Low-Power Distributed Acoustic Sensor Network for Autonomous Wildlife Monitoring Using LoRa and AI for Digital Twin

Gonzalo de Miguel, Miguel Zaragoza-Esquerdo, Alberto Ivars-Palomares, Sandra Sendra, Jaime Lloret
Instituto de Investigación para la Gestión Integrada de Zonas Costeras, Universitat Politècnica de València
C/Paranimf, 1, 46730 Grao de Gandia, Spain
email: gdemig1@posgrado.upv.es, mizaes2@epsg.upv.es, aivapal@epsg.upv.es, sansenco@upv.es, jlloret@dcom.upv.es

***Abstract*—Biodiversity loss driven by climate change, habitat degradation, and anthropogenic pressures demands efficient wildlife monitoring solutions. Conventional methods are often costly, invasive, and limited in spatial or temporal coverage. Acoustic monitoring provides a non-intrusive alternative but faces challenges related to high data volumes, limited power availability, and restricted communication bandwidth in remote deployments. This paper presents a low-power distributed acoustic sensor network for autonomous wildlife monitoring, with emphasis on bird species. Each node combines an ESP32 microcontroller, a high-sensitivity digital microphone, and a Long Range (LoRa) transceiver to capture and transmit event-triggered audio. Real-time Fast Fourier Transform (FFT) analysis detects relevant acoustic activity, triggering Adaptive Differential Pulse Modulation (ADPCM) compression and LoRa-based transmission to a central receiver. The backend decodes the audio, applies the BirdNET Artificial Intelligence (AI) model for species identification, and stores results in a MongoDB database with web-based visualization. Experimental validation demonstrates high detection reliability for species with distinctive calls, confirming the system's scalability, energy efficiency, and suitability for long-term biodiversity monitoring in remote environments without continuous connectivity.**

Keywords-. LoRa, acoustic sensors, wildlife monitoring, bioacoustic, low-power IoT, BirdNET, FFT, ADPCM compression, environmental sensing, edge computing.

I. INTRODUCTION

In recent decades, biodiversity conservation has become a global priority. Impacts resulting from climate change, urbanization, intensive agricultural expansion, and noise pollution are causing a drastic decline in many animal species, such as birds and insects. These species groups are essential to ecological balance due to their role in pollination, biological pest control, and forest regeneration [1].

Therefore, passive and noninvasive wildlife monitoring has become a fundamental research and environmental management tool. Traditional wildlife monitoring methods, such as camera trapping or manual censuses, present significant limitations regarding coverage, cost, impact, and dependence on the human factor. Faced with these restrictions, distributed acoustic sensors have proven to be an effective alternative for detecting animal presence through their vocalizations or sounds associated with their activity [2]. Acoustic technology allows species to be detected even in low visibility conditions or at night, greatly expanding observation time windows.

However, one of the main challenges facing acoustic monitoring systems is data processing and transmission. Continuous audio recording generates large volumes of information, which algorithms for artificial intelligence must efficiently manage for storage, transmission, or processing. Furthermore, these systems must operate in remote areas without electrical infrastructure or conventional connectivity.

In this scenario, Low-Power Wide-Area Network (LPWAN) technologies, such as LoRaWAN have emerged as promising solutions. LoRaWAN enables data transmission over long distances (up to several kilometers) with minimal power consumption, utilizing Europe's free 868 MHz spectrum. Unlike other alternatives such as Sigfox or Narrowband Internet of Things (NB-IoT), LoRaWAN stands out for its flexibility, low cost, and open ecosystem, which facilitates its adoption in academic and industrial settings [3].

This paper proposes designing and implementing a distributed network of autonomous acoustic sensors for wildlife detection, focusing primarily on birds. The system comprises nodes based on ESP32 microcontrollers and high-sensitivity digital microphones (such as the INMP441), capable of performing real-time spectral analysis using FFT. Only in the presence of acoustic activity within the expected frequency ranges does the system trigger audio recording and compression, which is then transmitted in fragments via LoRa to another node, which then transmits to its web server. On the server, the Python backend is responsible for receiving and assembling the audio fragments, decoding them, and analyzing them using the locally running BirdNET tool. BirdNET has demonstrated high accuracy in species classification using convolutional neural networks trained with millions of acoustic recordings from birds worldwide. After species identification, the data is stored in a MongoDB database, from which interactive dashboards are generated for visualization and temporal and geographic analysis. The added value of this project lies in the combination of four key elements:

- real-time acoustic detection.
- energy efficiency.

