Street Art & Urban Creativity | Vol. 11, No. 6, 2025 | ISSN 2183-9956
Street Art & Urban Creativity. Scientific Journal
https://doi.org/10.62161/sauc.v11.5990

AI-BASED PREDICTION MODELS FOR URBAN PARKING AVAILABILITY
A CASE STUDY OF VALENCIA

FRANCESC DOMENECH, VICENTE CAROT, CRISTINA PORTALÉS, INMACULADA COMA
[1] UNIVERSITY OF VALENCIA, SPAIN

KEYWORDS
Urban parking prediction
Smart City
Artificial Intelligence
Machine Learning
Recurrent neural networks

ABSTRACT
Efficient parking access is crucial for urban mobility in smart cities. This study presents a pilot system predicting public parking occupancy in Valencia, Spain, using municipal sensor data. We developed and compared recurrent neural network architectures (RNN, LSTM, GRU), achieving accurate forecasts with performance variations across locations and times. Explainable AI methods provided model interpretability and insights into variable influence. Results indicate that baseline recurrent models yield low MAEs, while Bayesian hyperparameter optimisation offers only marginal gains, highlighting the practicality of straightforward recurrent approaches for urban parking prediction.

Received: 16 /08 /2025
Accepted: 28 /09 / 2025

1. Introduction

The increase in urban density and growing number of vehicles pose significant challenges for parking management in large and medium-sized cities such as Valencia. Inefficient search for parking spaces increases congestion, unnecessary fuel consumption and pollutant emissions (Shoup, 2005). As part of the smart city movement, municipalities are exploring data-driven solutions and emerging technologies (Batty et al., 2012). Local councils have developed guides, such as *Transformative Mobility* (2024), to synthesize best practices in parking management for sustainable mobility. They have also promoted technological initiatives, such as the *Smart Parking* project at PCT Cartuja (Junta de Andalucía & Telefónica, 2021), which leverage real-time occupancy monitoring through cameras and inform users via mobile apps and display panels.

One of the possible applications is to optimize the use of available parking spaces and improve the experience for drivers.

The city of Valencia has a system that allows the number of available spaces to be monitored in real time in a set of public car parks. This information is available on the Valencia City Council's open data platform and is displayed on the information panels distributed in different parts of the city, helping manage urban mobility more effectively.

This paper presents a comprehensive evaluation of deep learning algorithms for predicting parking availability in Valencia. The proposed models forecast which car parks are likely to have spaces available at specific days and times, using historical data, behavioural patterns, and machine learning techniques. Their implementation can support more efficient mobility management and contribute to broader sustainability and quality-of-life goals in metropolitan areas.

Building on historical usage data provided by the Valencia City Council, we develop a scalable system capable of predicting the availability of public car parks using deep learning methods. The study compares the performance of three recurrent neural network (RNN) architectures—basic RNN, LSTM, and GRU—in both baseline and optimised forms, the latter obtained through Bayesian hyperparameter tuning. To enhance interpretability, a surrogate model is employed to identify the variables with the greatest influence on predictions.

Once the validity of the proposed solution has been verified, the new parking occupancy data generated in real time will be integrated to continuously optimize its predictions through automated deployments and training.

This article is organized as follows. Firstly, a review of the state of the art in techniques for predicting the availability of parking spaces is carried out in Section 2. Section 3 details the methodology followed for the analysis of the data and the application of the prediction models. Section 4 analyses and discusses the results obtained in the different models and compares results with previous works. Section 5 applies an AI explainability model to the results obtained. Finally, the conclusions and future work are shown in Section 6.

2. State of the art

The development of systems to predict the availability of parking spaces is part of the evolution of smart cities, which aims to optimise urban services through the use of advanced information and communication technologies (Batty et al., 2012) and artificial intelligence systems.

In recent years, various approaches have been explored to predict the availability of parking spaces, combining IoT sensors, historical data analysis and artificial intelligence (AI) techniques. AI techniques for predicting parking availability leverage a variety of machine learning (ML) and deep learning (DL) models to improve accuracy and efficiency. These models use a variety of datasets, including historical occupancy rates and contextual factors, to forecast available parking spaces.

