franch, N. Acosta, E. López, N. González, E. Zerba, G. Tarelli, and H. Masuh, "Post-control surveillance of triatoma infestans and triatoma sordida with chemically-baited sticky traps," *PLoS Neglected Tropical Diseases*, vol. 6, 2012.

[38] L. Abrahan, M. Cavallo, and I. Amelotti, "Impact of involving the community in entomological surveillance of triatoma infestans (klug, 1834) (hemiptera, triatominae) vectorial control," *Parasites Vectors*, vol. 14, 2021.

Journal of Artificial Intelligence and Computing Applications (2024) 2(2): 39–47

Journal
of
Artificial Intelligence
and
Computing Applications

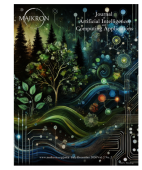

Exploration Paper

An Exploratory Application of Empirical Mode Decomposition and Recurrent Neural Networks for Meteorological Time Series Prediction

Jesús Humberto Sarabia-Osorio[1,2]

[1]Sarchentu Capital, Bonampak SM 06 LT 01, 77500, Cancún, Quintana Roo, México
[2]International Institute for Intelligent Technologies

ABSTRACT

Accurate short-term weather forecasting remains a critical yet complex task, particularly in tropical regions where high variability and abrupt climatic shifts can have immediate impacts on agriculture, infrastructure, and public safety. Traditional statistical methods often struggle to capture the non-linear and multi-scale nature of meteorological time series, limiting their effectiveness in localized forecasting scenarios. To address this challenge, this paper presents an exploratory prototype that combines Empirical Mode Decomposition with Recurrent Neural Networks, specifically Long Short-Term Memory (LSTM) architectures. Daily data on temperature, humidity, and atmospheric pressure from Mérida, Yucatán (2000–2018) were decomposed into Intrinsic Mode Functions, which served as input features for training separate LSTM models. The hybrid system achieved promising results, particularly for temperature and humidity, capturing key short-term patterns while highlighting limitations in pressure forecasting. These findings suggest that EMD-based preprocessing can enhance neural sequence models in dynamic forecasting contexts, offering a pathway toward more adaptive, data-driven approaches in weather-sensitive applications.

Keywords: empirical mode decomposition, recurrent neural networks, weather forecasting

1. Introduction

Weather forecasting remains one of the most challenging tasks in data science due to the inherently dynamic and complex nature of atmospheric systems [1, 2]. In regions with tropical climates, such as Mérida in southeastern Mexico, the difficulty is compounded by high variability, localized phenomena, and abrupt transitions in key meteorological variables [3, 4]. Accurately predicting conditions like temperature, humidity, and pressure is crucial not only for day-to-day planning but also for sectors such as agriculture, infrastructure management, and public safety [5]. However, traditional forecasting models often struggle to cope with the non-linear and non-stationary behavior exhibited by real-world weather data [6, 7]. In this context, the development of novel methods capable of identifying patterns in chaotic time series has become increasingly relevant, particularly in light of growing interest in localized climate-sensitive decision-making [8].

Despite advances in statistical modeling, traditional forecasting methods often fall short when dealing with the intricacies of real-world meteorological data [9, 10].

E-mail address: jsarabia@sarchentu.com
https://doi.org/10.5281/zenodo.15098399
 ISSN pending

These methods typically rely on assumptions of linearity and stationarity that rarely hold in practice [11], especially when variables exhibit abrupt shifts, irregular cycles, or long-term dependencies. As a result, they tend to oversimplify the underlying dynamics, leading to suboptimal predictions [12]. Furthermore, standard time series techniques often fail to account for the hierarchical nature of temporal patterns—some of which unfold over hours, while others emerge over weeks or months [13]. This limitation becomes particularly evident in tropical regions, where weather systems are influenced by both seasonal cycles and transient atmospheric disturbances [14]. Capturing such multi-scale dependencies requires more flexible, adaptive approaches capable of learning directly from the data [15].

To address these challenges, this project explores a hybrid approach that combines Empirical Mode Decomposition (EMD) [16] with Recurrent Neural Networks (RNNs), specifically Long Short-Term Memory (LSTM) architectures [17, 18]. EMD is an adaptive signal processing technique that decomposes complex, non-linear time series into a finite set of oscillatory components known as Intrinsic Mode Functions (IMFs). These IMFs capture meaningful temporal structures at different scales, effectively isolating patterns while reducing noise and irregularities present in the original data. Once the meteorological variables are decomposed into IMFs, they can be used as input features for training an LSTM model. This type of RNN is particularly well-suited for sequential data, as it retains information over long periods and is capable of modeling temporal dependencies [19]. By combining the noise-filtering capabilities of EMD with the sequence learning strength of LSTM networks, the proposed method aims to improve short-term forecasting accuracy in a data-driven and adaptive manner.

This work presents the design and implementation of a prototype system that integrates EMD and RNNs for the short-term forecasting of key meteorological variables in Mérida, Yucatán. Specifically, the project focuses on predicting daily values of temperature, humidity, and atmospheric pressure using historical data provided by the Comisión Nacional del Agua (CONAGUA). The goal is not to develop a production-grade forecasting tool, but rather to explore the feasibility and performance of the proposed hybrid approach in a real-world context. As such, the system operates as an early-stage proof of concept, emphasizing methodological exploration, performance evaluation, and identification of practical challenges in data preprocessing, model training, and result interpretation.

This project represents an exploratory effort to prototype an AI-based solution for a real-world forecasting challenge. Rather than offering a finalized or production-ready system, it emphasizes the early stages of implementation, highlighting both the potential and limitations of applying advanced learning techniques to meteorological data. By experimenting with the integration of EMD and recurrent neural architectures, the work contributes practical insights into how data-driven models can be developed and adapted for complex time series tasks. Although focused on a localized use case, the approach offers a foundation for future applications in fields such as agriculture, disaster preparedness, and environmental monitoring, where improved short-term weather prediction can lead to more informed and timely decisions.

2. Project Description

The proposed system follows a modular architecture designed to process raw meteorological time series and generate short-term forecasts through a hybrid machine learning pipeline. At its core, the system integrates EMD as a preprocessing step with a RNN based on LSTM units for prediction. The pipeline begins by taking daily meteorological variables as input—specifically temperature, humidity, and atmospheric pressure—collected over a multi-year period. These raw time series are then subjected to EMD to extract IMFs that isolate oscillatory patterns and reduce noise. The resulting IMFs serve as enhanced input features for the LSTM model, which is trained to learn sequential dependencies and generate forecasts for the following day. This sequential flow—from raw data to decomposition to sequence modeling—defines a structured yet flexible prototype aimed at improving prediction accuracy in complex, real-world weather scenarios.

The dataset used in this project consists of daily meteorological records for the city of Mérida, Yucatán, spanning from the year 2000 to 2018. The data was obtained from the CONAGUA, the national authority responsible for weather monitoring in Mexico. Three primary variables were selected for analysis: temperature (in degrees Celsius), relative humidity (percentage), and atmospheric pressure (in hPa). Each variable was recorded at consistent daily intervals, resulting in a structured time series with over 6,500 data points per variable. Prior to modeling, the raw data underwent a preprocessing stage that included the removal of missing or anomalous values, normalization of scales, and conversion into CSV format for compatibility with machine learning tools. The dataset was then divided into training and validation subsets, maintaining chronological order to preserve the integrity of the temporal structure essential for time series forecasting. The seasonal patterns and short-term variability of the three variables are illustrated in Figure 1.

EMD was applied individually to each meteorological variable to extract its underlying oscillatory modes in the form of IMFs. This adaptive, data-driven technique decomposes a non-linear and non-stationary time series into a set of components that represent localized fluctuations at different temporal scales. The decomposition was performed using the PyEMD library, a Python implementation that automates the sifting process without requiring predefined basis functions. For each variable, the EMD process yielded a series of IMFs ranked from high-frequency (short-term variations) to low-frequency (long-term trends). To balance model

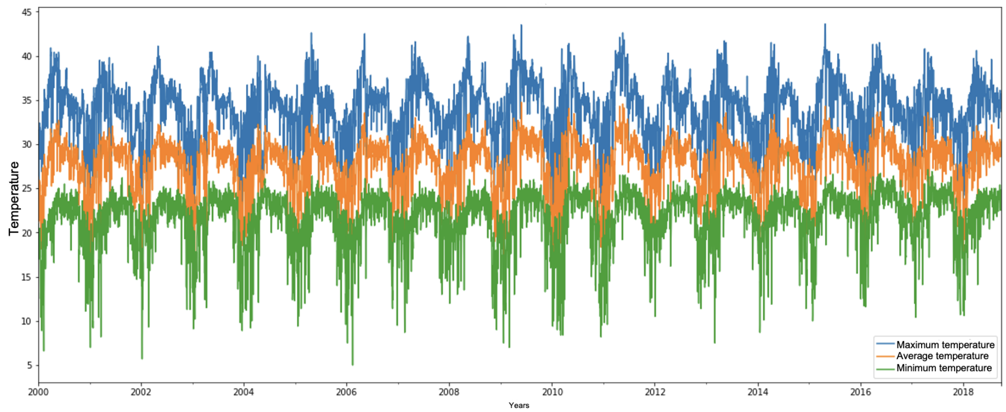

(a) Daily temperature (°C)

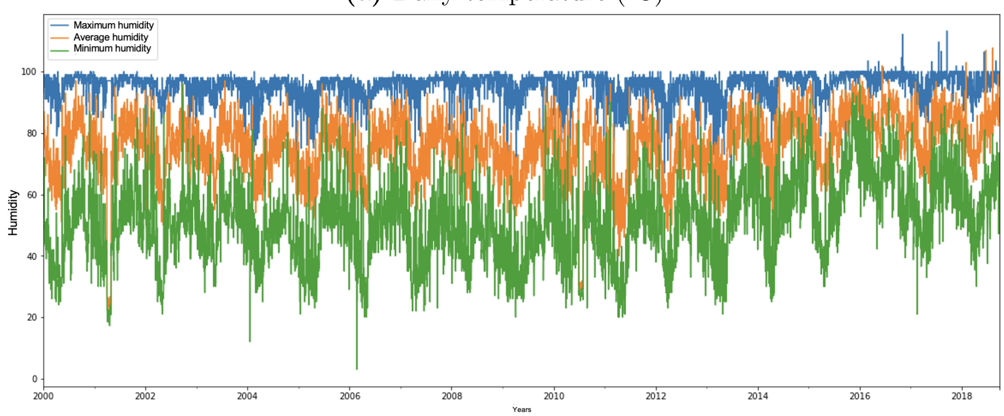

(b) Daily relative humidity (%)

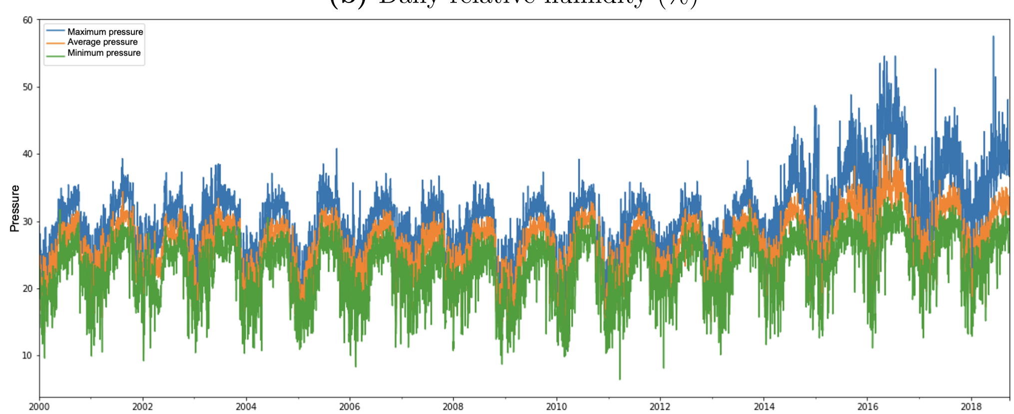

(c) Daily atmospheric pressure (hPa)

Figure 1. Meteorological Time Series for Mérida (2000–2018). Daily values of the three meteorological variables used in this study: (a) temperature, (b) relative humidity, and (c) atmospheric pressure, as recorded by CONAGUA. These raw time series illustrate seasonal variability and short-term fluctuations that make forecasting a challenging task.

complexity and performance, the five most informative IMFs were selected for each variable and used as input features for the forecasting model. This step allowed the neural network to learn from a cleaner, multi-scale representation of the data, reducing the impact of noise and enhancing its ability to capture relevant temporal patterns. An example of the decomposition applied to the temperature variable is shown in Figure 2.

