Journal of Electrical Engineering, Vol. 77, No. 4, 2026, pp. 460-471
https://reference-global.com/journal/JEE

Multimodal beehive weight forecasting using IoT telemetry and acoustic signal processing

Sebastian Górecki[1*], James Brusey[2]

Accurate forecasting of beehive weight is an important challenge in precision apiculture, as it enables continuous assessment of colony development, nectar-flow dynamics, and potential abnormal events. This paper presents a multimodal Internet of Things (IoT) telemetry framework for short-term beehive weight forecasting using data collected by the Intelligent Hives monitoring system. The proposed framework integrates load-cell measurements, environmental sensing, acoustic telemetry, edge data acquisition, and cloud-based analytics. The dataset contains more than 3.4 million telemetry records acquired from over 250 instrumented beehives operating in real-world apiaries between 2022 and 2025. Several forecasting approaches, including multiple linear regression, gradient boosting, and feedforward neural networks, were evaluated using MSE, RMSE, and R^2 metrics. The experimental results indicate that historical weight measurements constitute the dominant source of predictive information, while environmental and acoustic telemetry provide complementary contextual information for multimodal hive monitoring. Furthermore, lightweight regression-based models were found to offer a favorable trade-off between forecasting accuracy and computational complexity, supporting their deployment on resource-constrained IoT platforms. The presented results demonstrate the practical applicability of multimodal telemetry, signal processing, and machine learning techniques for intelligent monitoring and predictive analytics in precision agriculture.

Keywords: precision apiculture, beehive monitoring, IoT, acoustic telemetry, hive-weight forecasting, machine learning, edge computing, predictive analytics

1 Introduction

Honey bees (Apis mellifera) play a vital role in maintaining biodiversity and ensuring global food security through pollination services. The recent decline in bee populations, driven by climate change, habitat loss, diseases, and pesticide exposure, has emphasized the need for improved tools to monitor and support colony health [1-4]. In this context, smart monitoring technologies offer promising avenues to enable more proactive and data-driven management of beekeeping operations [5-7].

Hive weight is a critical indicator of colony status, reflecting nectar inflow, honey production, population dynamics, and the onset of swarming [8]. Traditional methods of monitoring weight based on periodic manual inspections lack temporal resolution and are inadequate for early anomaly detection. The integration of Internet of Things (IoT) devices and edge computing platforms now makes it feasible to collect and transmit continuous telemetry from distributed apiaries in near real-time [5, 6, 9].

Among the many physiological and behavioral signals observed in hives, acoustic signatures represent a particularly rich yet underutilized data source. The buzzing sound of a hive encodes valuable information on colony activity, brood status, and queen presence [10-13]. While past work has explored sound classification and anomaly detection [10, 11, 14], few studies have directly integrated acoustic telemetry into predictive modeling frameworks for key biological metrics such as hive weight.

This study investigates the use of supervised learning models to forecast short-term hive weight changes based on multimodal sensor inputs. The dataset consists of several million time-stamped observations collected via custom-built BeeHUB monitoring devices installed in over 250 hives across diverse environmental conditions. Telemetry includes internal temperature and humidity, external weather data, acoustic features (e.g., dominant frequency, amplitude), and precision weight measurements from strain gauges.

[1] Institute of Electronics, Lodz University of Technology, Stefana Zeromskiego 116, 90-924 Lodz, Poland
[2] Centre for Computational Science and Mathematical Modelling, Coventry University, Priory Street, Coventry, West Midlands, CV1 5FB, United Kingdom
* research@sgorecki.eu

https://doi.org/10.2478/jee-2026-0044, Print (till 2015) ISSN 1335-3632, On-line ISSN 1339-309X

It is hypothesized that multimodal telemetry may improve forecasting robustness and provide complementary information beyond historical hive-weight measurements. The main contributions of this work are as follows:

- A real-world telemetry dataset integrating environmental and acoustic measurements from over 250 smart hives is presented.
- Multiple regression models, including linear regression, gradient boosting, and neural networks, for forecasting hive weight are designed and compared.
- The importance of input features is analyzed, clarifying the role of acoustic data and highlighting directions for further exploration.
- The implications of these results for real-time hive monitoring, anomaly detection, and decision support in precision apiculture are discussed.

The study highlights the feasibility of forecasting hive dynamics using multimodal signals and contributes to the growing body of research in AI-powered agricultural monitoring systems.

2 Related work

Predictive modeling of hive weight has recently gained attention through advanced machine learning approaches. Anwar et al. introduced *WE-Bee* (2022), a bidirectional LSTM encoder-decoder with attention, which combines in-hive sensor data (temperature, humidity, CO_2, acoustics) and external factors (weather, season, hive size) to estimate daily weight changes [15]. They demonstrated that WE-Bee could reliably capture multi-week weight trends and emphasized that diverse training data and feature selection are crucial for real-world robustness. A follow-up model, *Apis-Prime* (2023), used two parallel self-attention encoders on a 2170-day sensor dataset to predict daily per-frame weight, achieving a mean error of 19.7 g and reducing the input feature set from 36 to 23 [16]. LeLouarn et al. (2024) curated a large dataset (758k records of weight, temperature, and video) and applied univariate time-series forecasting models (ARIMA, CNN, LSTM) to predict hive weight 6-24 hours ahead [17]. They found that ARIMA performed comparably to neural networks and that all models successfully captured rising or falling trends.

Despite these advances, weight prediction remains underexplored compared to classification tasks. Braga et al. (2022) observed that most studies focus on anomaly detection rather than forecasting [18]. They proposed a multi-output neural network for short-term prediction of temperature, humidity, and hive weight using telemetry data. Notably, they highlighted a lack of prior studies performing true predictive modeling in apiculture. This work addresses this gap by providing a structured framework for multimodal prediction of hive weight.