- optimized LoRa communication.
- intelligent local processing.

This architecture allows the system to be deployed in rural or protected areas without constant maintenance or connection to mobile or Wi-Fi networks. Furthermore, its modular and open design facilitates scalability and adaptation to different ecological contexts.

The need for this type of solution is evident in the face of ecosystem management and protection challenges. In European countries like Spain, the decline of endemic species, such as the lesser grey shrike (Lanius minor) and the black stork (Ciconia nigra), requires new monitoring tools that allow for rapid and precise action. Ultimately, this project responds to a real need to improve environmental monitoring systems through low-cost, highly efficient, and minimally intrusive technologies. It provides a viable, replicable, and sustainable solution for researchers and administrators.

The main objective of this work is to design, implement, and validate a low-power distributed acoustic sensor network capable of autonomously detecting and identifying wildlife species—particularly birds—in remote environments without continuous connectivity, and to integrate the collected data into a digital twin for environmental monitoring. The proposed system combines energy-efficient hardware, optimized communication via LoRa, and artificial intelligence-based bioacoustic analysis to provide a scalable, low-cost, and minimally intrusive solution that feeds real-time information into a virtual replica of the monitored ecosystem.

The paper is structured as follows. Section II reviews related work in acoustic monitoring and low-power communication technologies. Section III details the proposed system architecture and operation. Section IV presents and discusses the experimental results. Finally, Section V provides the conclusions and outlines directions for future work.

II. RELATED WORK

Numerous solutions have been developed in recent years for environmental monitoring and bioacoustics detection of wildlife, leveraging Low-Power Wide-Area Networks (LPWAN) such as LoRaWAN due to their low consumption and broad coverage.

A notable reference is the work by FentonSigla [4], which presents a distributed acoustic monitoring system characterized by energy efficiency and flexibility. Based on ESP32 nodes and INMP441 microphones, its architecture shares similarities with this project. However, it only extracts derived acoustic parameters (such as Sound Pressure Level (SPL) or direction of arrival) instead of transmitting audio fragments for detailed analysis.

Other contributions explore complementary approaches, such as the Internet of Things (IoT) architecture by Mohandass et al. [5] for animal health monitoring and intrusion detection, or the work of Ojo et al. [6], which experimentally evaluates LoRa propagation in forest environments. Likewise, Martínez Rach et al. [7] designed a ZigBee-based bioacoustics sensor to detect the red palm weevil, focusing on pest monitoring with high accuracy in acoustic recognition.

Beyond connectivity and hardware, several methods have been proposed for acoustic activity detection and wildlife classification. Traditional threshold-based triggering [8] and more advanced spectral analysis using FFT [10] allow event-driven audio capture, although both are susceptible to false activations under noisy conditions. Recent works have also incorporated Machine Learning at the edge (TinyML), such as Tinybird-ML [9], capable of performing syllable-level bird song analysis with low-power consumption. Similarly, call density estimation methods [10] directly model the occurrence of vocalizations without relying solely on threshold events, increasing robustness in complex soundscapes. Another line of research uses animal-borne soundscapes loggers [11], enabling classification and transmission directly from tags attached to animals, particularly for underwater soundscapes.

One of the most widely adopted tools in classification models is BirdNET [13], an Artificial Intelligence (AI) based system trained with millions of recordings worldwide. BirdNET applies convolutional neural networks to spectrograms for species identification and has demonstrated high accuracy even in noisy conditions. Its ability to operate locally, without dependence on cloud services, makes it especially suitable for autonomous monitoring projects such as the one presented here. Alongside BirdNET, lightweight embedded classifiers based on spectrograms and Mel-Frequency Cepstral Coefficients (MFCCs) [10] have also been explored. However, they remain limited by the computational and memory constraints of low-power devices.

Overall, prior work highlights both feasibility and the challenges of combining low-power communication with bioacoustics analysis. The present study builds upon these contributions by integrating real-time acoustic activity detection, efficient LoRa transmission, and BirdNET-based classification into a distributed sensor network designed for long-term wildlife monitoring.