The forecasting component of the prototype was built using a RNN with LSTM units, designed to model the temporal dynamics of the decomposed meteorological data. Each input sequence to the LSTM consisted of five IMFs per variable, spanning a historical window of five consecutive days. The model was trained to predict the value of the corresponding meteorological variable for the next day, establishing a one-day forecasting horizon. Separate LSTM models were trained for temperature, humidity, and pressure, allowing each network to specialize in the dynamics of its respective variable. The architecture included a single LSTM layer followed by a dense output layer, using the Adam optimizer and mean squared error as the loss function. The use of IMFs as input features represents a distinctive departure from conventional raw-sequence modeling, providing the network with richer, frequency-aware representations that enhance its ability to learn from complex time series data.

The design of the prototype was guided by the need to address the limitations of traditional models in capturing the complexity of meteorological time series. EMD was selected as a preprocessing step due to its capacity to handle non-linear, non-stationary signals without requiring a priori assumptions about the data's structure. This made it particularly well-suited for tropical weather patterns, which often display irregular and multiscale behavior. LSTM networks were chosen over other machine learning models because of their proven effectiveness in modeling long-term dependencies in sequential data. One of the innovative aspects of the prototype lies in the use of selected IMFs—especially those representing higher-frequency components—as input features, allowing the network to focus on short-term fluctuations that are critical for daily forecasting. Although the system remains in an exploratory stage, the modularity of the design and the coupling of signal decomposition with deep learning present a promising direction for future refinement and application.

3. Implementation and Results

The implementation of the prototype was carried out using Python as the primary development environment, leveraging a combination of specialized libraries for signal processing and deep learning. The PyEMD library was used to perform EMD on each meteorological variable, producing a set of IMFs that were stored and managed using NumPy and Pandas for efficient manipulation. For model development and training, the Keras API within TensorFlow was employed to build and compile the LSTM architectures. Each variable—temperature, humidity, and pressure—was processed separately through the EMD step, generating five IMFs per variable, which were then fed into independent LSTM models. The training pipeline included data normalization, reshaping of input sequences, and batching for efficient training. The experiments were conducted on a standard consumer-grade laptop with 16 GB of RAM and no dedicated GPU, with each model requiring approximately 10 to 20 minutes to train, depending on sequence length and hyperparameter settings.

Each LSTM model was trained independently for the three target variables—temperature, humidity, and atmospheric pressure—using a consistent training procedure tailored to the temporal structure of the data. Input sequences were generated using a sliding window of five previous days' IMFs to predict the value for the next day. The models were trained over 100 epochs with a batch size of 32, using an 80/20 split between training and validation data. Dropout layers with a rate of 0.2 were incorporated to mitigate overfitting, and the Adam optimizer was used with a default learning rate. Mean squared error (MSE) served as the loss function and primary performance metric. Training progress was monitored via loss curves, and although no early stopping mechanism was applied, performance was assessed visually through plots of training and validation loss to detect signs of underfitting or overfitting. This manual tuning process provided insight into the model's learning behavior across variables and informed minor adjustments in model architecture and input formatting.

The models achieved promising results in short-term forecasting, with performance varying slightly across the three meteorological variables. For temperature prediction, the LSTM model trained on EMD-derived IMFs yielded the best accuracy, reaching a mean squared error (MSE) of 0.75 on the validation set. Humidity forecasts followed closely with an MSE of 1.08, while pressure predictions proved more challenging, resulting in a slightly higher error of 1.32. Qualitative assessments using line plots of predicted versus actual values revealed that the models were particularly effective at capturing short-term fluctuations and general trends in temperature and humidity. In contrast, pressure exhibited more irregular patterns that the model struggled to anticipate consistently, likely due to its lower variability and subtler temporal shifts. An example of the forecast performance for temperature is shown in Figure 3, where the EMD-LSTM model's predictions closely track the actual values. Despite some limitations, the results demonstrate that the hybrid approach is capable of learning meaningful representations from decomposed time series and producing forecasts that align well with observed data, especially for highly dynamic variables.

During implementation, several practical challenges emerged that shaped the development process and highlighted areas for future refinement. One of the most notable difficulties was the occasional instability of the EMD algorithm when applied to noisy or irregular segments of the time series, which sometimes produced

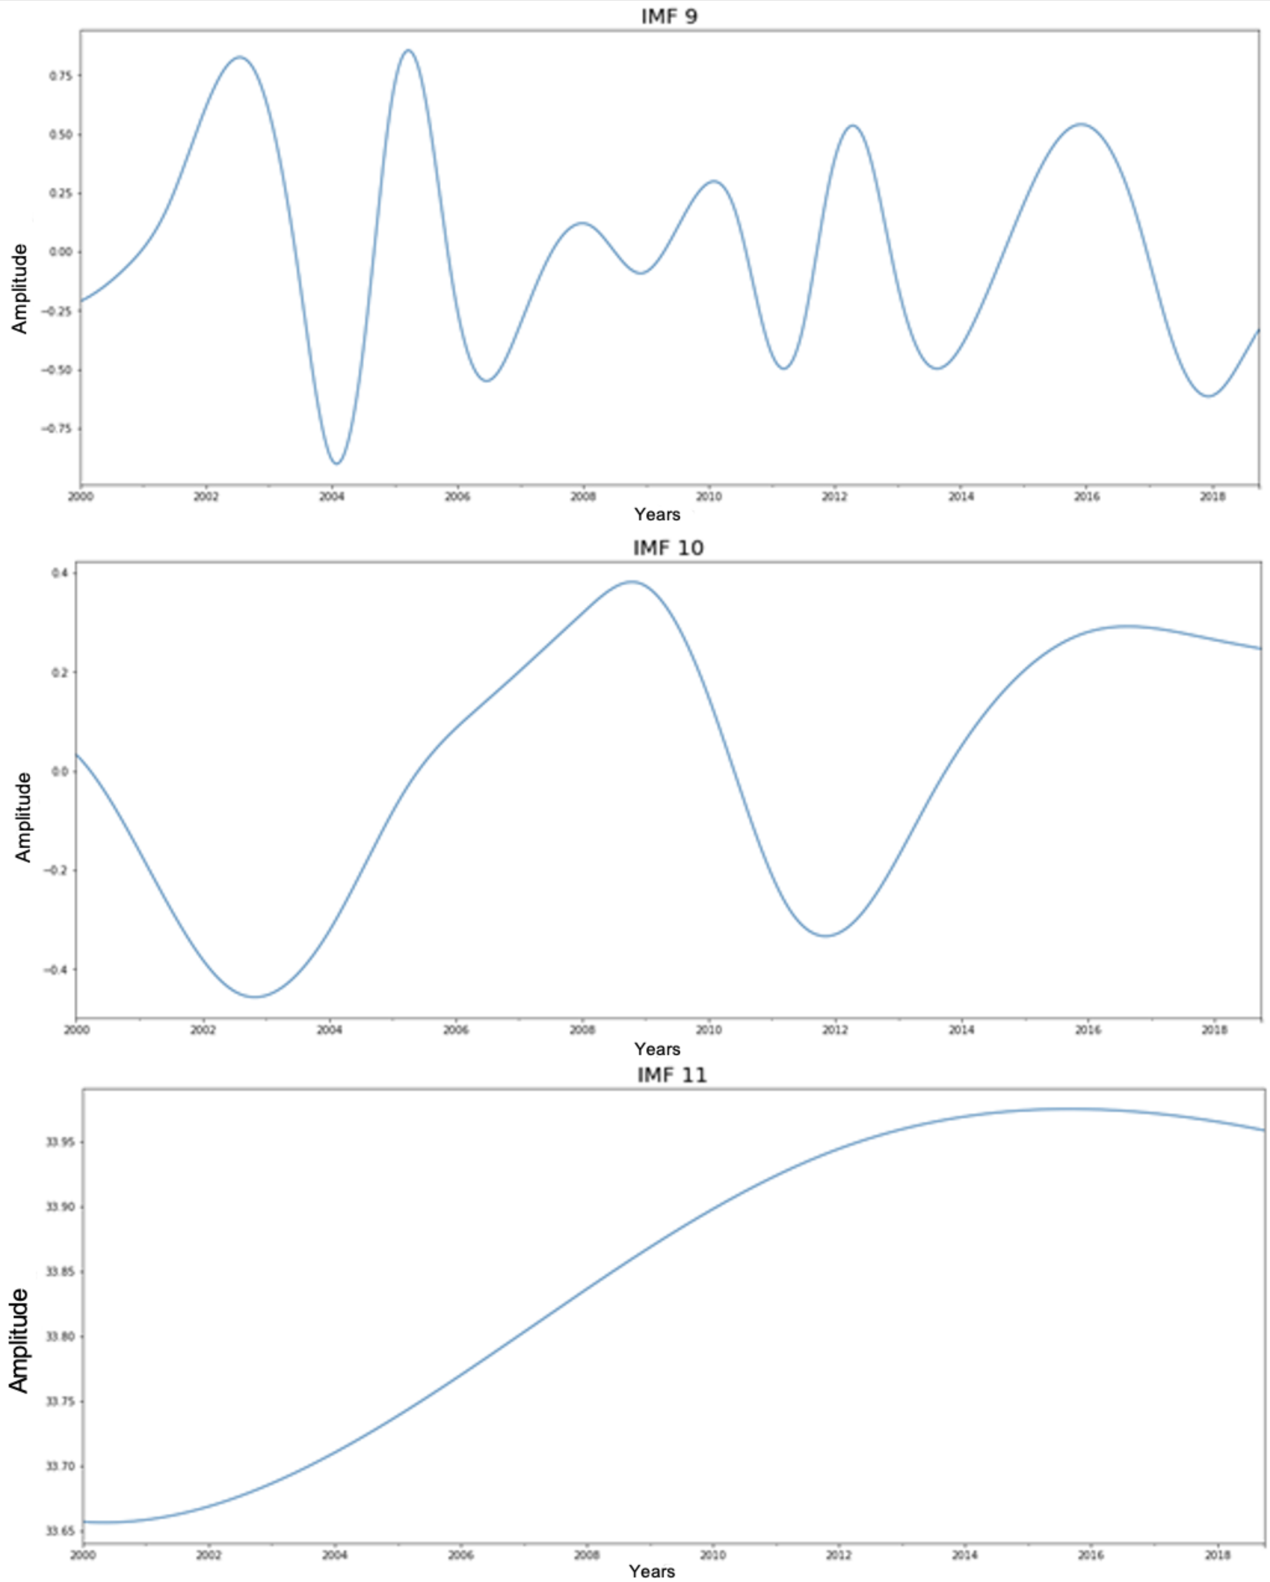

Figure 2. Intrinsic Mode Functions (IMFs) Extracted from Temperature Time Series via EMD. Decomposition of the temperature time series into IMFs using Empirical Mode Decomposition. The IMFs represent distinct oscillatory components at different frequency scales and serve as input features for the LSTM model, enabling the network to learn from multi-scale temporal patterns.