Different studies demonstrate the potential of time series analysis in urban contexts, including mobility from mobile phone data (Calabrese et al, 2011) or data collected from pervasive network infrastructure (Zheng et al, 2015). Pozo et al. (2022) developed a predictor of on-street parking in Madrid, Spain, exploiting occupancy data collected from smart parking infrastructures and Moreno (Moreno Esteban, 2021) conducted an occupancy prediction study for several car parks around train stations in the province of Barcelona, showing strong daily cyclic patterns in usage

Traditional machine learning models have been widely applied to parking availability prediction due to their simplicity and interpretability. These models include Decision Trees (DT), Random Forests (RF), K-Nearest Neighbors (KNN), and others. Some studies have shown that these simpler algorithms often outperform more complex models such as Multilayer Perceptron (MLP) in terms of prediction accuracy (Awan et al., 2020; Inam et al., 2022).

Deep learning models have gained significant attention in recent years due to their ability to learn complex patterns in data. Thus, Neural Networks have been used to predict parking space availability by learning the relationships between various factors such as time, weather, and traffic conditions. For instance, a study using data from the SFpark project in San Francisco demonstrated that Neural Networks outperformed other traditional models for time series forecasting like ARIMA and SARIMA in terms of mean squared error, especially when exogenous variables such as day type and time of day were considered (Sebatli,2023). Neural networks and temporal convolutional networks have been employed to capture both spatial and temporal features, leading to improved prediction outcomes (Chen et al., 2023; Zhang et al., 2024).

Much in the same way, architectures based on Recurrent Neural Networks (RNNs) and their variants such as Long Short-Term Memory (LSTM) (Hochreiter & Schmidhuber, 1997) and Gated Recurrent Units (GRU) (Cho et al., 2014) are well-suited for time series prediction tasks, making them a natural choice for parking availability prediction. LSTMs, in particular, have been shown to perform well in long-term time series prediction including problems related to parking occupancy prediction using historical data and traffic conditions (Barraco et al, 2021; Vieta, 2024; Yuen et al., 2021). Likewise, GRUs have been used to model the spatial relationships between parking lots and temporal dynamics of parking behaviors (Zhao & Zhang, 2024).

Despite these advances, several challenges remain. On the one hand, the quality and availability of historical and real-time data can be limited or heterogeneous (Zheng et al., 2014). On the other hand, the variability associated with contextual factors such as cultural events, changes in public transport supply or weather conditions remains difficult to accurately model (Yang et al., 2019).

The need for interpretability of AI models is another fundamental aspect to analyse when using these prediction systems. One effective strategy is the use of surrogate models, simple and inherently interpretable approximators (e.g., decision trees or linear regressions) trained to mimic the behaviour of a more complex "black-box" model. By analysing the surrogate's structure or feature importances, we gain insight into which inputs drive the original model's predictions.

In summary, the literature shows a clear trend towards the use of deep learning techniques to solve the problem of predicting parking availability in urban environments. The present work is part of this line, evaluating the use of RNN, LSTM and GRU on real data from the city of Valencia, and providing a systematic comparison between basic architectures and optimized versions. It also incorporates interpretability techniques to improve the understanding and applicability of the results in urban mobility management systems.

In the specific case of Spanish cities, the case of Valencia represents a particularly interesting real context, given the availability of public data and the progressive implementation of smart mobility systems.

3. Methodology

This section describes the methodology followed in the present work. First, the collected data is analysed and pre-processed to remove erroneous or missing values. Next, an exploratory analysis of the data is performed. Finally, the dataset is used to train a predictive model based on neural networks. The following subsections provide a more detailed description of these steps.

3.1. Data collection and description

The Valencia City Council, on its open data portal publishes real-time information on the availability of spaces in the city's public car parks. In addition, although the history of use of the car parks is not published in the open data portal, the city council through the Smart City office stores this information. Thus, the information used as the starting point for this work are:

- CSV file containing parking space data for each lot, covering the period specific to each facility
- CSV file with the identifier and name of the car parks.

These files contain information on 20 car parks, which are used in the analysis detailed below. For clarity of presentation, this paper shows results for only one or two car parks, although the method is applied to the entire set.