III. SYSTEM PROPOSAL

This section describes the design and implementation of the proposed distributed acoustic sensor network for autonomous wildlife monitoring. It begins with an overview of the system's architecture, detailing the hardware components, communication modules, and operating principles. Then, it explains the end-to-end workflow, from audio capture and activation strategies to compression, segmentation, and LoRa-based transmission. Subsequent subsections address the reception, decoding, and artificial intelligence-based bioacoustic analysis, followed by the storage and visualization of results.

A. System Overview

The proposed solution consists of the design and implementation of a distributed network of energy-efficient wireless acoustic sensors capable of capturing sounds emitted by wildlife, identifying relevant events in real time, and transmitting the audio fragments to a LoRa-based remote processing infrastructure. The information is processed with artificial intelligence tools to determine the detected species and is stored in a cloud-based MongoDB database, allowing for subsequent visualization and analysis through interactive dashboards.

The developed system is based on a low-power, low-cost architecture, designed to operate autonomously in natural environments. Each sensor node comprises three main elements: an ESP32 microcontroller, an INMP441 digital microphone, and a LoRa communication module with its corresponding antenna. Their technical characteristics and function within the system are described below.

This section describes the overall operation of the system, from audio capture to results visualization. The process is divided into three main blocks: the transmitter, the receiver, and the backend, as shown in Figure 1.

The transmitting node captures audio in short windows and applies real-time FFT spectral analysis to detect acoustic activity in the target band. When a significant event is detected, the recording of the entire fragment is triggered, which is then compressed using the ADPCM algorithm. Once compressed, the file is fragmented and transmitted to the receiving node via LoRa.

At the receiver, the fragments are reassembled to reconstruct the original file. Once completed, the file is temporarily stored and automatically sent via WiFi to a web server for processing.

In the backend, the ADPCM file is decoded into WAV format. The audio is then analyzed using BirdNET to identify animal species based on their vocalizations. The results obtained and the original fragment are stored in a MongoDB database. Finally, all this information is accessible through a web panel that allows users to view, filter, and query the detected acoustic events in a structured manner.

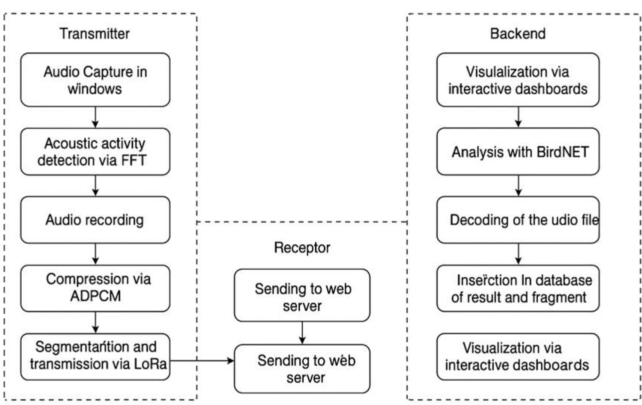

Figure 1. Block diagram of the complete system operation, from audio capture to its analysis and insertion into the database.

B. Acoustic activity detection and audio capture

The first key component of the proposed system is the process of capturing ambient audio by the sensor nodes. Since audio fragments contain critical information for detecting animal species, their recording must be selective, energy-efficient, and accurate enough to ensure its subsequent use in bioacoustic analysis processes. To this end, the nodes implement an intelligent activation strategy, meaning they are not continuously recorded but are activated only when they detect relevant acoustic activity. This decision was made after comparing activation methodologies.

After evaluating these approaches, FFT spectral analysis was selected as the activation strategy because it offers an appropriate balance between accuracy, energy efficiency, and feasibility of implementation on an ESP32-based platform. Table 1 below shows an estimated comparison between the different acoustic activation methodologies [12].

TABLE 1. QUANTITATIVE COMPARISON OF METHODOLOGIES

Method	Precision	Consumption (mA)	Latency
Sound threshold	20–40%	1–5	1 ms
FFT	60–80%	10–20	50–100 ms
TinyML	85–95%	50–100	200–500 ms
Multiple sensors	70–90%	5–15	10–50 ms

Once an acoustic event is detected in the band of interest, the node begins recording an audio fragment. To do this, an INMP441 digital microphone is connected to the ESP32 via the Inter-IC Sound (I2S) interface, allowing high-quality sampling at 16 kHz with 24-bit resolution. The recording duration is set to 20 seconds.