Acoustic monitoring of bee colonies has also gained traction. Uthoff et al. (2023) reviewed vibration-based sensing to detect queen presence and swarming activity, reporting that microphones and accelerometers can distinguish queenless hives, but noted small sample sizes and inconsistent feature pipelines [19]. Passive acoustic features, such as dominant frequency and amplitude modulations, have been used to infer behavioral states, including swarming, queen loss, and mite infestation [10]. Mel-frequency cepstral coefficients (MFCCs) are among the most common features used in bee sound studies [20], while others have used Hilbert–Huang and wavelet transforms to capture the non-stationary nature of bee sounds. Classification models range from SVMs and random forests to convolutional and recurrent neural networks [5].

The application of IoT and AI in beekeeping has been thoroughly surveyed. Turyagyenda et al. (2025) outlined a taxonomy of precision beekeeping systems integrating wireless sensors and cloud/edge intelligence [5]. They identified high potential in integrating telemetry with AI but pointed out that most systems remain cloud-centric, with limited on-device processing. Similarly, Danieli et al. (2024) reviewed smart hive platforms offering remote access to temperature, humidity, and weight data, and emphasized the role of usability and energy autonomy in field deployments [21].

Finally, several prototypes of intelligent hives have been developed. Ntawuzumunsi et al. (2021) proposed *SBMaCS*, a solar-powered IoT beehive with motion, flame, and weight sensors for remote control and alerts [22]. Dsouza et al. (2023) presented *HiveLink*, a platform supporting MQTT telemetry and anomaly detection [23]. Bratek and Dziurdzia (2021) focused on custom-designed, energy-efficient load-cell systems for long-term hive weight tracking [24]. This work builds on these foundations by combining regression forecasting based on multimodal telemetry to predict hive weight and support anticipatory diagnostics in beekeeping.

To the best of the authors' knowledge, few studies have evaluated multimodal hive-weight forecasting using a telemetry dataset exceeding three million observations collected from operational IoT deployments over multiple years.

Table 1. Comparison of representative beehive weight forecasting studies

Study	Method	Dataset	Horizon
WE-Bee [15]	BiLSTM	2170 days	Daily
Apis-Prime [16]	Transformer	2170 days	Daily
LeLouarn et al. [17]	ARIMA/ LSTM	758k records	24 h
This work	LR/GB/ FNN	3.4M+ records	Short-term

Table 1 highlights the scale of the dataset used in this study compared with previous beehive forecasting research. While recent works primarily focused on deep-learning architectures, the present study evaluates multiple forecasting approaches using a substantially larger multimodal telemetry dataset collected from operational smart hives.

Compared with previous studies, the present work focuses on short-term forecasting using a substantially larger telemetry dataset collected from operational IoT deployments. In contrast to approaches based exclusively on deep learning, both linear and nonlinear models were evaluated to assess the trade-off between forecasting accuracy and implementation complexity.

3 Methodology

This section details the methodology employed for forecasting beehive weight based on the collected multimodal telemetry data. It outlines the problem formulation, the predictive models selected for evaluation, and the experimental setup including dataset splitting, model training, and performance evaluation metrics.

3.1 Problem formulation

The objective of this study is to predict the short-term future weight of a beehive using historical time-series data. This task is formulated as a supervised regression problem. Given a sequence of observations of environmental conditions $(T_{env}, H_{env}, R_{env}, S_{env}, W_{env}, AQ_{env})$, internal hive conditions (T_{int}, H_{int}), and bee behavior metrics $(A_{features})$ up to time t, along with past hive weight values $(W_{hive}(t), W_{hive}(t-1), \ldots)$, the goal is to predict the hive weight at a specific future time point, $W_{hive}(t+k)$, where k is the prediction horizon (e.g., $k = 1$ for predicting the next recorded time step). The input to the models at time t is a feature vector $\mathbf{f}_t$ comprising the engineered features derived from the raw sensor data and external APIs, including lagged variables and calculated deltas, as described in Sections 3.2 and 3.6.

3.2 BeeHUB monitoring infrastructure

The telemetry data used in this study were acquired using the BeeHUB monitoring platform developed for precision apiculture. Each monitored hive was equipped with an ESP32-based sensing node connected to weight, temperature, humidity, and acoustic sensors. The communication subsystem was based on the SIM7000 cellular modem supporting LTE-M and NB-IoT connectivity. Measurements were transmitted to a cloud-based analytics platform using MQTT.

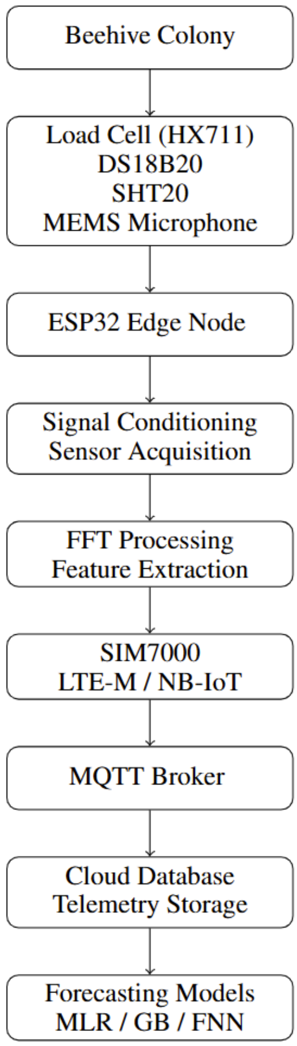

Fig. 1. Architecture of the BeeHUB monitoring system. Multimodal telemetry from hive sensors is processed by an ESP32-based edge node, transformed into compact acoustic and environmental features, transmitted via LTE-M/NB-IoT, and used for cloud-based hive weight forecasting.