3.2. Data preparation

The file with the information on parking spaces contains 13 million records, corresponding to the 20 car parks. It is common for large amounts of data to have erroneous, invalid or missing values. Therefore, the first step is to carry out an initial inspection and data cleaning.

The data file provided contains the following columns (table 1):

Table 1. Dataset information

name	Meaning
_id	Car park id (alphanumeric)
entityId	parking id
entityType	indicates that the parking lot is of type parking
availableSpotNumber	total number of available parking spaces
availableSpotPercentage	percentage of available seats
totalSpotNumber	total number of parking spaces
idParking	Car park identifier
recvTime	time of data recording
TimeInstant	time when the data is received

Source: Own elaboration, 2025.

The dataset undergoes the following processing steps:

1. All redundant columns or those containing a large amount of missing data are removed, as they are not used in the analysis, leaving: availableSpotNumber, totalSpotNumber, and idParking.
2. The necessary type conversions are performed: for example, correcting records that refer to the same parking lot, but have different identifiers due to erroneous registration.
3. Another corrected inconsistency involves records where the total number of parking spaces varies for the same location.

4. If there are few erroneous records for a car park, they are corrected, as they are interpreted as transmission errors, and if there are many, the data for that car park is discarded.

Subsequently, a function is defined to display the time range for each car park (see Figure 1), as well as the periods with missing data, which are linearly interpolated. Erroneous records are also identified, that is, those that availability is greater than the total number of parking spaces or availability is negative

Figure 1. Example of data collected for the car park with ID = 78 – Severo Ochoa

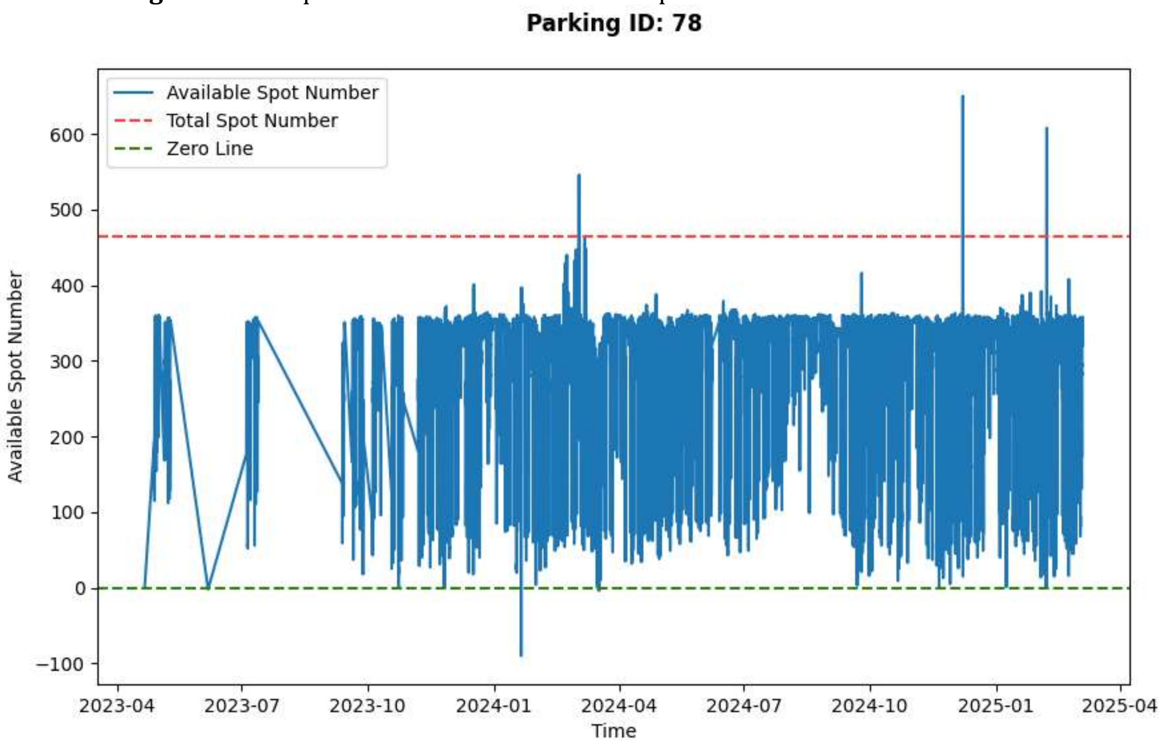

Source: Own elaboration, 2025.