Once the capture is complete, the audio fragment is saved to the ESP32's flash memory using the SPI Flash File System (SPIFFS). This non-volatile memory, accessible like a small virtual disk, allows files to be preserved even after reboots or power losses. Throughout the process, the node continues monitoring the environment to verify the persistence of acoustic activity, thus avoiding storing empty or redundant fragments.

The resulting file represents the basic information unit of the system, which will subsequently be compressed and transmitted using LoRa technology for remote analysis.

C. Audio compression and segmentation

Since LoRa technology presents strict limitations regarding bandwidth and maximum packet size (for example, 51 bytes per packet in the European band with SF12), it is essential to apply data compression techniques to reduce the amount of information before transmission. In this project, the ADPCM algorithm was chosen, widely used in embedded applications due to its low computational cost and good compromise between compression and fidelity. ADPCM is a differential coding technique that predicts the value of the next audio sample based on the previous one and transmits only the quantized difference. This difference is represented with fewer bits than a full sample. In this project, a 4-bit-per-

 ISBN: 978-1-68558-288-3

sample encoding is used, which reduces the file size by half compared to an 8-bit linear Pulse-Code Modulation (PCM), and up to 4 times compared to a 16-bit recording.

The algorithm is efficient enough to run in real time on an ESP32 without the need for additional coprocessors, and the resulting quality has proven sufficient for bioacoustics analysis tasks such as BirdNET classification, especially in environments without excessive noise.

Once the node has captured a 20-second fragment of digital audio, the buffer is processed by the ADPCM algorithm. The result is a compressed binary file that occupies approximately 160KB.

The main advantage of ADPCM in this context is its low CPU and RAM requirements, allowing for efficient real-time implementation without compromising system autonomy. Furthermore, its simple structure facilitates encoding and decoding both on the node and in the backend.

However, it also has some limitations. Its compression is not as efficient as that of codecs such as MP3 or Opus, and it is more sensitive to noise in signals with abrupt changes. Even so, it has been experimentally verified that files compressed with ADPCM maintain sufficient fidelity for BirdNET to correctly identify characteristic vocalizations of wild species.

This file is temporarily stored in memory and later segmented into blocks compatible with LoRa payload limitations. Segmentation is performed by ensuring that each packet contains a header with minimal information such as fragment number, node ID, and end-of-transmission flag. This allows the complete file to be reconstructed on the destination server even if packets are received out of order.

D. Transmission of compressed audio via LoRa

A point-to-point wireless communication system based on LoRa technology was implemented to transmit compressed audio fragments, using two Heltec WiFi LoRa 32 V2 boards. These boards operate on the 433 MHz band, which allows for more flexible experimental use as they are not subject to the duty cycle restrictions inherent to LoRaWAN. The modulation was configured with a Spreading Factor of 7, a bandwidth of 250 kHz, and a coding rate 4/5, optimizing the balance between transmission speed and channel robustness.

The compressed file, approximately 160 kB for 20 seconds of audio, is fragmented into blocks of 220 bytes of data plus a 2-byte header. Each packet includes a sequence identifier and an end-of-transmission indicator, allowing for orderly reconstruction at the receiver. A sliding window of size three is used to improve efficiency, allowing multiple packets to be kept in flight without saturating the channel.

Based on a Heltec board, the receiver reconstructs the file over SPIFFS and acknowledges each packet using ACKs. If a packet is not acknowledged, the transmitter automatically resends it after a delay. Once all the fragments have been received, the receiving node verifies the file's integrity using a Secure Hash Algorithm (SHA-256) hash function and, if everything is correct, sends the file over WiFi to a web server for analysis.

This scheme has proven effective and robust in a laboratory environment, enabling reliable transmission without perceptible loss of quality and the need for LoRaWAN infrastructure [13].

E. Receiving, reassembling and decoding the file

Once all the compressed audio file fragments have been transmitted via LoRa, the receiving node stores them locally and reconstructs the complete file in. ADCPM format [14]. This reconstruction is based on the indices' order in each packet header, allowing the content to be assembled accurately even if the fragments arrive out of order or with an unavoidable delay.