(a) (b)

Fig. 2. Physical installation of the BeeHUB sensing system. (a) BeeHUB Queen device and hive-scale platform mounted underneath the beehive for continuous load-cell-based weight measurement. (b) Placement of the internal temperature, humidity, and acoustic sensing module between hive frames to capture representative microclimatic and acoustic conditions inside the colon.

The monitored signals included hive weight, internal temperature, internal humidity, acoustic activity, and external environmental variables. Figure 1 illustrates the overall data flow from the beehive to the forecasting module.

The physical installation of the sensing system is shown in Fig. 2. The hive-scale platform is mounted beneath the beehive to continuously measure colony mass, while the internal sensing module is placed between hive frames to capture temperature, humidity, and acoustic conditions close to the bee cluster. This configuration enables non-invasive monitoring of both external hive-level dynamics and internal colony microclimate.

Table 2. BeeHUB sensing hardware

Component	Function
ESP32	Edge processing and local control
SIM7000	LTE-M/NB-IoT communication
HX711	Load-cell interface
Load cell	Hive weight measurement
DS18B20	Temperature sensing
SHT20	Humidity sensing
MEMS microphone	Acoustic monitoring
Li-Ion battery	Power supply
Solar panel	Energy harvesting

Table 3. Summary of telemetry variables

Variable	Description
Weight	Hive mass (kg)
Temperature	Internal hive temperature (°C)
Humidity	Internal relative humidity (%)
Frequency	Dominant acoustic frequency (Hz)
Amplitude	Acoustic signal amplitude
Weather	External meteorological variables
AQI	Air quality indicators

The measurement subsystem was designed to provide continuous, non-invasive monitoring of colony activity while minimizing disturbance to the bees. The resulting telemetry infrastructure enables long-term acquisition of synchronized multimodal measurements suitable for signal-processing and predictive-analytics tasks.

The proposed system can be viewed as a distributed telemetry and signal-processing platform in which heterogeneous measurements are transformed into compact descriptors suitable for transmission through low-power cellular networks.

The physical installation of the sensing system is shown in Fig. 2. The hive-scale platform is mounted beneath the beehive to continuously measure colony mass, while the internal sensing module is placed between hive frames to capture temperature, humidity, and acoustic conditions close to the bee cluster. This configuration enables non-invasive monitoring of both external hive-level dynamics and internal colony microclimate.

3.3 System architecture

The BeeHUB monitoring platform follows a distributed Internet of Things (IoT) architecture consisting of sensor nodes, edge-computing devices, wireless communication modules, and cloud-based analytics services. Each monitored beehive is equipped with a sensing unit containing a load cell, temperature sensor, humidity sensor, and MEMS microphone. Sensor measurements are acquired by an ESP32-based controller responsible for local preprocessing, feature extraction, and telemetry management. Acoustic signals are sampled locally and transformed using Fast Fourier Transform (FFT) algorithms to derive compact descriptors such as dominant frequency and spectral activity indicators. This reduces communication requirements while preserving diagnostically relevant information. Processed telemetry is transmitted through a SIM7000 cellular modem using LTE-M or NB-IoT communication. Data packets are forwarded via MQTT to a cloud platform where long-term storage, visualization, and predictive analytics are performed. The modular architecture enables scalable deployment of large numbers of instrumented hives while maintaining low energy consumption and reliable telemetry acquisition under field conditions. The proposed architecture reduces communication overhead by transmitting compact features instead of raw acoustic signals, thereby improving energy efficiency and enabling practical large-scale deployments.

3.4 Predictive models

To assess different approaches to modeling the complex relationships between input features and hive weight, three distinct supervised machine learning models were selected for evaluation: Multiple Linear Regression as a simple baseline, and Gradient Boosting and Feedforward Neural Networks as more complex models capable of capturing non-linear patterns.

3.4.1 Multiple linear regression (MLR)

Multiple Linear Regression was used as a foundational baseline model. It assumes a linear relationship between the dependent variable (future hive weight) and the independent variables (engineered features). The model estimates coefficients (β_j) for each feature (f_j) by minimizing the sum of squared differences between observed and predicted values:

$$W_{hive}(t+k) = \beta_0 + \sum_{j=1}^{p} \beta_j f_j(t) + \epsilon$$

where p is the number of features, β_0 is the intercept, and ϵ is the irreducible error. MLR provides interpretability and a simple benchmark to evaluate the performance gain achieved by more sophisticated models.

3.4.2 Gradient boosting

Gradient Boosting is a powerful ensemble technique that builds a prediction model in a stage-wise fashion, typically using decision trees as weak learners. In each stage, a new tree is trained to predict the residuals (errors) of the combined model from the previous stages. The predictions of all trees are then summed to produce the final prediction. Libraries such as XGBoost [25] or LightGBM [26] are popular implementations known for their speed and high predictive accuracy on structured data. Gradient Boosting models can capture complex non-linear relationships and interactions between features automatically.

3.4.3 Feedforward neural network (FNN)

A Feedforward Neural Network is a type of artificial neural network characterized by connections that move in only one direction, from the input layer, through one or more hidden layers, to the output layer. Each layer consists of neurons (nodes) that apply an activation function to a weighted sum of inputs from the previous layer. For this regression task, an FNN architecture with several dense hidden layers, using activation functions like ReLU, and a single output neuron with a linear activation was employed. FNNs are universal function approximators and are capable of learning highly complex, non-linear mappings from input features to the target variable, provided sufficient data and appropriate architectural design.

3.5 Experimental setup

The experimental evaluation was designed to rigorously assess the performance of the selected models on the hive weight forecasting task.

3.6 Acoustic signal processing

Acoustic telemetry was processed into compact descriptors suitable for transmission and forecasting. The dominant acoustic frequency was derived from FFT-based processing performed before storage in the telemetry database.

$$X(k) = \sum_{n=0}^{N-1} x(n) e^{-j2\pi kn/N},$$

where $x(n)$ denotes the discrete acoustic signal sample, N is the FFT length, and $X(k)$ is the complex spectral coefficient for frequency bin k.