This data cleansing feature is applied for each car park, and all records associated with each car park where there are long periods of missing data are deleted, as imputing such long periods can compromise the quality of the results. Those in which there is little or incorrect data are also deleted.

With the analysis carried out so far, the main problems in the data received are detected:

- Records of the same car park associated with different identifiers, although it should be the same.
- Records of the same car park where the number of total parking spaces varies.
- Erroneous registrations (where the number of available spaces is negative or exceeds the total number of parking spaces).
- Records with the data types incorrectly converted.

Based on the analysis carried out so far, the main issues in the received data are identified:

- Several records associated with the same time slot, in this case the average of these records is calculated.
- Time slots without records (missing values, NA). Here it is imputed using a weighted average between the previous 4 records and that same record, at the same time, in the previous week.

- If many missing values appear at the beginning of the series, data will be used when the null values are significantly reduced, so as not to compromise the result of the predictions, since imputing long periods can be dangerous.

Applying these corrections for each car park removes records that may contain erroneous data.

In summary, at the end of this stage, the data is obtained without irregularities: a single piece of data per hour from the beginning to the end of the time series. An output file that maintains 3 variables is obtained: the car park id, the percentage of available spaces in a range between 0 and 100, and the time instant to which the record refers.

3.3. Exploratory data analysis and periodogram

This stage provides an initial understanding of how the data are distributed, enabling the detection of repetitive temporal patterns and the identification of when and how occupancy changes over time. To this end, two techniques are applied: exploratory data analysis (EDA) and a periodogram to study the data in the frequency domain.

In the EDA, the percentage of available parking spaces (ranging from 0 to 100) is calculated for each car park, with the data aggregated by day of the month, day of the week, and month of the year.

As an example, the results of the exploratory data analysis for two representative car parks is presented.

For car park ID 6 (Figure 2), availability decreases during weekends, suggesting use for shopping and leisure activities, while an increase in August reflects reduced demand during the holiday season. In contrast, car park ID 7 shows increased availability both at weekends and in August, indicating predominant use during working hours, with higher occupancy on weekdays and lower demand during periods of reduced professional activity.