When the End-Of-Transmission (EOF) packet is detected, the file is considered complete and is saved in the receiving node's SPIFFS file system. At that point, the file is automatically sent to a web server via WiFi, where it is decoded.

The backend, developed in Python, converts the ADCPM file into an audio file in WAV format format. To achieve this, a decoder is implemented that reverses the ADPCM compression process, reconstructing a 16-bit, 16kHz linear PCM signal. This transformation is essential to ensure compatibility with acoustic analysis tools such as BirdNET, Audacity, or Sonic Visualizer.

The decoding process is fully automated and is part of the system's continuous processing flow. This integration ensures that each recording transmitted via LoRa can be reliably stored and analyzed, maintaining the fidelity necessary for subsequent acoustic classification based on artificial intelligence.

F. Bioacoustic analysis

Once the audio file has been reconstructed and converted to the appropriate format, the next step is automatically identifying the species in the recording. To do this, BirdNET [15] was used, an artificial intelligence tool developed by the Center for Conservation Bioacoustics at Cornell University, in collaboration with the Technical University of Chemnitz. This platform is specifically designed to recognize bird vocalizations, although it can also detect other types of fauna in more advanced versions.

BirdNET works by converting the audio into spectrograms, which visually represent how the signal's energy is distributed over time and at different frequencies. From this representation, a convolutional neural network model, pre-trained with millions of recordings, can identify characteristic patterns associated with different species.

One of BirdNET's most significant advantages is that it can operate locally, without relying on cloud services. This makes it especially useful in projects like this one, which seek to maintain system autonomy and minimize the need for a permanent connection. For this work, BirdNET-Analyzer was used, a version optimized for execution on personal computers that is easily integrated into automated analysis flows.

BirdNET was chosen for several reasons. On the one hand, it is a tool widely validated in scientific work, with excellent results even in noisy environments or low-quality recordings. On the other hand, it is specifically oriented toward the acoustic analysis of wildlife, which fits perfectly with the objective of this project. Unlike other, more generic audio classifiers, BirdNET returns very detailed information: the common and scientific names of the species, the exact time it was detected, and the confidence level of the prediction.

Furthermore, as an open-source project with clear documentation and an active community, its integration has been relatively simple and offers room for improvement for future versions. However, for it to function correctly, the input files must meet certain format conditions, which have been considered from the early design stages, both in audio capture and compression and decoding.

G. Storing and displaying results

Once the audio file has been reconstructed and converted to the appropriate format, the next step is automatically identifying species. For this purpose, the system integrates BirdNET-Analyzer [14], executed locally on the backend server.

The tool processes the decoded audio fragments by converting them into spectrograms and applying a convolutional neural network inference. The backend records the species name, confidence score, and time stamps for each detected vocalization, storing these results with the original audio fragment in the database.

This integration allows the proposed architecture to benefit from a widely validated AI model while maintaining local autonomy, without needing cloud-based services. Furthermore, the modular design of the backend enables future integration of alternative classifiers (e.g., TinyML models or call density estimation methods) to complement BirdNET or extend recognition to other taxa.

IV. RESULTS

Throughout the development of the system, multiple tests have been performed to validate the correct operation of each module and to assess the performance of different activation methods for acoustic event detection. These tests, described in the corresponding sections, focus primarily on the evaluation of FFT-based activation and the subsequent classification of bird species using BirdNET under control conditions. The validation process employed recordings from the Xeno-Canto platform [16], played back near the microphone to simulate realistic field scenarios. Tests were conducted with three bird species—Eurasian Nightjar (Caprimulgus europaeus), Eurasian Blackbird (Turdus merula), and Mallard (Anas platyrhynchos)—using 10 or 11 audio fragments per species. These species were selected to represent different levels of vocal distinctiveness: the Eurasian Nightjar has a highly characteristic and continuous call, the Eurasian Blackbird produces more common and melodically variable songs, and the Mallard emits short, low-frequency quacks that can be like other waterfowl sounds. This diversity allows the evaluation of the system under varying degrees of classification difficulty.

Next, we will discuss the accuracy and confidence results. For the Eurasian Nightjar, the top 1 classification accuracy was 72.7 % (8 out of 11 recordings correctly identified). Confidence scores for correct detections were generally high but not uniformly near 1.0, with some variability across Figure 2 recordings. This indicates that, even after compression, segmentation, transmission, and reconstruction, the call retains enough spectral fidelity for reliable recognition in most cases, though a fraction of recordings still leads to misclassification.