Audio was sampled at 8192 Hz. To accommodate the computational and energy constraints of the ESP32 edge node, acoustic processing was performed using consecutive 2 s signal segments. For each segment, an FFT of length 4096 samples was computed and dominant acoustic descriptors were extracted.

Approximately ten consecutive 2-second segments were analyzed during a single acquisition cycle, resulting in an effective observation window of approximately 20 s. The extracted dominant frequency and acoustic amplitude descriptors were subsequently aggregated and stored in the telemetry database, where they were used as forecasting features.

This approach substantially reduces storage and transmission requirements compared with continuous raw-audio acquisition while preserving the most relevant acoustic information for telemetry and predictive analytics.

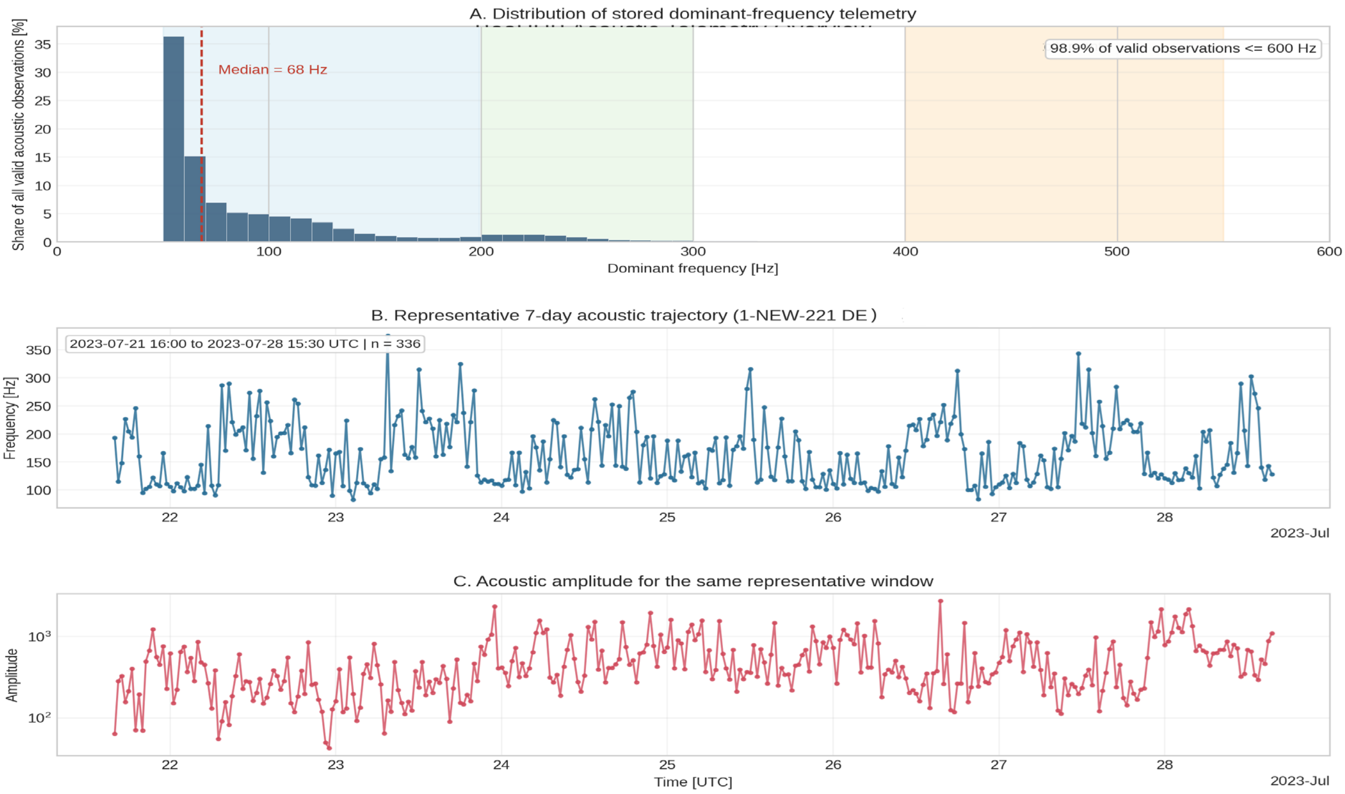

Fig. 3. BeeHUB acoustic telemetry overview based on stored device-level acoustic features. Panel A shows the distribution of dominant frequency across valid acoustic observations, with most measurements concentrated below 600 Hz. Panels B and C show a representative 7-day trajectory of dominant frequency and acoustic amplitude for one hive, illustrating the temporal variability of the acoustic features used by the forecasting pipeline.

Because raw acoustic waveforms were not stored in the telemetry database, the forecasting pipeline used compact acoustic descriptors extracted at the device level. The available acoustic telemetry consisted primarily of dominant frequency and amplitude features. Figure 3 summarizes these stored acoustic descriptors and illustrates their temporal variability in a representative hive.

The observed concentration of dominant frequencies below 600 Hz is consistent with previously reported acoustic characteristics of healthy honeybee colonies. Temporal fluctuations visible in the representative hive indicate that acoustic telemetry captures dynamic behavioral activity and therefore constitutes a potentially useful complementary information source for intelligent hive monitoring systems. The large number of acoustic observations confirms the long-term operational capability of the BeeHUB monitoring infrastructure and demonstrates the feasibility of collecting acoustic telemetry at scale under real-world field conditions.