Figure 2: EDA analysis for 2 car parks

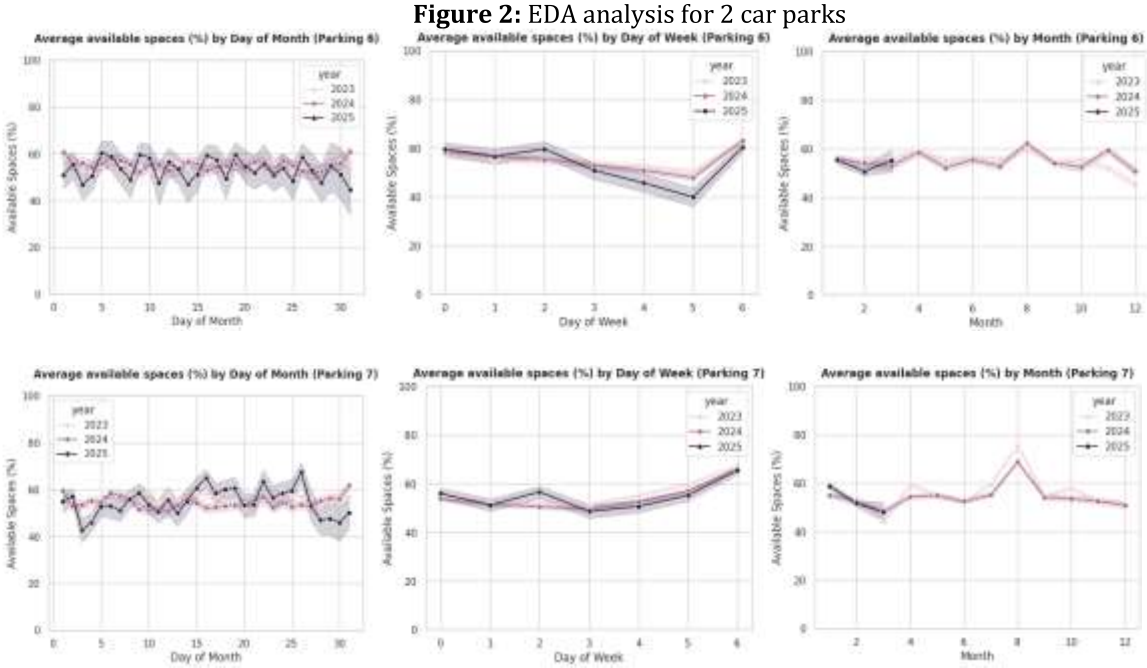

Source: Own elaboration, 2025.

Although exploratory data analysis provides valuable insights, it may be insufficient for detecting periodic behaviours. To address this, a periodogram was computed (Figure 6), transforming the time-domain data into the frequency domain to assess the significance of each frequency component. Each bar on the chart shows which frequency is most important in the time series. If the bar is high, it means that this frequency stands out and can indicate that there are patterns that are repeated, every day, every noon, etc.

For car park ID 6, notable cycles occur at 24, 28, and 12 hours, potentially corresponding to daily routines, irregular patterns such as part-time work or school schedules, and work-shift transitions, respectively.

Figure 3: Periodogram for parking with ID = 6

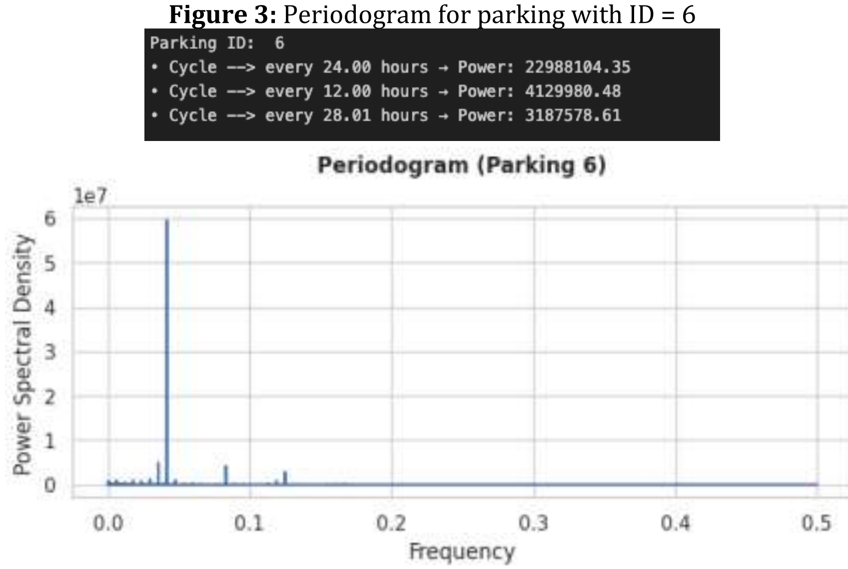

Source: Own elaboration, 2025.

3.4. Dataset split and feature engineering

The cleaned and pre-processed dataset is used to train a predictive model for car park occupancy. The target variable is the percentage of available parking spaces, rescaled from 0–100 to 0–1 to improve convergence and accelerate training. Three categorical variables—day of the week, month, and weekend indicator—are incorporated to provide contextual information. These categorical features are encoded using embedding layers, which enabled the model to learn dense vector representations and capture relationships between categories.

Recurrent Neural Networks (RNNs) are selected due to their ability to process sequential data and leverage past information to predict subsequent values. The model input consisted of 24-hour time windows, which were determined from the periodic patterns identified during the exploratory data analysis. Each window contained the 24 hours preceding a given time and was used to predict availability at that time.

The hold-out method is applied to divide the dataset into training (70%), validation (15%), and test (15%) subsets. The test set covers an identical time range for all car parks to ensure performance comparability, while the training and validation subsets were individually adjusted according to each car park's data availability.

The model is trained using data from car parks with IDs 6, 7, 8, 13, 34, 75, 77, and 78, employing rescaled availability values and categorical embeddings within 24-hour input sequences.