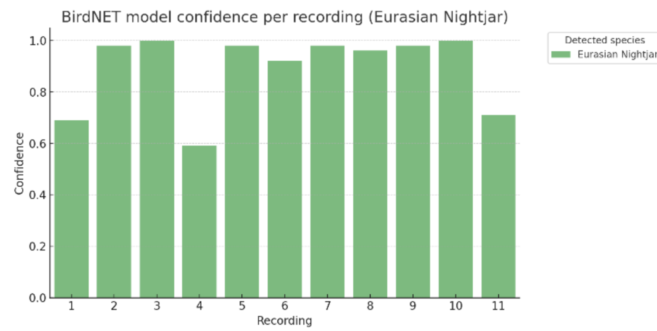

Figure 2. Species detected by recording and confidence level for Eurasian Nightjar.

In contrast, species with more common or less distinctive vocalizations, such as the Eurasian Blackbird or the Mallard, show a greater dispersion in confidence levels and, in some cases, lower results, as seen from Figures 3 and 4. This is consistent with the difficulty of automatically identifying sounds overlapping with many other species.

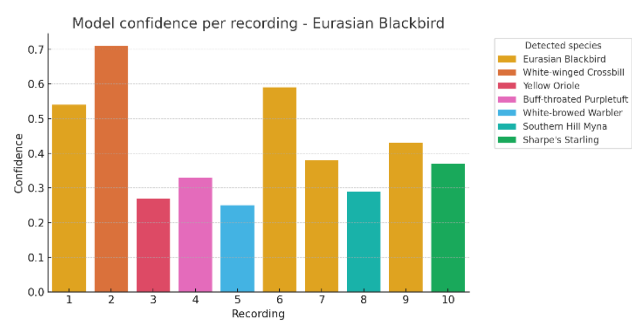

Figure 3. Species detected by recording and confidence level for Eurasian Blackbird.

Figures 3 and 4 show the results of the BirdNET model in ten analyses performed on recordings of the Eurasian Blackbird and Mallard, respectively. Each bar represents a species detected by the model in a specific recording, with its corresponding confidence level. Unlike the Eurasian Nightjar, these recordings show greater dispersion of results, with several species identified as possible candidates. In many recordings, the Eurasian Blackbird, like the Mallard, Anas platyrhynchos, appears with medium or low confidence, while in others it is outperformed by acoustic similar or commonly occurring species. This behavior highlights the system's sensitivity to song characteristics and the fidelity of the transmission process, especially in species with less distinctive vocalizations.

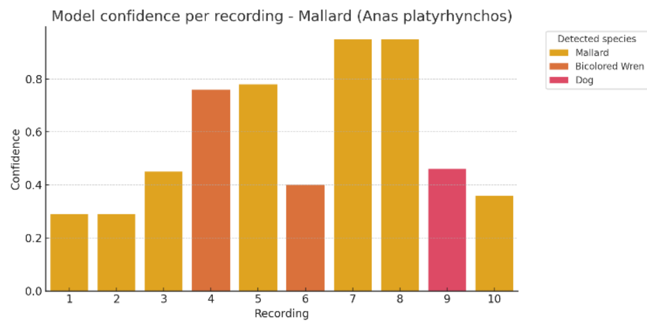

Figure 4. Species detected by recording and confidence level for Mallard Anas Platyrhynchos.

To gain a broader view of the system's performance, a comparative graph was created representing the BirdNET model's reliability for multiple bird species common in urban and natural environments in Spain. This comparison is shown in Figure 5. The graph uses violin diagrams to represent the complete distribution of confidence values obtained by the BirdNET model in the different recordings analyzed for each species. This type of representation allows us to observe the median confidence, the variability, and the density of values. The wider the curve in each area, the greater the concentration of detections in that confidence range. For the Eurasian Nightjar and Common Nightingale species, both with very distinctive and melodic vocalizations, the system presents consistently high confidence values, close to 1.0.