3.6.1 Dataset splitting

The dataset used in this study contained telemetry records collected from more than 250 instrumented beehives deployed across multiple geographic regions. The forecasting models were trained and evaluated using historical observations obtained from the complete telemetry dataset. The telemetry records were split chronologically into training, validation, and testing subsets. This split was applied at the time-series level to preserve temporal ordering and avoid data leakage. A typical split ratio, such as 70% for training, 15% for validation, and 15% for testing, was applied, ensuring that the testing data represented a future time period unseen during training and validation. This chronological split is crucial for time-series forecasting to avoid data leakage and provide a realistic assessment of the model's ability to predict future values.

Table 4. Dataset deployment overview

Metric	Value
Monitored hives	250
Countries	8
Observation period	Jan. 2022 – Feb. 2025
Records	3,445,496
Sampling interval	15 min

3.6.2 Model training and hyperparameter tuning

Models were trained on the training subset and evaluated on the chronologically separated test subset. Hyperparameter tuning for the Gradient Boosting model and the Feedforward Neural Network architecture (e.g., number of layers, number of neurons per layer, learning rate, regularization) was performed using the validation set. Techniques such as grid search or random search were employed to find optimal hyperparameters that minimized the chosen evaluation metric on the validation data. The Linear Regression model, having no hyperparameters to tune in this context, was trained directly on the training data. Training was implemented using standard machine learning libraries (e.g., scikit-learn [27] for MLR and potentially Gradient Boosting, TensorFlow [28] or PyTorch [29] for the FNN).

3.6.3 Evaluation metrics

The performance of all models was evaluated on the held-out test set using standard regression metrics to quantify the accuracy of the weight forecasts.

- **Mean Squared Error (MSE)**: Measures the average squared difference between the predicted and actual values.

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2$$

where N is the number of test instances, y_i is the actual hive weight, and $\hat{y}_i$ is the predicted hive weight. Lower MSE indicates better accuracy.

- **Root Mean Squared Error (RMSE)**: The square root of the MSE, providing an error measure in the original units of the target variable (kilograms).

$$RMSE = \sqrt{MSE}$$

- ***R*-squared (R^2)**: Indicates the proportion of the variance in the dependent variable that is predictable from the independent variables. An R^2 value closer to 1 suggests that the model explains a larger portion of the variance in hive weight.

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2}$$

where $\bar{y}$ is the mean actual hive weight in the test set.

Comparison of these metrics across the three models, and particularly evaluating models incorporating different feature sets (e.g., with and without acoustic features), allowed for a comprehensive understanding of their respective strengths and the value of multimodal data for hive weight forecasting.

4 Results

This section presents the empirical results obtained from evaluating the predictive models for short-term beehive weight forecasting. The performance metrics for the trained models are reported, and the importance of input features, including those derived from acoustic signals, is analyzed.

4.1 Predictive performance

The trained models were evaluated on the held-out test dataset using the metrics defined in Section 3.6.3. Table 5 summarizes the performance of the baseline Multiple Linear Regression (MLR) model, the Gradient Boosting (GB) model, and the Feedforward Neural Network (FNN).

The results indicate that increasing model complexity did not necessarily improve forecasting performance. The comparable performance of linear and nonlinear models suggests that short-term hive-weight dynamics are largely governed by strong temporal dependencies that can be effectively captured using relatively simple predictive models.

Table 5. Predictive performance of hive weight forecasting models on test data

Model	MSE (kg²)	RMSE (kg)	R^2
Linear regression	0.1807	0.4251	0.9992
Gradient boosting	0.1918	0.4380	0.9992
Feedforward neural network	0.7481	0.8649	0.9969

Interestingly, the simple linear regression model achieved the best forecasting performance, slightly outperforming both gradient boosting and the feedforward neural network. This suggests that short-term hive-weight dynamics in the analyzed dataset are primarily driven by strong temporal dependencies rather than highly nonlinear interactions. From an engineering perspective, this is important because simpler models are easier to interpret and more suitable for future implementation on resource-constrained edge devices.

As shown in Tab. 5, the baseline linear regression model achieved a mean squared error (MSE) of 0.1807, corresponding to a root mean squared error (RMSE) of approximately 0.4251 kg on the test data. The R^2 score of 0.9992 indicates that the MLR model was able to explain a very high percentage of the variance in hive weight fluctuations in the test set. The achieved RMSEof approximately 0.42 kg represents sufficient precision to detect subtle hive events such as nectar inflow variations, early stages of swarming, or minor changes in colony activity. Considering that daily nectar inflow may reach several kilograms and that colony-related weight changes associated with foraging activity, resource accumulation, or early swarming events often exceed 0.5 to 1 kg, the obtained forecasting error is sufficiently small for practical decision-support applications in precision apiculture.