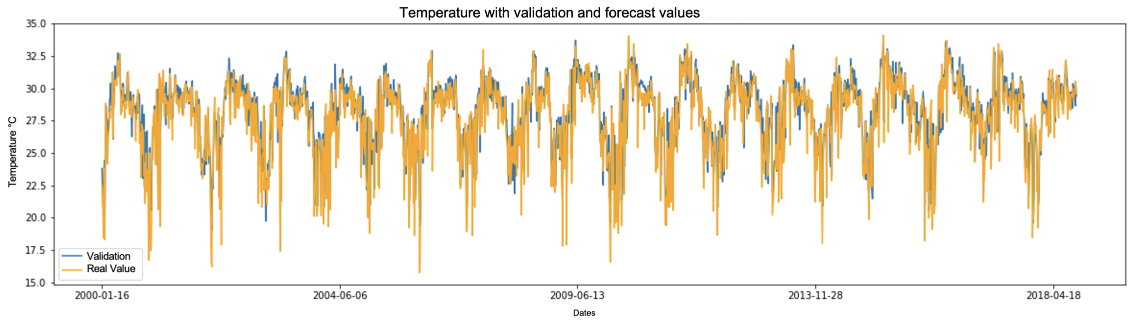

Figure 3. Forecast Results Using the EMD-LSTM Model for Temperature. Comparison between actual and predicted temperature values on the validation set using the EMD-LSTM model. The hybrid approach effectively captures short-term trends and daily fluctuations, demonstrating the potential of IMF-based neural forecasting for dynamic meteorological variables.

overlapping or distorted IMFs. This required manual inspection and, in some cases, adjustments to the preprocessing pipeline to ensure the decomposition yielded consistent and meaningful components. On the modeling side, tuning LSTM networks with a relatively small dataset proved to be delicate; minor changes in window size, batch configuration, or dropout rate significantly impacted validation performance. Additionally, while temperature and humidity exhibited clear temporal dependencies, pressure showed less predictable behavior, making it harder for the model to generalize. These limitations suggest the need for more robust IMF selection criteria, possible data augmentation strategies, and exploration of ensemble or multi-variate architectures in future iterations to enhance stability and performance across all variables.

4. Discussion and Potential Impact

The results obtained from this prototype suggest that the combination of EMD and LSTM networks offers a viable strategy for short-term weather forecasting, particularly in environments characterized by high temporal variability. By isolating frequency components through EMD and feeding them into a memory-based neural architecture, the system was able to effectively capture and reproduce short-term fluctuations in temperature and humidity. This indicates that decomposing time series into simpler, interpretable components can significantly enhance the learning capacity of sequence models, especially when working with relatively small datasets. The consistency in trend alignment between predicted and actual values further supports the notion that hybrid architectures can offer meaningful improvements over traditional end-to-end learning on raw sequences. These observations point toward the potential of refining and scaling this approach for broader forecasting scenarios, where data complexity and noise often limit the effectiveness of standard methods.

By leveraging EMD to preprocess complex time series, the system reduces noise and exposes latent temporal structures that are more readily learnable by recurrent networks like LSTMs. This layered approach demonstrates that coupling domain-agnostic signal processing methods with deep learning can improve forecasting outcomes without requiring specialized features or handcrafted inputs. Moreover, the architecture's modularity and adaptability suggest that it could be extended to a wide range of applications involving time-dependent data—such as energy demand forecasting, financial market analysis, or patient monitoring in healthcare. The ability to distill relevant patterns from noisy signals using this hybrid strategy opens up promising avenues for AI systems designed to operate under real-world constraints and data imperfections.

Accurate short-term weather forecasts hold substantial value across multiple sectors, particularly in regions where sudden climatic variations can have immediate consequences [20]. In agriculture, timely predictions of temperature and humidity can inform irrigation schedules, pest control measures, and crop protection strategies, ultimately improving yields and resource efficiency [21]. Urban planning and infrastructure management also benefit from short-range forecasts, which can aid in traffic regulation, drainage planning, and energy distribution during periods of extreme weather [22]. In the realm of public safety, early warnings based on localized forecasts can enhance emergency preparedness and response to heatwaves, storms, or unexpected weather shifts [23]. Even at a prototype stage, the system developed in this project illustrates the feasibility of building lightweight, adaptable tools that translate raw meteorological data into actionable insights, laying the groundwork for more robust decision-support systems tailored to the needs of specific communities or institutions.

While the current prototype demonstrates the potential of the EMD-LSTM approach, several avenues remain open for future improvement and exploration. Expanding the dataset to include more recent years or data from additional meteorological stations could enhance the model's generalizability and robustness. Further research is also needed to refine the IMF selection process, potentially incorporating automated criteria or relevance-based filtering to improve signal quality and model input. Exploring multivariate models that can jointly predict multiple variables, or ensemble approaches that combine different architectures, may yield more comprehensive forecasts. Additionally, comparing the LSTM performance with alternative architectures such as Gated Recurrent Units (GRUs) [24] or Transformer-based models [25] could offer insight into trade-offs between complexity and accuracy. Finally, integrating external data sources—such as satellite imagery, radar data, or environmental indices—could enrich the feature space and support more context-aware predictions, bringing the system closer to real-world operational deployment.

This project set out to explore the feasibility of using a hybrid approach that combines EMD with LSTM networks for short-term weather forecasting in a tropical urban setting. Motivated by the challenges posed by non-linear and non-stationary meteorological data, particularly in regions like Mérida, the goal was to develop a prototype capable of learning meaningful temporal patterns from raw climate observations. The core methodology involved decomposing daily time series data into IMFs to isolate key oscillatory components, and then using these components as input features for LSTM models trained to predict temperature, humidity, and atmospheric pressure. The project remained exploratory in nature, focusing on the practical steps and limitations involved in building a working AI-based forecasting system from real-world data.

The results of the prototype highlight the potential of the EMD-LSTM combination to improve forecasting accuracy, particularly for highly dynamic variables such as temperature and humidity. The use of intrinsic mode functions as input features allowed the models to learn from cleaner, frequency-resolved representations of the data, leading to lower prediction errors and better alignment with observed short-term patterns. While the model encountered some limitations in predicting pressure, overall performance metrics demonstrated that the hybrid approach outperformed what might be expected from standard modeling on raw time series alone. These findings underscore the practical value of integrating signal processing and neural modeling in applied AI contexts, especially when working with complex, real-world datasets. By addressing a concrete forecasting problem in a localized environment, the project contributes to the broader effort of translating AI research into functional solutions with societal relevance.

As an exploratory effort, this project demonstrates the early-stage viability of combining empirical decomposition and deep learning for meteorological forecasting, while also revealing important directions for continued development. The prototype serves as a foundation for further experimentation with larger and more diverse datasets, more sophisticated decomposition strategies, and expanded model architectures capable of handling multivariate interactions. Future research could aim to operationalize the system as part of decision-support platforms, particularly in climate-sensitive sectors such as agriculture, disaster risk management, or urban planning. Additionally, integrating external data sources and comparing performance across alternative neural models would help refine the system's predictive capabilities and adaptability. By advancing from concept to implementation, this work opens a pathway toward scalable, AI-driven forecasting solutions tailored to the unique demands of specific regions and applications.

CRediT authorship contribution statement

Jesús H. Sarabia-Osorio: Conceptualization, Methodology, Software, Validation, Formal analysis, Data curation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used ChatGPT in order to improve readability. After using this tool, the author reviewed and edited the content as needed and took full responsibility for the content of the publication.

Declaration of competing interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The author gratefully acknowledges the Comisión Nacional del Agua (CONAGUA) for providing access to the meteorological data used in this study.

References

[1] U. Huiningsumbam, A. Jain, and N. Verma, "Artificial neural network for weather forecasting: A review," in *2020 IEEE International Conference on Technology, Engineering, Management for Societal impact using Marketing, Entrepreneurship and Talent (TEMSMET)*. IEEE, 2020, pp. 1–7.

[2] M. Hasan, K. Islam, M. S. Rahman, and S. Li, "Weather forecasting using artificial neural network," in *Recent Findings in Intelligent Computing Techniques*, 2019, pp. 171–180.

[3] R. I. Rasel, N. Sultana, and P. Meesad, "An application of data mining and machine learning for weather forecasting," in *Proceedings of the International Conference on Engineering Research and Applications*, 2017, pp. 169–178.

[4] S. Navarro-Tec, M. G. Orozco-del Castillo, J. C. Valdiviezo-Navarro, D. R. Ordaz-Bencomo, M. R. Moreno-Sabido, and C. Bermejo-Sabbagh, "Análisis del crecimiento urbano y su relación con el incremento de temperaturas en la ciudad de mérida utilizando imágenes satelitales." *Res. Comput. Sci.*, vol. 147, no. 7, pp. 285–294, 2018.

[5] A. P. Rodrigues, R. Fernandes, and P. Vijaya, "A study on the evaluation of different regressors in weather prediction," in *2022 International Conference on Artificial Intelligence and Data Engineering (AIDE)*. IEEE, 2022, pp. 13–18.

[6] P. Mahajan, C. Nawale, S. Kini, and K. Shinde, "Weather forecasting using neural network," *International journal of engineering research and technology*, vol. 5, 2018.

[7] M. Sharma, L. Mathew, and S. Chatterji, "Weather forecasting using soft computing and statistical techniques," *International Journal of Advanced Research in Electrical, Electronics and Instrumentation Energy*, vol. 3, pp. 11 285–11 290, 2014.

[8] S. B. Nuthalapati and A. Nuthalapati, "Accurate weather forecasting with dominant gradient boosting using machine learning," *International Journal of Science and Research Archive*, 2024.

[9] M. Raghuwanshi, Y. Katre, A. Sahu, D. Sharma, A. Udapure, and C. Lonarkar, "Weather prediction with machine learning," *International Journal of Innovative Science and Research Technology (IJISRT)*, 2024.