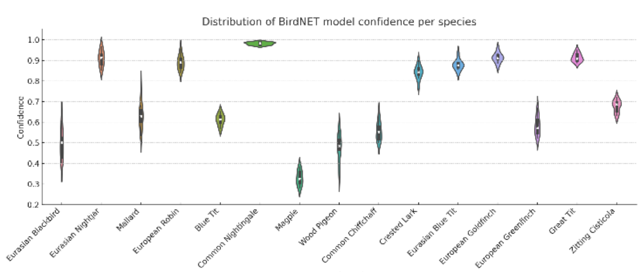

Figure 5. Distribution of BirdNET model confidence by species.

In contrast, species such as the Eurasian Blackbird, the Magpie, or the Wood Pigeon, whose songs are more common, less distinctive, or louder, present more variable and generally lower confidence levels. One of the conclusions drawn from the tests is that part of the loss in detection reliability, especially in species with less distinctive calls, is due to the implemented audio compression. To reduce the size of the transmitted fragments and adapt to the limitations of the LoRa channel and the need for higher transmission speeds, a compression scheme based on the ADPCM codec was chosen. While maintaining reasonable quality for human vocal frequencies and simple song patterns, this introduces degradations that affect the spectral integrity of certain birds' songs. It has been observed that, in complementary tests conducted with the same audio fragments but without applying prior compression, the system's reliability increases slightly, confirming that compression, although necessary for channel efficiency, sometimes negatively impacts the identification capacity of the artificial intelligence model. Furthermore, it should be noted that the system relies on pre-trained AI models (BirdNET), whose extensive database does not always offer uniform performance for all species. It is possible that some of the birds used in the tests are not sufficiently represented in the model's training set, contributing to lower reliability in certain circumstances.

These factors, combined with the acoustic characteristics of each species, explain the differences observed in the quality of the detections and should be considered when interpreting the system's results represented in Figure 6.

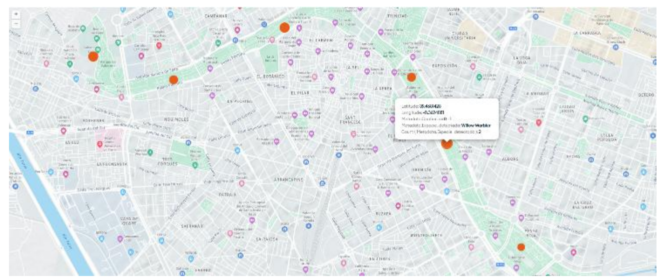

Figure 6. Location map of the detected species.

V. CONCLUSIONS AND FUTURE WORK

This work has presented the design, implementation, and validation of a low power distributed acoustic sensor system that automatically detects wildlife, focusing on birds. The integration of accessible technologies such as ESP32, digital microphones, and LoRa communication, combined with advanced artificial intelligence models (BirdNET), has enabled the development of an efficient and autonomous environmental monitoring solution. Additionally, the system features a modular design that facilitates expansion, integration with additional sensors, and advanced analysis through web platforms.

The results indicate that the system can detect and identify species with distinctive calls under real conditions, maintaining acceptable performance despite limitations imposed by ADPCM compression and the constraints of the LoRa channel. Compression, necessary to optimize transmission, introduces degradations that affect the detection of less distinctive vocalizations, representing a challenge to be addressed.

It has been shown that the AI models and dataset used do not offer uniform coverage for all species, affecting reliability in some instances.

Future directions include optimizing compression algorithms, incorporating edge AI inference in sensor nodes to further reduce data transmission, and deploying the system in natural environments powered by renewable energy to evaluate autonomy and robustness. Additionally, future work should focus on a more comprehensive evaluation of the entire system beyond model accuracy. This includes experiments under real deployment conditions with a network of low-cost, energy-efficient sensors, reporting key performance metrics such as LoRa packet loss rates, battery lifetime (closely tied to local processing and transmission loads), and overall system reliability in the field.

In summary, this project represents a significant advance toward accessible, scalable, and automated biodiversity conservation and monitoring systems, providing innovative tools that could be integrated into large-scale environmental programs.

ACKNOWLEDGMENT

This work has been funded by the Ministry of Science, Innovation and Universities through the State Research Agency (AEI) with the projects PID2020-114467RR-C33/AEI/10.13039/501100011033, TED2021-131040B-C31.