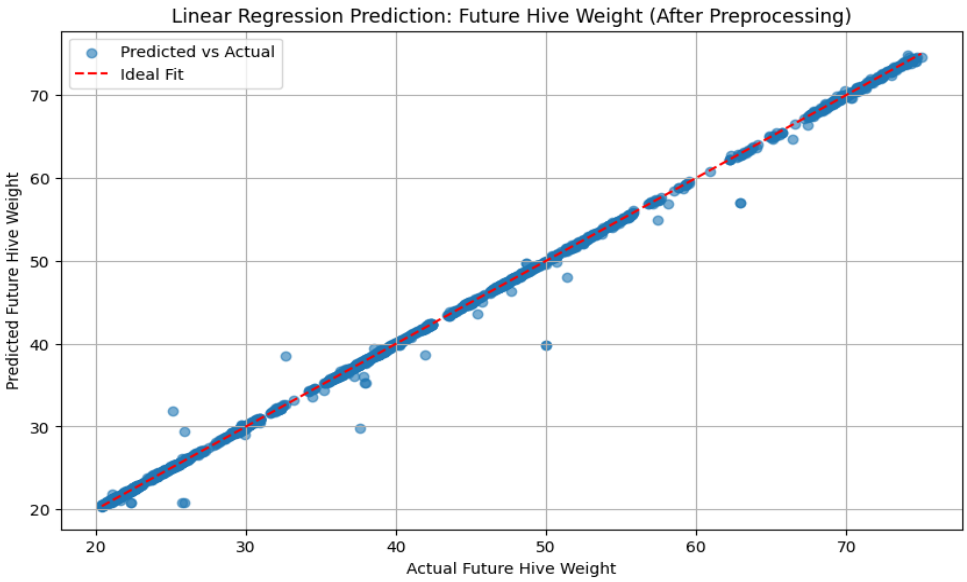

Fig. 4. Linear regression prediction

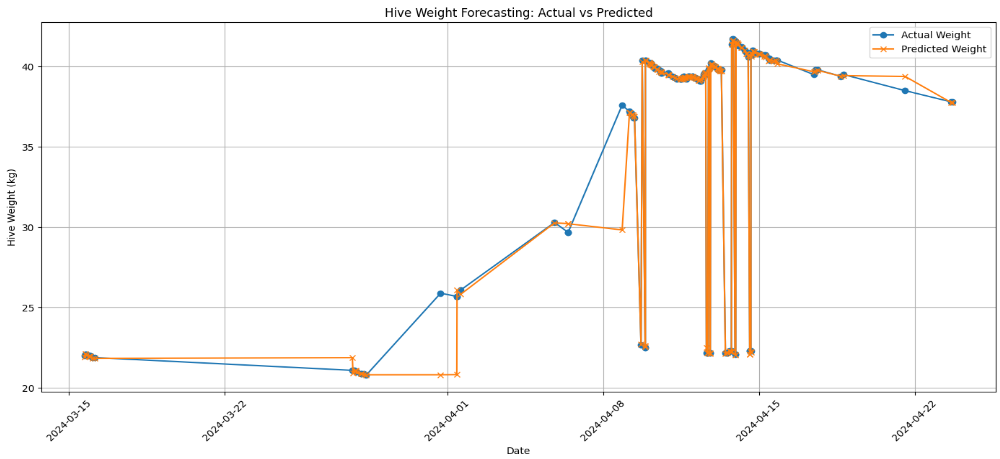

Fig. 5. Example time-series comparison of actual and predicted hive weight over a selected period (7 days)

Evaluation of the gradient boosting and feedforward neural network models revealed performance comparable to or slightly worse than the MLR baseline for this specific dataset and short-term forecasting task. The Gradient Boosting model achieved an MSE of 0.1918 and an R^2 of 0.9992. The Feedforward Neural Network showed slightly lower performance, resulting in an MSE of 0.7481 and an R^2 of 0.9969. While all models achieved very high R^2 values, indicating strong overall predictive power, the smaller MSE and RMSE values for linear regression and gradient boosting suggest they captured the short-term fluctuations with slightly higher accuracy than the tested FNN architecture. The comparable performance between MLR and GB suggests that for this short-term prediction horizon, a large portion of the predictable variance is captured by linear relationships or simpler interactions, which both models handle effectively.

4.2 Feature importance analysis

To understand the contribution of different input signals to the predictive models, a feature importance analysis was conducted using the Gradient Boosting model, which provides a robust measure of feature relevance based on its internal structure. Figure 6 illustrates the top features contributing to the model's predictions.

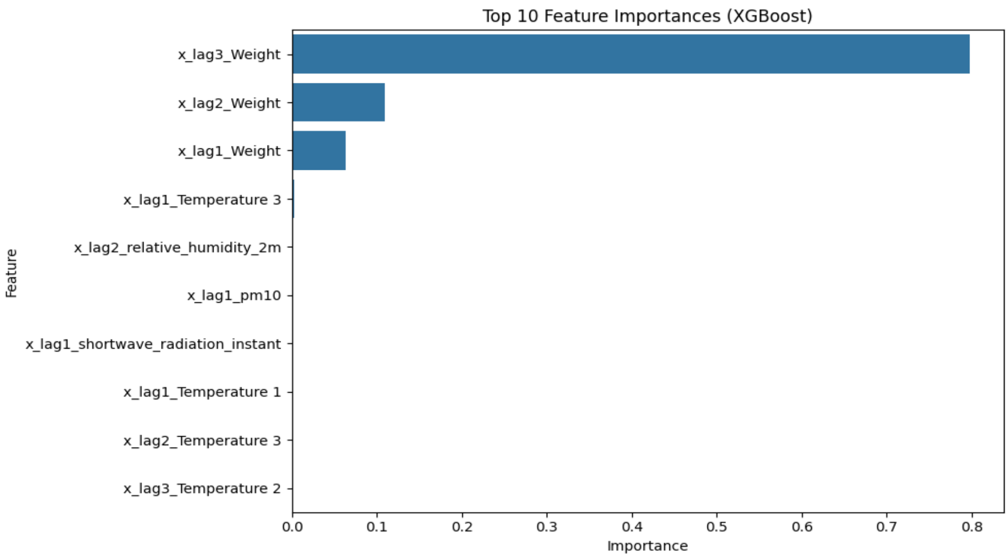

Fig. 6. Feature importance obtained from the gradient boosting model. Although linear regression achieved the highest forecasting accuracy, gradient boosting was used to estimate relative feature contributions because tree-based models provide a natural feature-importance measure.

The analysis confirms that lagged values of hive weight (specifically x_lag3_Weight, x_lag2_Weight, and x_lag1_Weight) are overwhelmingly the most important predictors for short-term forecasting, underscoring the strong auto-regressive nature of the time series. These three lagged weight features alone account for a very large proportion of the total feature importance, contributing a combined total of approximately 97% (with individual contributions of roughly 80%, 11%, and 6% respectively).

Among the features other than lagged weight, lagged internal hive temperatures (such as those from sensors Temperature 3 and Temperature 2 at previous time steps), lagged external environmental parameters (like relative humidity at 2 meters and shortwave radiation instant at previous time steps), and lagged air quality (like PM10 at a previous time step) showed the highest importance, ranking within the top 10 features. These findings reflect the influence of recent internal hive conditions, external weather, and air quality on short-term weight changes.