[10] E. A. Suárez-Gallareta, J. J. Hernández-Gómez, G. Cetzal-Balam, M. G. Orozco-del Castillo, M. R. Moreno-Sabido, and R. A. Silva-Aguilera, "Sistema hıbrido basado en redes neuronales artificiales y descomposición modal empırica para la evaluación de la interrelación entre la irradiancia solar total y el calentamiento global," *Res. Comput. Sci*, vol. 147, no. 5, pp. 319–332, 2018.

[11] S. Bhatia, B. B. Naib, N. Goel, B. Agarwal, B. Mallikarjun, and M. Arvindhan, "A new era in weather forecasting: Harnessing the potential of machine learning," in *2024 IEEE International Conference on Computing, Power and Communication Technologies (IC2PCT)*, 2024, pp. 1355–1361.

[12] P. Das, P. Parmar, S. Sahoo, A. Saluja, and S. Pande, "An intelligent regression approach for weather forecasting system using machine learning," in *2024 1st International Conference on Cognitive, Green and Ubiquitous Computing (IC-CGU)*, 2024, pp. 1–6.

[13] P. R. P. G. Hewage, M. Trovati, E. Pereira, and A. Behera, "Deep learning-based effective fine-grained weather forecasting model," *Pattern Analysis and Applications*, vol. 24, pp. 343–366, 2020.

[14] N. Ng, H. Gopalan, V. Raghavan, and C. Ooi, "Day-ahead forecasting for the tropics with numerical weather prediction and machine learning," in *2022 17th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, 2022, pp. 125–130.

[15] P. Nazareth, A. K. Konnur, A. M. Chavan, and K. M. D. Shetty, "Weather forecasting using machine learning approach," in *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 2024, pp. 1–7.

[16] G. Chen, S. Liu, and F. Jiang, "Daily weather forecasting based on deep learning model: A case study of shenzhen city, china," *Atmosphere*, 2022.

[17] M. Ndiaye and M. Mboup, "Empirical mode decomposition with envelope extraction and lstm for univariate time series forecasting," in *2023 International Conference on Machine Learning and Applications (ICMLA)*, 2023, pp. 459–464.

[18] M. G. Orozco-del Castillo, E. C. Orozco-del Castillo, E. Brito-Borges, C. Bermejo-Sabbagh, and N. Cuevas-Cuevas, "An artificial neural network for depression screening and questionnaire refinement in undergraduate students," in *International Congress of Telematics and Computing*. Springer, 2021, pp. 1–13.

[19] S. Zhou, B. J. Bethel, W. Sun, Y. Zhao, W. Xie, and C. Dong, "Improving significant wave height forecasts using a joint empirical mode decomposition–long short-term memory network," *Journal of Marine Science and Engineering*, 2021.

[20] V. Kumar, Shaktibala, and S. Khan, "Importance of weather prediction for sustainable agriculture in bihar, india," *Archives of Agriculture and Environmental Science*, vol. 2, no. 2, pp. 105–108, 2017. [Online]. Available: https://journals.aesacademy.org/index.php/aaes/article/view/02-02-08

[21] P. Stana and M. Vladu, "The importance of the accuracy of public weather forecasts for farmers in the context of climate change in the management of agro-technical decisions," *Annals of the University of Craiova - Agriculture Montanology Cadastre Series*, 2020.

[22] Y. Gong, Y. Zhang, F. Wang, and C.-H. Lee, "Deep learning for weather forecasting: A cnn-lstm hybrid model for predicting historical temperature data," in *Proceedings of the 5th International Conference on Signal Processing and Machine Learning*, 2024. [Online]. Available: https://www.ewadirect.com/proceedings/ace/article/view/17393

[23] J. Xu, Z. Wang, X. Li, Z. Li, and Z. Li, "Prediction of daily climate using long short-term memory (lstm) model," *International Journal of Innovative Science and Research Technology (IJISRT)*, 2024.

[24] W. A. Respaty, C. Hong, N. K. Putra, F. Kurniadi, and Riccosan, "Weather prediction in jakarta: An analysis of climate data and regional influences using lstm and gru," in *2023 International Conference on Data Science and Its Applications (ICoDSA)*, 2023, pp. 408–413.

[25] Y. Arifin, I. Sonata, Maryani, and E. P. Gunawan, "Weather prediction in agriculture yields with transformer model," *Procedia Computer Science*, 2024.

Journal of Artificial Intelligence and Computing Applications (2024) 2(2): 48–59

Journal
of
Artificial Intelligence
and
Computing Applications

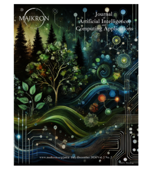

Research Article

An Explainable Clustering-Based Approach for Cyber Situational Awareness on Masquerade Attacks Detection

Nelva N. Almanza-Ortega[1], **Joaquin Perez-Ortega**[1], **Sergio M. Martinez-Monterrubio**[2], **and Juan A. Recio-Garcia**[3,*]

[1]National Technological Institute of Mexico, México
[2]International University of La Rioja, Spain
[3]University Complutense of Madrid, Spain

ABSTRACT

Masquerade attacks pose a significant challenge in cybersecurity, as intruders mimic legitimate user behavior to evade detection. In dynamic, data-intensive environments, traditional intrusion detection systems often struggle to provide both timely and interpretable results, limiting their usefulness for effective Cyber Situational Awareness (CSA). This article presents a clustering-based approach for detecting masquerade attacks using OK-Means—a variant of K-Means optimized for faster convergence—combined with a nearest neighbor classifier and noise reduction techniques. The proposed Intrusion Detection System (IDS) reduces computational overhead while enhancing explainability, leading to more reliable and transparent Cyber Threat Intelligence (CTI) decisions.

Keywords: cyber situational awareness (CSA), masquerade attack detection, explainable machine learning

1. Introduction

In the last three decades, humanity has witnessed a revolution in Information Technology (IT), that has liberated the need for a digital transformation in all economic sectors in cyberspace. It has become the mainstay of growth and prosperity in the world economy. As a result, different nations have been forced to implement legislation and regulations to protect their cyber assets and digital markets. Nowadays, it is mandatory to become aware of the cyber situation for the proper performance of command and control tasks. Cyber situational awareness (CSA) contributes to mission-centric Cyber Threat Intelligence (CTI) providing support for informed decisions required to maintain a safe and secure IT environment. According to [1], "situation awareness is the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and projection of their status." However, the dramatic increase in IT complexity and the large amount of raw data generated by such systems complicates the correct implementation of this model.

Current approaches for CTI from the point of view of CSA are based on Machine Learning (ML) techniques, like decision trees, support vector machines, neural networks, etc. [2]. These methods are based on the generation of a predictive model able to identify and classify threats. However, they have two significant drawbacks. Firstly, as CSA applies explicitly to dynamic environments, the ML predictive models must be re-trained as the complexity of IT systems and their

E-mail address: jareciog@fdi.ucm.es
https://doi.org/10.5281/zenodo.15098302

associated risks evolve. Secondly, most of these ML algorithms are not suitable for big-data scenarios where a continuous and massive data stream provides the information to be analyzed.

The problem of applying static predictive models can be addressed through the use of lazy-learning ML methods such a the k nearest-neighbors (k-NN). This method does not require a precomputed model because all the computation is deferred until classification. Moreover, k-NN has the advantage that classification is approximated locally, and therefore, only a subset of the data should be required in order to obtain the prediction. This makes, k-NN the right approach for processing big data streams if we can select the local subset of data required to classify a potential threat. However, as k-NN is an instance-based approach, it will obtain a poor performance if we compare a new input sample to the whole large-scale datasets of collected information. In this sense, its combination with clustering techniques stands as a possible alternative to improve its applicability in big data scenarios at the expense of lowering the performance of the classification. The primary assumption is that the k-nearest neighbors of a threat lie in the same cluster. Thus, the k-NN search can be efficiently performed in two steps: (1) reaching the nearest clusters; and (2) finding the k-nearest neighbors within the selected clusters. In a big data scenario, this would eventually save a massive amount of instance comparisons.

Typically, clustering methods such as K-Means take a considerable time to compute clusters. This is not a problem in static environments. However, in the heterogeneous and dynamic environments where CSA is required, it is an additional issue to be addressed: as the environment evolves, clusters must be recomputed to avoid the loss of performance. To address this problem, in this work we present an efficient approach for the detection of suspicious events based on the combination of k-NN and a novel clustering strategy, OK-Means [3], that decreases the cost of recomputing clusters as the environment evolves. Concretely, we will demonstrate the benefits of our approach for Masquerade Detection. The objective of such systems is to raise an alert when computer behavior differs to a certain extent from standard computer behavior, as profiled from a history of computer sessions. As the amount of information to be potentially logged is vast (user actions, files accesses, etc.), this kind of detection requires methods able to manage big data.

OK-Means is specifically useful in big data realms as it uses a criterion to balance the processing time and the solution quality when the number of instances is significant. Instead of trying to improve the initialization or classification steps as the majority of the known strategies aimed to improve the performance of k-mean, this algorithm applies in the convergence step. This way, it achieves a decrease in computing time of about a factor 4/100, yielding solutions whose quality reduces by less than 2%.

As we will present in the results of this paper, our experimental evaluations show that the combination of OK-Means and k-NN decreases the computing time without a significant quality loss when applied to the detection of treads in a CSA scenario. Moreover, the second major contribution of this paper is the proposal of explanatory strategies to improve the CSA by allowing the cybersecurity analyst to understand the outcomes of the intrusion detection system. Explainable AI (XAI) is nowadays a significant research challenge and is driven by the evidence that many AI applications lack trust on behalf of their users. The running hypothesis is that by building more transparent, interpretable, or explainable systems, users will be better equipped to understand and therefore trust the intelligent system [4]. This hypothesis can be directly extrapolated to CSA, where black-box machine learning algorithms such as neural networks or Bayesian networks are commonly applied for intrusion detection. Here, the CTI is limited by the lack of transparency of the CSA system as the cybersecurity analyst is not able to obtain a clear perception of the causes that led to an intrusion alert. Therefore, we have chosen k-NN and clustering not only because they are able to deal with big data, but also because they are white-box methods that can be introspected to provide explanations about the causes of a potential threat.

The paper runs as follows. Section 2 presents the related work. Section 3 describes our approach based on the application of OK-Means to large datasets and its combination with a k-NN classifier. The experimental evaluation and associated results is presented and discussed in Section 4 and 5. Section 6 presents visual explanatory strategies and finally, Section 7 concludes the paper.

2. Background

The Computer Emergency Response Team (CERT) defines malicious insider as "an employee, contractor, or business partner who has or had authorized access to an organization's network, system or data, and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information system" [6].