REFERENCES

[1] A. Farina and S. James, “The ecoacoustic event: A conceptual framework for the ecological role of sounds,” *Biological Conservation*, vol. 195, pp. 80–87, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0006320716300118

[2] E. Vidaña-Vila, J. Navarro, C. Borda-Fortuny, D. Stowell, and R. M. Alsina Pagès, “Low-cost distributed acoustic sensor network for real-time urban sound monitoring,” *Electronics*, vol. 9, no. 12, p. 2119, 2020. [Online]. Available: https://www.mdpi.com/2079-9292/9/12/2119

[3] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, “Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios,” *IEEE Wireless Communications*, 2016. [Online]. Available: https://arxiv.org/abs/1510.00620

[4] S. Fenton, “An ultra-low energy solution for large-scale distributed audio monitoring,” *IEEE Access*, vol. 11, pp. 2156–2168, 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10005289

[5] S. Mohandass, S. Sridevi, and R. Sathyabama, “Animal health monitoring and intrusion detection system based on LoRaWAN,” *Turkish Journal of Computer and Mathematics Education*, vol. 12, no. 2, pp. 2397–2403, 2021. [Online]. Available: https://www.proquest.com/openview/81b1732640fc8e67432d8dc6f5cd6bf2/1?pq-origsite=gscholar&cbl=2045096

[6] M. O. Ojo, D. Adami, and S. Giordano, “Experimental evaluation of a LoRa wildlife monitoring network in a forest vegetation area,” *Future Internet*, vol. 13, no. 5, p. 115, 2021. [Online]. Available: https://www.mdpi.com/1999-5903/13/5/115

[7] M. Martínez Rach, H. Migallón Gomis, O. López Granado, M. Pérez Malumbres, A. Martí Campoy, and J. J. Serrano Martín, “On the design of a bioacoustic sensor for the early detection of the red palm weevil,” *Sensors*, vol. 13, no. 2, pp. 1706–1729, 2013. [Online]. Available: https://www.mdpi.com/1424-8220/13/2/1706

[8] J. Juodakis and S. Marsland, “Wind-robust sound event detection and denoising for bioacoustics,” *arXiv preprint*, arXiv:2110.05632, 2021. [Online]. Available: https://arxiv.org/abs/2110.05632

[9] L. Schulthess, S. Marty, M. Dirodi, M. D. Rocha, L. Rüttimann, R. H. R. Hahnloser, and M. Magno, “Tinybird-ML: An ultra-low power smart sensor node for bird vocalization analysis and syllable classification,” *arXiv preprint*, arXiv:2407.21486, 2024. [Online]. Available: https://arxiv.org/abs/2407.21486

[10] J. Liang, I. Nolasco, B. Ghani, H. Phan, E. Benetos, and D. Stowell, “All thresholds barred: Direct estimation of call density in bioacoustic sound event detection,” *Frontiers in Bird Science*, vol. 4, p. 1380636, 2024. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fbirs.2024.1380636/full

[11] T. Noda et al., “Animal-borne soundscape logger as a system for edge classification of sound sources and data transmission for monitoring near-real-time underwater soundscape,” *Scientific Reports*, vol. 14, p. 6394, 2024. [Online]. Available: https://doi.org/10.1038/s41598-024-56439-x

[12] J. D. Rojas and P. A. Dayton, “Optimizing acoustic activation of phase change contrast agents with the activation pressure matching method: a review,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 64, no. 1, pp. 264–272, Jan. 2017.

[13] M. Zaragoza-Esquerdo, L. Parra, S. Sendra, and J. Lloret, “LoRa video streaming in rural wireless multimedia sensor networks,” in *Proc. 2024 19th Int. Symp. Wireless Communication Systems (ISWCS)*, Jul. 2024, pp. 1–6.

[14] T. Nishitani, I. Kuroda, M. Satoh, T. Katoh, and Y. Aoki, “A CCITT standard 32 kbit/s ADPCM LSI codec,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 35, no. 2, pp. 219–225, Feb. 1987.

[15] BirdNET Team, “BirdNET-Analyzer: A tool for bird species identification from audio recordings,” GitHub, 2023. [Online]. Available: https://github.com/birdnet-team/BirdNET-Analyzer

[16] Xeno-Canto Foundation, “Xeno-Canto: Sharing bird sounds from around the world,” 2024. [Online]. Available: https://www.xeno-canto.org.