Acoustic features contributed to the predictive models, although their relative importance was lower than that of lagged hive-weight measurements and selected environmental variables. This observation suggests that acoustic telemetry may provide complementary contextual information describing colony activity rather than acting as a primary driver of short-term weight dynamics.

However, acoustic telemetry remains valuable from a diagnostic perspective because acoustic activity is known to reflect colony behavioral states, queen presence, swarming activity, and stress conditions. Consequently, acoustic sensing may be more suitable for behavioral monitoring and anomaly detection than for direct short-term weight prediction.

The consistent dominance of lagged weight, coupled with contributions from multimodal features, suggests that combining historical weight with relevant internal and external factors provides a strong basis for predictive modeling.

5 Conclusion, limitations, and future work

5.1 Conclusion

This study demonstrated the feasibility and effectiveness of using multimodal telemetry data for accurate short-term forecasting of beehive weight. By integrating historical weight, environmental parameters, internal hive conditions, and acoustic telemetry obtained from the Intelligent Hives Monitoring System (BeeHUB), supervised machine learning models were developed and evaluated, achieving high predictive accuracy ($R^2 >$ 0.99 and RMSE around 0.42 kg). This precision supports practical applications in precision apiculture, including real-time colony monitoring and early anomaly detection.

A notable observation from the feature importance analysis was that lagged hive-weight measurements dominated the forecasting models, while environmental and acoustic telemetry provided complementary contextual information. Although acoustic features contributed less strongly to short-term weight prediction, they remain valuable for behavioral monitoring, anomaly detection, and colony-state assessment. These findings suggest that acoustic sensing may be more informative for behavioral analytics than for direct short-term weight prediction.

The employed methodology, combining engineered time-series features with multiple regression models, proved robust and effective. Overall, this study lays a solid foundation for developing intelligent apiary management systems, leveraging predictive modeling to optimize hive health and productivity.

5.2 Limitations

Despite the encouraging results, several limitations exist:

Firstly, external environmental data sourced from public APIs may not fully capture microclimatic conditions specific to individual apiary sites, potentially introducing discrepancies.

Secondly, the acoustic features evaluated (dominant frequency, spectral density) represent a narrow subset of possible characteristics. Future studies could benefit from exploring advanced acoustic analyses (e.g., MFCCs, wavelet transforms) or applying deep learning methods optimized for audio data.

Lastly, the scope was limited to short-term weight forecasting. Predicting hive dynamics over longer periods introduces complexities tied to seasonal variability, colony development phases, and significant events such as swarming, which were beyond this study's focus.

5.3 Future work

Several promising areas for further investigation are identified based on the current findings:

- **Scalability and generalization study**: Expanding evaluations to include numerous diverse hives across varying geographic and climatic conditions to rigorously test model robustness and transferability.

- **Enhanced multimodal feature integration**: Exploring sophisticated deep-learning architectures (e.g., LSTMs, Transformers) capable of integrating complex multimodal sequences. Advanced acoustic feature extraction or direct audio modeling could further improve predictive insights.
- **Real-time anomaly detection**: Developing systems that leverage prediction errors or specialized detection algorithms to provide early warnings of critical hive events such as swarming, theft, or disease outbreaks.
- **Long-term forecasting and event prediction**: Extending predictive models to forecast longer-term hive trends by incorporating seasonal variables, historical colony data, and external sources like satellite-derived phenological metrics.
- **Edge deployment and TinyML**: Optimizing predictive models for deployment on resource-constrained devices through methods like quantization and pruning to facilitate real-time, edge-based diagnostics and reduce communication overhead.
- **Causal discovery and intervention analysis**: Applying causal inference techniques to identify and analyze potential causal relationships within hive telemetry, enabling deeper insights into how environmental changes or management interventions affect colony outcomes. It is important to note that the present study focused exclusively on predictive modeling; causal relationships should be explicitly investigated using appropriate inference methods.
- **Advanced acoustic signal processing**: Future studies should investigate MFCC, wavelet, and spectrogram-based representations of hive sounds, as well as end-to-end deep-learning approaches for acoustic telemetry analysis.

Exploring these directions can significantly advance precision apiculture, empowering beekeepers with enhanced tools to sustainably manage honeybee health and productivity.

References

[1] D. Goulson, E. Nicholls, C. Botías, and E. L. Rotheray, “Bee declines driven by combined stress from parasites, pesticides, and lack of flowers,” Science, vol. 347, no. 6229, 2015.

[2] E. Genersch, “Honey bee pathology: Current threats to honey bees and beekeeping,” Applied Microbiology and Biotechnology, vol. 87, no. 1, pp. 87–97, 2010.

[3] S. Smoliński and A. Glazaczow, “Causal network linking honey bee (apis mellifera) winter mortality to temperature variations and varroa mite density,” Science of The Total Environment, vol. 954, pp. 176245, 2024.

[4] G. Cilia and A. Nanetti, “Challenges and advances in bee health and diseases,” Veterinary Sciences, vol. 10, no. 4, pp. 253, 2023.

[5] A. Turyagyenda, A. Katumba, R. Akol, M. Nsabagwa, and M. E. Mkiramweni, “IoT and machine learning techniques for precision beekeeping: A review,” AI, vol. 6, no. 2, pp. 26, 2025.

[6] K. Jeong, H. Oh, Y. Lee, H. Seo, G. Jo, et al., “IoT and AI systems for enhancing bee colony strength in precision beekeeping: A survey and future research directions,” IEEE Internet of Things Journal, vol. 12, no. 1, pp. 362–389, 2025.

[7] H. Hadjur, D. Ammar, and L. Lefèvre, “Analysis of energy consumption in a precision beekeeping system,” Proceedings of the 10th international conference on the internet of things, 2020.