The importance of the above definition is that it delimits the term insider to a particular context since if only the definition of this concept is taken, the notion of belonging to a group could lead to infer the existence of a physical limit. However, with the advance of computer systems and telecommunications added to the evolution of organizations, this type of limit in practice is diffuse and difficult to identify. Therefore, it is no longer enough that the person belongs to the organization, but that the person has the authorization to interact with the organization's systems. In this sense, Bishop [7] has proposed an alternative definition that allows having a greater detail of the previous concept when considering the following aspects: (1) There is an entity (i.e., a person) that, by its level of trust has the power to violate one or more rules of a given security

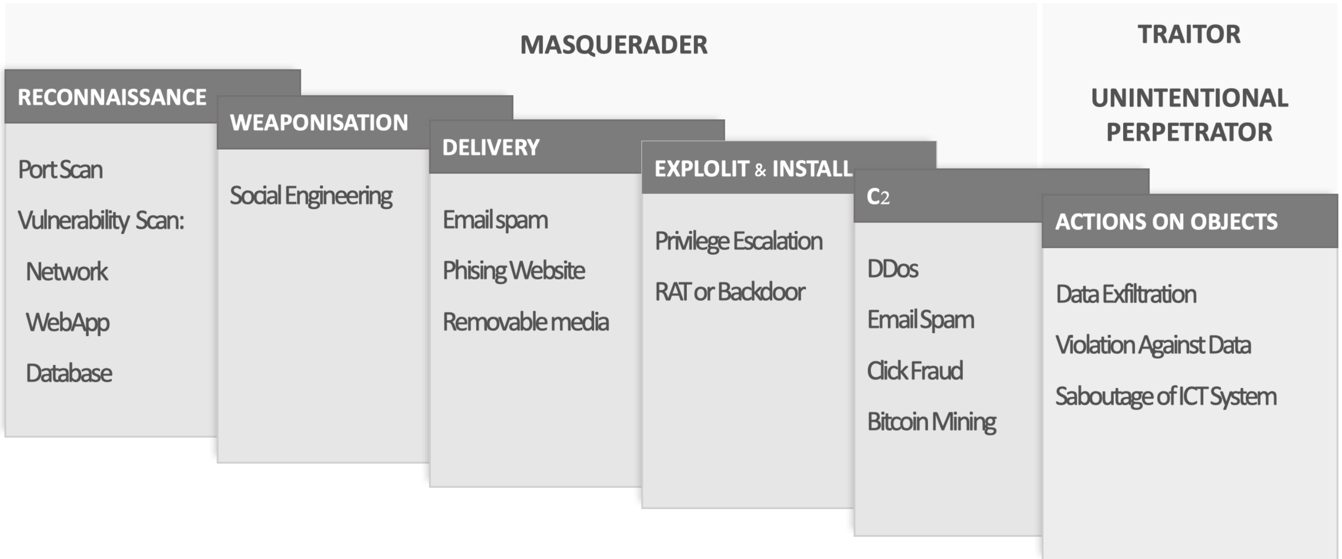

Figure 1. Taxonomy of the Insider Types and Specific Insider Threats [5].

policy. (2) The violation of the security policy is carried out using legitimate accesses. (3) A violation of the access control policy occurs when unauthorized access is obtained.

Figure 1 illustrates the most relevant types of insider and relevant threats [5]. This way, the most commonly seen insider threats are 1) data exfiltration, 2) violations against data integrity or availability, and 3) sabotage of IT systems. Technically, traitors and unintentional perpetrators can fulfill these threats straightway. A masquerader may pose the same threats via an intrusion campaign that consists of social engineering, eavesdropping, packet sniffing, malware delivery, installation, etc.

CSA is used to safeguard sensitive data, sustain fundamental operations, and protect national infrastructure from both malicious insiders and external attackers. It is a multifaceted and well-studied phenomenon, which can be looked upon from several different perspectives [?]. The need for situation awareness is essential to understand the organization's environment and accurately to predict and respond to potential problems that might occur. CSA involves three key areas: computing and network components, threat information, and mission dependencies [8]. Achieving this level of situation awareness requires an investment in data collection, data management, and analysis to maintain an ongoing picture of how the computer systems, networks, and users are operating in an organization. In threat awareness, the crucial facts are to identify and track internal incidents and suspicious behavior. Understanding these critical dependencies will anticipating and avoiding situations.

The process of situational awareness can be viewed as a three-phase process [9]: (1) Situation perception. Perception gains awareness about the status, attributes, and dynamics of relevant elements within the enterprise networks. (2) Situation comprehension. Comprehension of the situation encompasses how analysts combine, correlate, and interpret information. (3) Situation projection. Projection of the situation into the near future encompasses the ability to make predictions based on the knowledge acquired through perception and comprehension.

Here, it is essential to note that the comprehension phase is directly influenced by the capability of the IT systems to explain their internal processes. Both comprehension and interpretability are key aspects of Explainable AI. The goal of Explainable Artificial Intelligence (XAI) is "to create a suite of new or modified machine learning techniques that produce explainable models that, when combined with effective explanation techniques, enable end-users to understand, appropriately trust, and effectively manage the emerging generation of Artificial Intelligence (AI) systems" [10]. However, few works relate CSA and XAI. For example, Marino et al. [11] present an adversarial approach to generate explanations for incorrect classifications made by data-driven IDS, and Fujii et al. [12] propose an explainable AI intrusion detection system through the combination of deep tensor and knowledge graph.

3. Method

The global structure of the proposed method for masquerade detection is depicted in Figure 2. It consists of the following components:

Cyber sensing. The perception step is performed by User Activity Monitoring (UAM) sensors installed in every machine of the IT system. In this case, sensors focus on the elements required for the masquerade detection: commands, file ac-

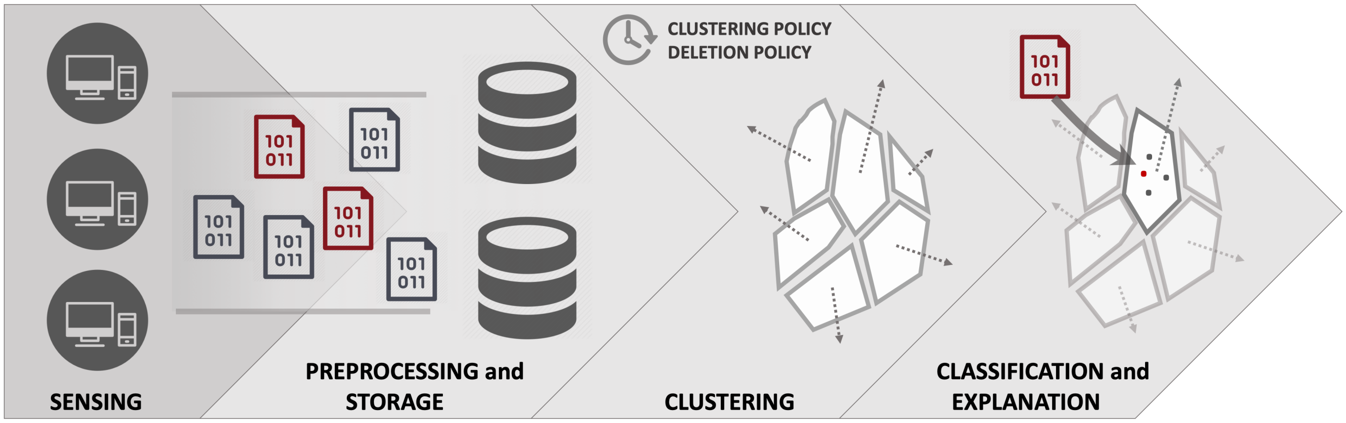

Figure 2. Global structure of the proposed method for masquerade treats detection.

cesses, etc.

Preprocessing and storage. The preprocessing and storage of the collected data must be implemented using the most suitable big-data solution such as Hadoop or Spark. See [13] for a comprehensive comparison.

Clustering policy. The clustering of the collected data must be executed periodically, according to a defined policy, in order to update the clusters used to classify potential threats. These policies can be based on the amount of new data collected or the increase of the error ratio of the classification.

Data deletion policy. As masquerade attacks evolve, it is necessary to define a forget policy in order to delete old data that does not contribute to the correct classification of the masquerade attacks. This policy must define periodical deletions of unuseful data based on noise reduction algorithms such as Repeated Edited Nearest Neighbour (RENN) [14], Blame-Based Noise Reduction (BBNR) [15] or Conservative Redundancy Reduction (CRR) [16].

Classification. This component performs the classification of a potential threat using a two steps process: (1) reaching the nearest clusters using their prototype; and (2) finding the k-nearest neighbors within the cluster.

Explanation. Explainability is one of the significant challenges in machine learning nowadays. In our case, the attack alerts raised by the system must be justified to improve user acceptance and provide appropriated counter-measures. Therefore, our method includes several visual explanation strategies to let the expert introspect the potential attack and understand its nature and severity.

The critical process within this method is the clustering of the dataset. As it is a very time-consuming task, we propose the OK-Means algorithm in order to decrease the classification time. The main schema consists of using the cluster's prototypes to filter and select those clusters containing the most similar records to the potential threats. Then, only the records from the selected clusters are used for the classification. However, this process has an obvious impact on classification performance. Figure 3 shows the expected performance of this approach according to our previous results [17]. Although the processing time follows a linear progression as the number of selected clusters grows, the accuracy may not increase proportionally, slowing down when the number of records is too high. This way, it is necessary to find a trade-off between the processing time and the prediction performance in order to be able to process as many masquerade attacks as possible with an acceptable success rate.

3.1 The OK-Means clustering algorithm

The clustering problem consists of dividing a set of n objects into two or more non-empty subsets or clusters, such that the objects in the same cluster have similar attribute values and have attribute values different from those of the objects in other clusters [18].

The OK-Means clustering algorithm is an improvement of K-Means [19], which by using a new stop or convergence criterion, allows reducing the processing time considerably at the expense of a small reduction of the solution quality. The new stop criterion consists in halting OK-Means when the number of objects that change cluster membership in an iteration is smaller than a threshold U. The value of U expresses an optimal compromise between the computational effort and the solution quality, and it is calculated by applying the Pareto Principle [3, 20].

According to [21, 22] the type of problems that are solved by K-Means belong to the NP-hard problems for $k \geq 2$ or $d \geq 2$. The complexity of K-Means is $O(nkdr)$ [21, 22], while the complexity of OK-Means is $O(nkdr\alpha)$, where r denotes the number of iterations and α is the ratio of the number of iterations of OK-Means

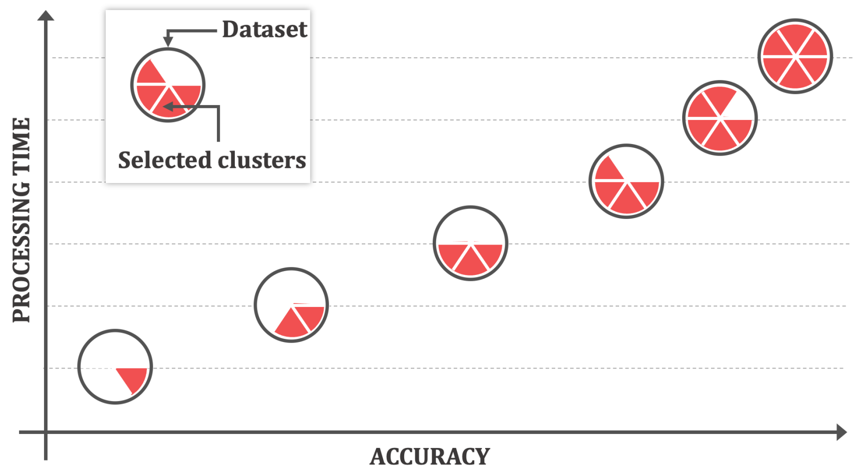

Figure 3. Expected time/performance as the number of selected clusters raises.

and the number of iterations of K-Means. In the experiments carried out, an average value of $\alpha = 0.0389$ was obtained, which shows that OK-Means significantly reduces the computational complexity of K-Means [3, 20].