3.5. Modelling and evaluation

The models are trained using three neural network architectures (Figure 4):

- Recurrent Neural Networks (RNN): The most basic architecture, consisting of a hidden layer that feeds back on itself at each time step. RNNs handle short sequences effectively but have limited long-term memory retention.
- Long Short-Term Memory (LSTM): An RNN variant designed for long-term information retention. It incorporates three control gates—forget, input, and output—and an internal memory cell to regulate information flow. While capable of capturing long-term dependencies, LSTMs are more computationally expensive.
- Gated Recurrent Unit (GRU): A lighter alternative to LSTM, with fewer parameters and faster training. It uses an update gate and a reset gate to control information flow, achieving competitive performance with reduced complexity.

The objective is to compare the basic version of each architecture, without modifications (single layer, 32 neurons per layer, no optimizations) with an optimized version obtained through Bayesian hyperparameter optimization, which was selected over random search for its greater efficiency in exploring the search space. In machine learning, this basic, unmodified version of a model it is usually called "vanilla model". We will use this term to refer to a baseline before adding custom features or optimizations.

Figure 4: Architecture of the neural networks used

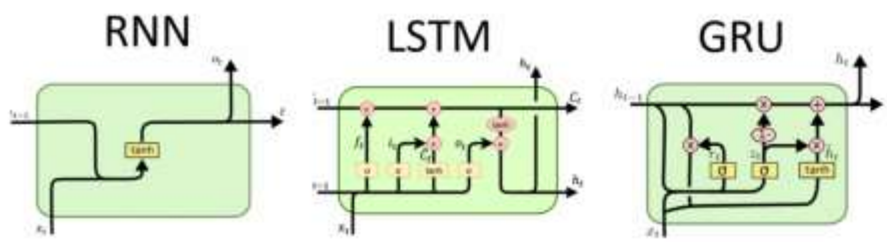

Source: Hasan, 2020.

The hyperparameter tuning process for the optimized models considers adjustments to the learning rate, number of layers, and number of neurons per layer. The tested hyperparameter configurations for each approach are shown in Table 2

Table 2. parameters proposed in each architecture

Parameter	Value	Optimal
Optimizer	Adam	Adam
Batch Size	32	32
Learning rate	0.001	[0.005-0.02]
Weight decay	0.05	0.05
Dropout	0.2	0.2
Number of layers	1	[2,3,4]
Number of neurons per layer	32	[16,32,64]

Source: Own elaboration, 2025.

4. Results and discussion

4.1. Loss Function Evolution

In predictive modelling, the loss function quantifies the discrepancy between a model's predictions and the actual observed values, guiding the optimization process during training. By minimizing the loss, the model iteratively adjusts its parameters to improve both accuracy and generalization. In the context of parking space prediction, the loss function serves as a key indicator of how effectively the model captures occupancy patterns and forecasts availability over time.

For Parking ID 6 with a recurrent neural network (RNN) architecture, the training and validation loss curves (Figure 5) exhibit a steep decline during the initial epochs, stabilizing at near-zero values (approximately 0.002–0.003). The minimal divergence between training and validation loss indicates stable convergence and negligible overfitting.

Two key factors likely contributed to this behaviour. First, the bounded range of input variables stabilized the neural network activation dynamics. Second, the implementation of Early Stopping terminated training when improvements plateaued, thereby reducing unnecessary computation and mitigating the risk of overfitting.

Figure 5: Evolution of training and validation loss for Parking ID 6 using the RNN model.