[8] B. Rector, G. Mercadier, and N. Holst, “Within-day variation in continuous hive weight data as a measure of honey bee colony activity,” Apidologie, vol. 39, pp. 694–707, 2008.

[9] W. G. Meikle and N. Holst, “Application of continuous monitoring of honeybee colonies,” Apidologie, vol. 46, no. 1, pp. 10–22, 2014.

[10] M. Abdollahi, P. Giovenazzo, and T. Falk, “Automated beehive acoustics monitoring: A comprehensive review of the literature and recommendations for future work,” Applied Sciences, vol. 12, pp. 3920, 2022.

[11] S. Ferrari, M. Silva, M. Guarino, and D. Berckmans, “Monitoring of swarming sounds in bee hives for early detection of the swarming period,” Computers and Electronics in Agriculture, vol. 64, pp. 72–77, 2008.

[12] M. Bencsik, Y. Le Conte, M. Reyes, M. Pioz, D. Whittaker, et al., “Honeybee colony vibrational measurements to highlight the brood cycle,” PLOS ONE, vol. 10, no. 11, pp. e0141926, 2015.

[13] A. De Simone, L. Barbisan, G. Turvani, and F. Riente, “Advancing beekeeping: IoT and TinyML for queen bee monitoring using audio signals,” IEEE Transactions on Instrumentation and Measurement, vol. 73, pp. 1–9, 2024.

[14] Y. Zhao, G. Deng, L. Zhang, N. Di, X. Jiang, et al., “Based investigate of beehive sound to detect air pollutants by machine learning,” Ecological Informatics, vol. 61, pp. 101246, 2021.

[15] M. A. Anwar, M. A.-S. Abdullah, and N. H. Huda, “WE-bee: A bidirectional LSTM encoder–decoder with attention for honey bee hive weight forecasting,” Artificial Intelligence in Agriculture, vol. 7, pp. 58–69, 2022.

[16] M. A. Anwar, M. A.-S. Abdullah, and N. H. Huda, “Apis-prime: A novel attention-based model for high-resolution honey bee hive weight forecasting,” IEEE Access, vol. 11, pp. 128844–128856, 2023.

[17] M. LeLouarn, P. Pustelnik, N. Courteille, and T. Lardy, “Forecasting hive weight using time series analysis, CNN and LSTM,” Computers and Electronics in Agriculture, vol. 217, pp. 108570, 2024.

[18] J. Braga, J. P. Moura, T. Sousa, R. Morais, J. Gomes, et al., “A neural network approach for short-term prediction of temperature, humidity, and weight in honey bee hives,” Computers and Electronics in Agriculture, vol. 197, pp. 106887, 2022.

[19] D. Uthoff, J. Hentschel, A. Zacepins, and J. Meitalovs, “Acoustic and vibration analysis to detect queen presence and swarming activity in honey bee colonies: A review,” Computers and Electronics in Agriculture, vol. 213, pp. 108240, 2023.

[20] M. Mehdi, S. M. Alirezaee, M. Pour-Abdollah, M. M. Aref, and P. Giovenazzo, “Acoustic signal analysis for automatic detection of queen honey bee condition,” Computers and Electronics in Agriculture, vol. 204, pp. 107540, 2023.

[21] M. Danieli, L. Comuzzo, E. De Zorzi, C. Zanon, and A. Cenedese, “A survey on smart hive platforms for precision beekeeping,” IEEE Access, vol. 12, pp. 11115–11139, 2024.

[22] J. Ntawuzumunsi, F. Kwizera, and J. Mugisha, “SBMaCS: A smart beehive monitoring and control system for beekeeping automation,” 2021 international conference on innovative trends in communication and computer technologies (ITCT), pp. 151–156, 2021.

[23] A. Dsouza, A. P, and S. Hegde, “HiveLink, an IoT based smart bee hive monitoring system,” 2023.

[24] R. Bratek and A. Dziurdzia, “A wireless system for continuous honey bee hive weight tracking,” Measurement, vol. 180, pp. 109573, 2021.

[25] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp. 785–794, 2016.

[26] G. Ke, Q. Meng, T. Wang, W. Ze, S. Cao, et al., “Lightgbm: A highly efficient gradient boosting decision tree,” Advances in neural information processing systems, vol. 30, 2017.

[27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, et al., “Scikit-learn: Machine learning in python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[28] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, et al., “TensorFlow: Large-scale machine learning on heterogeneous distributed systems,” 12th USENIX symposium on operating systems design and implementation (OSDI 16), pp. 265–283, 2016.

[29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, et al., “PyTorch: An imperative style, high-performance deep learning library,” Advances in Neural Information Processing Systems, vol. 32, 2019.

Sebastian Górecki received the B.Eng. degree in Electronics and Telecommunications and the M.Sc. degree in Computer Science from Lodz University of Technology, Poland, in 2018 and 2020, respectively, and the MBA degree in Digital Transformation from Kozminski University, Warsaw, Poland, in 2023. He is currently pursuing the Ph.D. degree in Automation, Electronics, Electrical Engineering, and Space Technologies at Lodz University of Technology. He is the founder of Intelligent Hives, where he develops Internet of Things (IoT) and edge artificial intelligence solutions for precision agriculture and environmental monitoring. His research interests include IoT, edge AI, cyber-physical systems, and precision agriculture.

James Brusey received the B.Ap.Sc. degree and the Ph.D. degree from RMIT University, Melbourne, VIC, Australia, in 1996 and 2003, respectively. He is currently a Professor of Computer Science with the Centre for Data Science, Coventry University, Coventry, U.K. His research interests include machine learning for cyber-physical systems, wireless sensor networks, reinforcement learning, edge intelligence, intelligent sensing systems, and decision-support technologies.

Received 19 June 2026