Let $N = \{x_1, ..., x_n\}$ denote the set of n points to be grouped by a closeness criterion, where $x_i \in \Re^d$ for $i = 1, ..., n$, and $d \geq 1$ is the number of dimensions (the objects' attributes). Further, let $k \geq 2$ be an integer and $K = \{1, ..., k\}$. For a k-partition $P = \{G(1), ..., G(k)\}$ of N, denote μ_j the centroid of group (cluster) $G(j)$, for $j \in K$, and let $M = \{\mu_1, ..., \mu_k\}$.

Thus, the clustering problem can be formulated as a constrained optimization one (see, for instance, [23]):

$$\boldsymbol{P}\text{: minimize } z(W, M) = \sum_{i=1}^{n} \sum_{j=1}^{k} w_{ij} d(x_i, \mu_j) \qquad (1)$$

$$\text{subject to } \quad \sum_{j=1}^{k} w_{ij} = 1, \text{ for } i = 1, \ldots, n,$$

$$w_{ij} = 0 \text{ or } 1, \text{ for } i = 1, \ldots, n, \text{ and } j = 1, \ldots, k,$$

where $w_{ij} = 1 \iff$ point x_i belongs to cluster $G(j)$, and $d(x_i, \mu_j)$ denotes the Euclidean distance between x_i and μ_j, for $i = 1, \ldots, n$, and $j = 1, \ldots, k$.

Algorithm 1 shows the pseudocode of the OK-Means algorithm [3, 20].

Algorithm 1 OK-Means

Step 1 Initialization. Produce points $\mu_1, \ldots, \mu_k$, as a random subset of N.

Step 2 Classification. For all $x \in N$ and $j \in K$, compute the Euclidean distance between points x and μ_j, namely, $d(x, \mu_j)$. Then, point (object) $x \in N$ is assigned to a cluster $G(\ddot{j})$ if $d(x, \mu_{\ddot{j}}) \leq d(x, \mu_j)$, for $\ddot{j}, j \in K$.

Step 3 Centroids. Determine the centroid μ_j of cluster $G(j)$, for $j \in K$.

Step 4 Convergence. If $\gamma_r \leq U$ stop the algorithm, otherwise perform another iteration starting from **Step 2**.

In **Step 1 (Initialization)**, the initial centroids are selected for generating the k clusters; in **Step 2 (Classification)**, the membership of an object is determined according to the smallest Euclidean object-centroid distance; in **Step 3 (Centroids)** the centroid of each cluster is calculated; and finally in **Step 4 (Convergence)**, the algorithm is halted when the number of objects that change cluster membership is smaller than the threshold value U.

The OK-Means algorithm was evaluated using a set of experiments with large-size data instances like those used in Big Data. For the experiments, instances of synthetic data were generated, and instances of real data were obtained from the UCI repository [24]. The experimental results show that the OK-Means algorithm reduces the processing time by a factor of 4% with a reduction of less than 2% of the quality of the solution.

4. Experimental evaluation

In order to prove the validity of our proposal, we have conducted an experimental evaluation for the Masquerade Detection scenario. We have chosen the Windows-Users and Intruder simulations Logs Dataset (WUIL), created by [25] that, instead of focusing on users' actions, is based on the objects that are subjects of those actions. Therefore, this dataset uses the concept of *locality*, the tendency of programs to cluster references to memory. Authors of the WUIL dataset define spatial locality as the property of the user to access files that are close to each other, and temporal locality as the property of the user to access the same file in the near future. As a practical implementation of these ideas, the WUIL dataset abstracts low-level events into a set of 16 temporal, spatial, and directional locality features, showing better performance than a purely action-based approach. This dataset has been collected through a User Activity Monitoring (UAM) sensor installed in several computers running different flavors of Windows OS. This sensing software gathers information for the date, the time, the file access path, and another kind of information.

Afterward, the information collected by the software sensors is preprocessed in order to extract the locality features. These features, as described in [26], are:

Spatial Locality Features are based on the idea that, while working, a user may access files that are close to each other. User events are recorded, including the object and timestamp of the interaction, and segmented into fixed time windows. Then, the file path of the object is used to compute several event distances, summarized in four spatial locality features.

The rationale behind these features is that event distances tend to be short for a legitimate user working on specific tasks, thus visiting objects close one another. An intruder, by contrast, may perform hops between far separated objects, while moving around looking for files of interest.

Temporal Locality Features are based on the intuition that while working, a regular user will frequently access the same files within a short period of time, whereas an intruder will traverse the file system looking for vulnerabilities. Authors define several features taking into account the user's access frequency and the elapsed time between two consecutive accesses to a given file.

Direction Features attempt at capturing where a user is heading at while interacting with the file system. Authors claim that an ordinary user is expected to browse over the file system following a prevailing direction and that the masquerader will have a strange direction pattern. This way, the WUIL dataset includes features describing the user is traveling between file system objects.

The WUIL dataset contains 54,649 instances where 52,884 are legitimate actions, and 1,765 are intruder attacks. The main goal of the evaluation is to measure the performance of our proposed methodology, simulating several data-demanding scenarios. Therefore, we have conducted a cross-validation evaluation using a progressive subsampling of the dataset but keeping this stratification ratio (96.77% legitimate, 3.33% attack). Thus, we have conducted 10-fold evaluations using 20%, 40%, 60%, 80%, and 100% of the instances in the dataset. For every experimentation, we have measured the performance of the system, setting-up several configuration values. Firstly, we have modified the number of clusters calculated by the OK-Means algorithm (denoted as C), and the number of selected clusters (according to the similarity to the prototype) used to filter the most similar records (sC). Finally, we have also tested several values for the k parameter of the nearest neighbor algorithm. The outcome of the classification is calculated through a majority-voting strategy.

As a result, for every subsampling of the dataset we have conducted 15 different evaluations configuring $k = 1, 3, 5$, $C = 4, 8$ and $sC = 1, 2, 4$. The applied evaluation metrics are time improvement (in order to obtain comparable results, all the experimental evaluations were run in the same computer under identical execution conditions) and recall for class "Attack" as explained next.

5. Results

Before analyzing the experimental evaluation, we must notice that the dataset is very unbalanced as 96.77% of the collected instances belong to the class "Legitimate". This is a coherent ratio as "Attacks" are strange events. However, it presents a challenge for the correct detection of attack attempts. If we analyze the accuracy metrics and type of errors for our classification goal, we find that precision is not relevant. Actually, a dummy classifier that always returns class "Legitimate" for any instance would obtain a 96.77% of precision value due to the unbalanced class ratio. Therefore, we must focus on recall as it denotes false negatives. In the case of the "Legitimate" class, recall measures those instances that being classified as "Attack" their actual class is "Legitimate". Here, we can consider it as a false alarm that must be analyzed by the system administrator, but it does not represent any risk for the system. Contrarily, false-negative classification for class "Attack" does represent a serious risk for the organization as these actions would not be identified by the masquerade detection system and reported as legitimate actions. This way, the results presented next will focus on the recall values obtained for class "Attack".

5.1 Baseline performance

Before evaluating the performance of the clustering approach, it is necessary to define the baseline to compare with. In this case, we have tested several classification algorithms over the whole dataset without any clustering optimization. Results are shown in Figure 4. Firstly we have tested two state-of-the-art classification methods such as Random Forest and Multilayer Perceptron, achieving a 0.67 and 0.64 recall value for class "Attack" respectively. These values could be assumed as the baseline performance for our detection system. Then, we have evaluated the k-NN classifier using majority voting and different values for the k parameter, obtaining a lower performance [0.68 - 0.58], but demonstrating that the 1-NN is the best option for this domain. Next, assuming that our goal is to decrease the false-negative rate, we have changed the majority voting strategy to

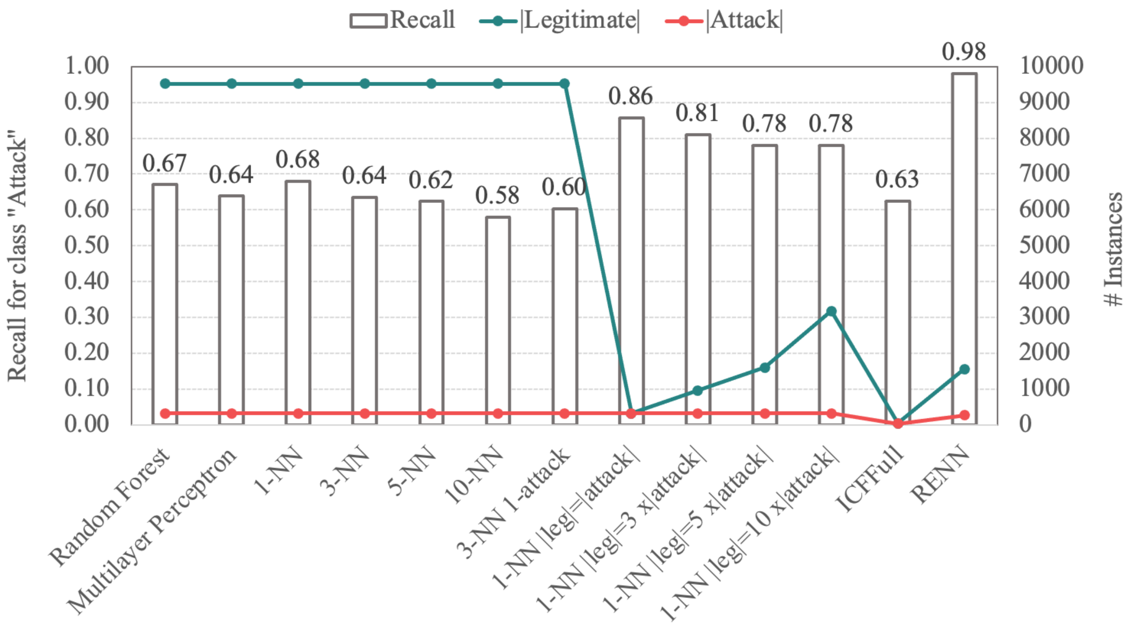

Figure 4. Baseline results using different classification strategies.

classify any instance as an attack if any of the three most similar neighbors is an attack. However, results were also disappointing, with a 0.6 recall value. The following strategies being evaluated consisted of applying an undersampling method for the predominant class "Legitimate". Therefore we randomly removed instances of this class according to the number of "Attack" instances, concretely, 1, 3, 5, and 10 times. Here, we noticed an improvement in the recall values, up to 0.86. From the observation that the removal of redundant instances from the dataset improved the performance of the system, we finally executed noise removal strategies, where RENN obtained a remarkable 0.98 recall value. RENN undersampling of the majority class is done by removing points whose class label differs from a majority of its k nearest neighbors. Removal is applied successively until it can remove no further points.

From this baseline evaluation, we can conclude that a noise removal method, concretely RENN, is a key component of the masquerade detection system in order to raise the performance. By combining this method with the clustering of the dataset, we will be able to obtain not only a very competitive classification but also a good time performance.

Next, we will analyze the impact of the clustering algorithm in the execution times.

5.2 Analysis of the execution time

The five subsamples of the dataset (20%, 40%, 60%, 80%, and 100%) were tested with 15 different configurations for the OK-Means and k-NN algorithms. Table 1 shows the time values achieved when executing 10-fold validation for every dataset subsample. As expected, the execution time decreases as the subset of clusters selected to retrieve instances (sC) decreases. Additionally, if the number of clusters (C) grows, each cluster has fewer instances, and therefore, the execution time also decreases.

From these results, we can observe that we could decrease execution time up to 85% approximately by using 8 clusters and selecting the most similar one (according to the prototype) to the query instance,

Obviously, the use of a subset of clusters in order to decrease execution time will make an impact on the classification performance, as explained next.

5.3 Performance analysis

From our preliminary analysis to obtain a baseline to compare with, we discovered the remarkable impact of the RENN algorithm regarding the recall. The next step consisted of combining this algorithm with the clustering strategy. Here, there are two possibilities. The first option consists of executing the RENN algorithm over the whole dataset and then perform the OK-Means clustering. Its alternative is to cluster first and then apply RENN to every clustering. Results are shown in Figure 5 (left). Surprisingly, the performance obtained when applying REEN to every cluster (labeled as "RENN intracluster" in Figure 5) was similar to the results obtained when this algorithm was not executed, and the complete dataset was clustered. This was the first indicator of the impact of the clustering quality in the performance results. As we will explain in Section 6, OK-Means is able to split legitimate and attack instances very efficiently. Therefore, the impact of the RENN algorithm in every cluster is minimized because noisy instances that may led to the miss-classification of the query are assigned to a different cluster. Although

Table 1. Execution time (in seconds) for different subsamples of the WUIL dataset under varying configurations of OK-Means (denoted as OK-NN when combined with k-NN). Results are shown for $k = 1$, two values of clusters ($C = 4$ and $C = 8$), and different numbers of selected clusters ($sC = 1, 2, 4$). Line 1 reports baseline k-NN execution times without clustering, and lines 2–5 show the percentage of time improvement (ΔT) relative to this baseline.

				Subsamples	20%	40%	60%	80%	100%
Line	Algorithm	k	C	sC	Time (seg.)	Time (seg.)	Time (seg.)	Time (seg.)	Time (seg.)
1	K-NN	1			159.02	350.14	548.94	746.17	941.87
					Δ T (%)	Δ T (%)	Δ T (%)	Δ T (%)	Δ T (%)
2	OK-NN	1	4	1	68.69	66.52	68.97	69.00	67.84
3	OK-NN	1	4	2	32.93	34.92	33.55	26.95	28.63
4	OK-NN	1	8	1	84.46	86.54	85.59	86.24	84.32
5	OK-NN	1	8	4	40.50	41.81	43.35	44.63	36.18

globally, it is a positive feature, if we join the instances of these clusters as the sC parameter grows, we will be recovering the original noisy classifications.

On the other hand, the results obtained when applying RENN to the complete dataset and later performing clustering were very satisfactory, as shown in Figure 5 (RENN complete series).

As there are several combination schemes for noise removal and clustering algorithms, we have also collected their execution times. Figure 5 (right) shows the results. As expected, the complete dataset without noise removal obtains the lowest values. Next, the application of RENN to every cluster obtains average times because the number of instances in each cluster is much lower than the complete dataset. Finally, applying RENN to the whole dataset and later performing clustering + classification obtains the worse execution times.

Once we can conclude the application of RENN to the complete dataset and its later clustering leads to the best recall values, we analyzed the stability of our approach for the different subsamples of the dataset. As Figure 6 reports, results were homogeneous when ranging from the 40% to the 100% percent of the instances. Performance for the 20% of the dataset was a bit lower, indicating that it may be an excessive subsampling. However, as a general result, we can conclude that our method can detect up to 99% of the masquerade attacks with a remarkable time performance.

6. Visual explanatory strategies

CSA consists of making informed decisions based on the comprehension of the environment and the meaning of an event [1]. Therefore, we propose the use of two explanatory strategies to let the security expert analyze both dimensions of the potential masquerade attack. These are the clustering analyzer and the attack introspection visual tools. Both are explained next.

6.1 Clustering analyzer

The clustering analyzer allows the cybersecurity analysts to evaluate the environment of attack through the visualization of the level of hazard of the cluster where the potential threat was classified. This classification is performed by comparing the potential threat to the prototypes of every cluster. Then, this analysis tool shows the proportional size and the Legitimate/Attack ratio of every cluster graphically.

This visual explanation is exemplified in Figure 7. In this case, we have visualized the resulting clusters using the 20% subsampling (left) and the complete dataset (right). This tool uses a hierarchical tree-map where the size (height) of every cluster represents the total number of belonging instances. Then every cluster is divided into the "Legitimate" (blue) and "Attack" (red) areas that are also proportional to the number of instances. Finally, the visualization highlights the "Attack" subcluster, where the potential threat has been classified.

As we can observe with this example, this tool lets the cybersecurity analyst evaluate the potential hazard of the alert. For example, in the first case (Figure 7 left) using the 20% subsampling, there is a potential attack classified in Cluster 4 that shows a minimal number of attacks, indicating a high possibility of false alarm. On the other hand, Figure 7 right shows an attack being classified in a cluster that contains a large number of past attacks, denoting its potential risk.

Finally, both figures demonstrate the excellent performance of the OK-Means algorithm graphically, which globally generates clusters that have either a majority of legitimate or attack instances, especially as the dataset grows. If the performance of the clustering was deficient, clusters should follow the 97% "Legitimate" and 3% "Attack" proportion of the dataset.

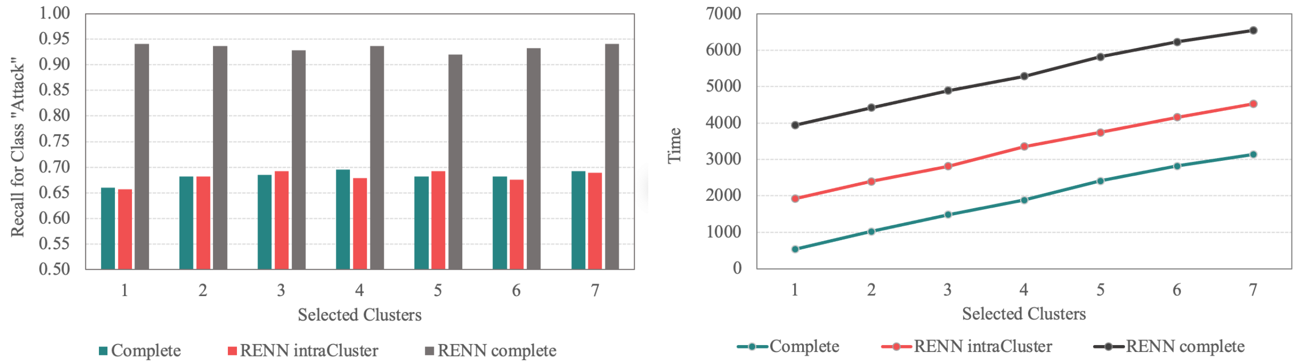

Figure 5. Recall (left) and processing time (right) using different noise removal approaches (dataset 20%).

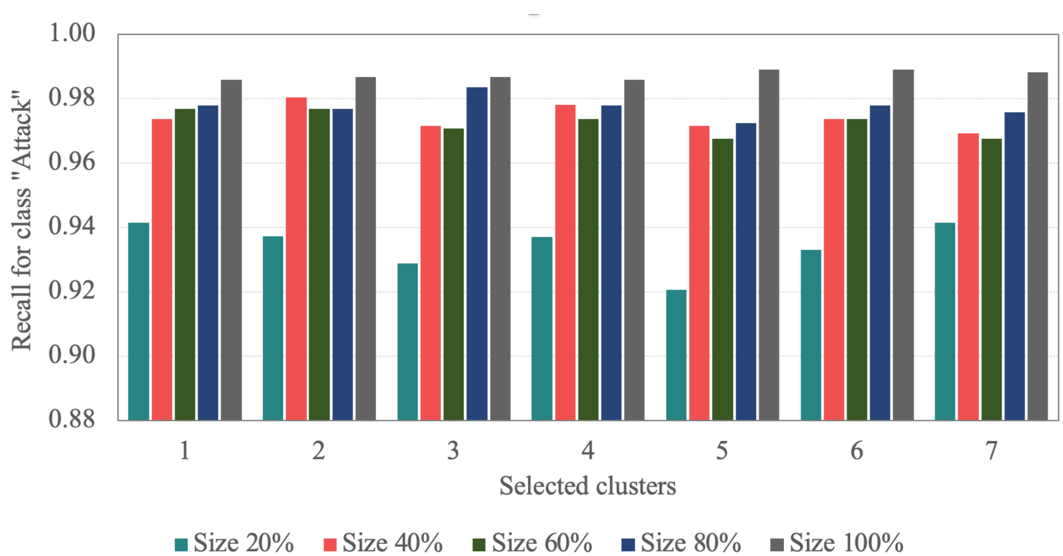

Figure 6. Recall when applying RENN to different subsamplings of the dataset

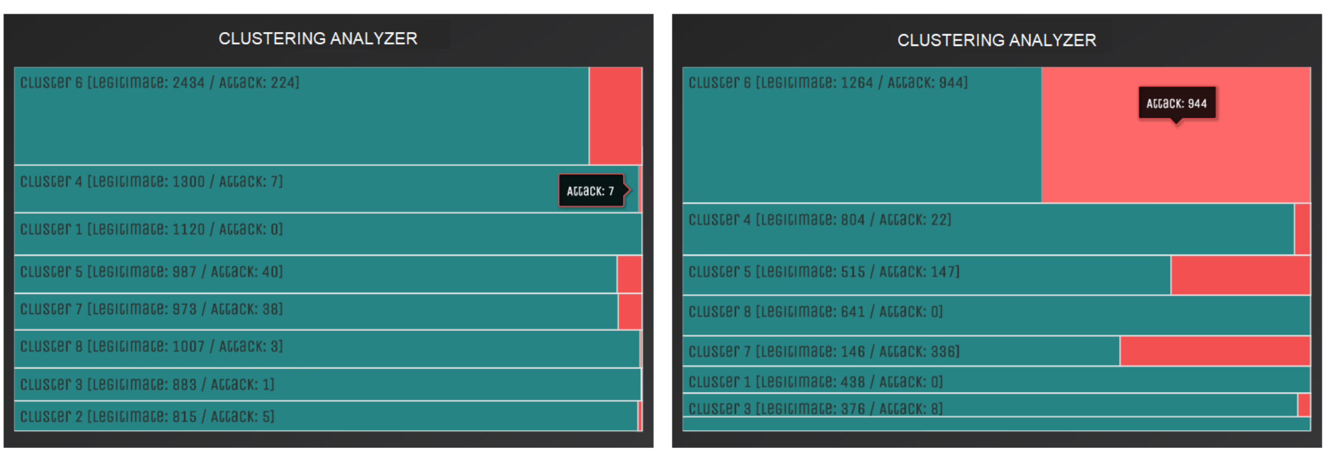

Figure 7. Clustering analyzer tool when visualizing the 20% subsampling (left) and the complete dataset (right).

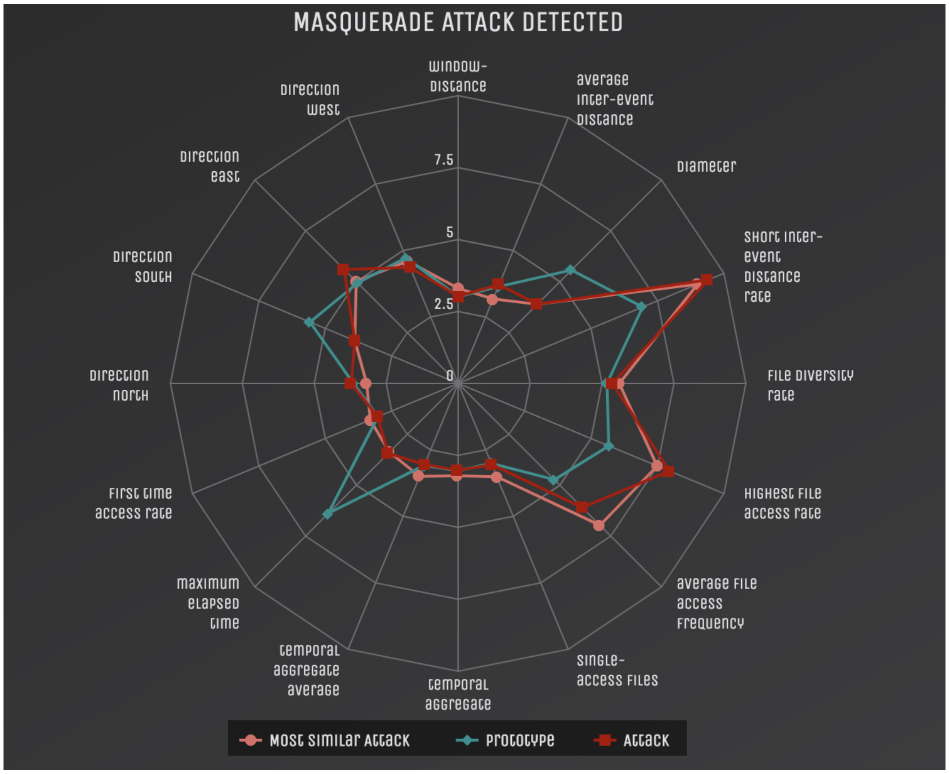

Figure 8. Screenshot of the attack introspection tool showing the features of the attack, the most similar attack that raised the alarm, and the cluster's prototype.

6.2 Attack introspection tool

The attack introspection tool complements the previous tool and allows the security analyst to understand the nature, meaning, and projection of a potential threat. In this case, the features of the attack are displayed graphically to compare them with the most similar attack that raised the alarm (note that we are applying 1-NN) and the cluster's prototype.

A screenshot of this tool is shown in Figure 8. In this case, we use a polar chart where every axis represents one instance's feature. As we are evaluating our approach with the WUIL dataset that uses 16 features to represent masquerade attacks, these are the features shown in the example.

We can observe that the prototype representing all instances in the cluster (blue line) contains average values. However, the potential attack (red line) and the most similar attack that raised the alert (orange) do have anomalous values for some features. In this case, the inter-event rate, the highest file access rate, and the average file access frequency are atypically high compared to the cluster average. These are clear indicators of an intrusion attack as the event, and file access values denote hops between far separated objects looking for files of interest. Additionally, the maximum elapsed time is also lower than average, denoting that the intruder is traversing the file system looking for vulnerabilities.

7. Conclusions

In this work, we have presented an efficient approach for the detection of malicious threats based on the combination of k-NN and a novel clustering strategy, OK-Means [3], that decreases the cost of recomputing clusters as the environment evolves. Concretely, we have demonstrated the benefits of our approach for Masquerade Detection, where execution times decreased up to 68%, and intrusion detection performance is able to detect up to 99% of the masquerade attacks. Moreover, we propose visual explanatory strategies to increase the cybersecurity analyst acceptance of the alerts raised by the IDS.

Data Availability

The dataset used to evaluate the system proposed in this paper is the Windows-Users and Intruder simulations Logs Dataset (WUIL), created by [25]. Availability can be inquired directly with the authors.

CRediT authorship contribution statement

Nelva N. Almanza-Ortega: Conceptualization, Methodology, Software, Validation, Formal analysis, Data curation, Visualization, Writing – original draft.

Joaquin Perez-Ortega: Conceptualization, Methodology, Software, Investigation, Formal analysis, Writing – original draft. **Sergio M. Martinez-Monterrubio:** Conceptualization, Writing – review & editing, Supervision. **Juan A. Recio-Garcia:** Conceptualization, Writing – review & editing, Supervision, Project administration.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT in order to improve readability. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was funded by the Spanish Committee of Economy and Competitiveness [TIN2017-87330-R], the Spanish Committee of Science and Innovation [PID2020-114596RB-C21/AEI/10.13039/501100011033] and the BOSCH-UCM Chair on Artificial Intelligence applied to Internet of Things [20896].

References

[1] M. Endsley, "Endsley, m.r.: Toward a theory of situation awareness in dynamic systems. human factors journal 37(1), 32-64," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 37, pp. 32–64, 03 1995.

[2] M. Conti, A. Dehghantanha, and T. Dargahi, "Cyber threat intelligence : Challenges and opportunities," *CoRR*, vol. abs/1808.01162, 2018. [Online]. Available: http://arxiv.org/abs/1808.01162

[3] J. Pérez-Ortega, N. N. Almanza-Ortega, and D. Romero, "Balancing effort and benefit of k-means clustering algorithms in big data realms," *PloS one*, vol. 13, no. 9, p. e0201874, 2018.

[4] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artificial Intelligence*, vol. 267, pp. 1 – 38, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0004370218305988

[5] L. Liu, O. De Vel, Q. Han, J. Zhang, and Y. Xiang, "Detecting and preventing cyber insider threats: A survey," *IEEE Communications Surveys Tutorials*, vol. 20, no. 2, pp. 1397–1417, Secondquarter 2018.

[6] D. Cappelli, A. Moore, and R. Trzeciak, *The CERT Guide to Insider Threats: How to Prevent, Detect, and Respond to Information Technology Crimes (Theft, Sabotage, Fraud).* Addison-Wesley Professional, 2012.

[7] M. Bishop, L. Coles-Kemp, D. Gollmann, J. Hunker, and C. W. Probst, Eds., *Insider Threats: Strategies for Prevention, Mitigation, and Response, 22.08. - 26.08.2010*, ser. Dagstuhl Seminar Proceedings, vol. 10341. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2010. [Online]. Available: http://drops.dagstuhl.de/portals/10341/

[8] J. Bone, *Cognitive hack : the new battleground in cybersecurity.*, 2017.

[9] M. Albanese and S. Jajodia, *Formation of Awareness.* Cham: Springer International Publishing, 2014, pp. 47–62. [Online]. Available: https://doi.org/10.1007/978-3-319-11391-3_4

[10] D. Gunning, "Darpa's explainable artificial intelligence (xai) program," 03 2019, pp. ii–ii.

[11] D. L. Marino, C. S. Wickramasinghe, and M. Manic, "An adversarial approach for explainable AI in intrusion detection systems," *CoRR*, vol. abs/1811.11705, 2018. [Online]. Available: http://arxiv.org/abs/1811.11705

[12] M. Fuji, H. Morita, K. Goto, K. Maruhashi, H. Anai, and N. Igata, "Explainable ai through combination of deep tensor and knowledge graph," *Fujitsu Scientific and Technical Journal*, vol. 55, pp. 58–64, 01 2019.

[13] S. Garcia, S. Ramirez-Gallego, J. Luengo, J. Benitez, and F. Herrera, "Big data preprocessing: methods and prospects," *Big Data Analytics*, vol. 1, 12 2016.

[14] D. Hand and B. Batchelor, "Experiments on the edited condensed nearest neighbor rule," *Information Sciences*, vol. 14, no. 3, pp. 171 – 180, 1978. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0020025578900403

[15] B. Smyth and E. McKenna, "Competence models and the maintenance problem," *Computational Intelligence*, vol. 17, no. 2, pp. 235–249, 2001, cited By 51. [Online]. Available: https://www.scopus.com/inward/record.uri?eid=2-s2.0-0035329732&doi=10.1111%2f0824-7935.00142&partnerID=40&md5=76bf7fe08a06a358f4ddab97a4d78047

[16] S. J. Delany and P. Cunningham, "An analysis of case-base editing in a spam filtering system," in *Advances in Case-Based Reasoning*, P. Funk and P. A. González Calero, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 128–141.

[17] A. Fornells, J. A. Recio-García, B. Díaz-Agudo, E. Golobardes, and E. Fornells, "Integration of a methodology for cluster-based retrieval in jcolibri," in *Case-Based Reasoning Research and Development, 8th International Conference*

on Case-Based Reasoning, ICCBR 2009, Seattle, WA, USA, July 20-23, 2009, Proceedings, ser. Lecture Notes in Computer Science, L. McGinty and D. C. Wilson, Eds., vol. 5650. Springer, 2009, pp. 418–433. [Online]. Available: https://doi.org/10.1007/978-3-642-02998-1_30

[18] R. Xu and D. Wunsch, *Clustering*. John Wiley & Sons, 2008, vol. 10.

[19] R. Jancey, "Multidimensional group analysis," *Australian Journal of Botany*, vol. 14, no. 1, pp. 127–130, 1966.

[20] N. N. Almanza-Ortega, "Desarrollo de heurísticas para la mejora del algoritmo k-means en las fases de clasificación y convergencia," Ph.D. dissertation, Centro Nacional de Investigación y Desarrollo Tecnológico, 2018.

[21] D. Aloise, A. Deshpande, P. Hansen, and P. Popat, "Np-hardness of euclidean sum-of-squares clustering," *Machine learning*, vol. 75, no. 2, pp. 245–248, 2009.

[22] M. Mahajan, P. Nimbhorkar, and K. Varadarajan, "The planar k-means problem is np-hard," *Theoretical Computer Science*, vol. 442, pp. 13–21, 2012.

[23] S. Z. Selim and M. A. Ismail, "K-means-type algorithms: A generalized convergence theorem and characterization of local optimality," *IEEE Transactions on pattern analysis and machine intelligence*, no. 1, pp. 81–87, 1984.

[24] M. Lichman, "Uci machine learning repository; 2013 [cited 24 aug 2017]. database: Repository of machine learning databases [internet]."

[25] J. B. Camiña, C. Hernandez-Gracidas, R. Monroy, and L. Trejo, "The windows-users and -intruder simulations logs dataset (wuil): An experimental framework for masquerade detection mechanisms," *Expert Systems with Applications*, vol. 41, no. 3, pp. 919 – 930, 2014, methods and Applications of Artificial and Computational Intelligence. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0957417413006349

[26] J. Camiña, R. Monroy, L. Trejo, and M. Medina-Pérez, "Temporal and spatial locality: An abstraction for masquerade detection," *IEEE Transactions on Information Forensics and Security*, vol. 11, pp. 2036–2051, 06 2016.