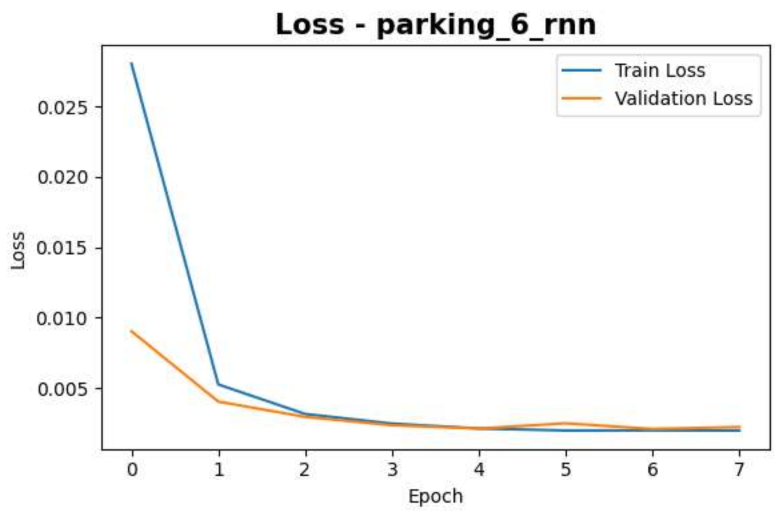

Source: Own elaboration, 2025

4.2. Baseline Model Performance: Mean Absolute Error (MAE)

The Mean Absolute Error (MAE) was selected as the evaluation metric due to its interpretability. MAE represented the average absolute deviation between predicted and actual values, regardless of direction, providing a direct measure of prediction accuracy. Importantly, in this work, the MAE was expressed as a percentage of occupied spaces rather than absolute counts. For example, a 2.4% MAE for car park 13 in the vanilla GRU model corresponded to approximately 5.49 spaces out of a total capacity of 229.

The MAE results for the baseline (vanilla) models were shown in Figure 6. Several observations emerged from this analysis. Certain parking facilities exhibited higher prediction errors, notably Parking 34, which reached 4.0% MAE with the GRU architecture. In contrast, car parks 13 and 8 yielded MAEs in the range of 2.2–2.3%, suggesting highly regular occupancy patterns. Regarding architectural comparison, LSTM and GRU delivered comparable performance, generally outperforming the simpler Multilayer Network (MN).

Figure 6: Mean Absolute Error by parking facility and model architecture (vanilla configuration).

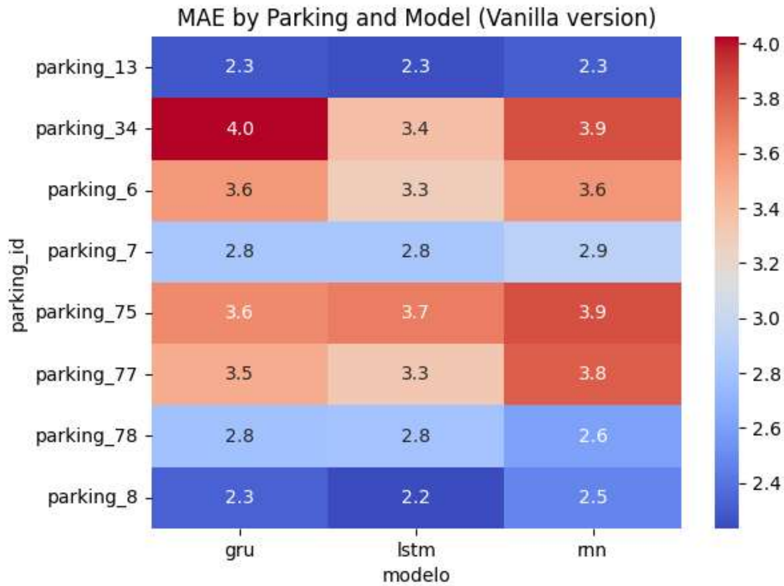

Source: Own elaboration, 2025.

4.3. Optimized Model Performance: Bayesian Hyperparameter Tuning

Bayesian optimization via Optuna was employed to refine the hyperparameters of the RNN, LSTM, and GRU architectures. The results (Figure 7) indicated overall improvement, as in most parking–model combinations the optimization yielded lower MAE values, confirming the benefit of hyperparameter tuning.

However, the gains were generally modest, reflecting the relatively regular and predictable nature of occupancy patterns in the dataset.

A notable exception was Parking 34 with the GRU model, which showed degraded performance, with the MAE increasing from 4.0% to 6.9%, indicating sensitivity to hyperparameter configurations and potential overfitting. Despite this, most optimized models achieved MAEs below 3%, meeting the threshold for real-time deployment.

Figure 7: Mean Absolute Error by parking facility and model architecture after Bayesian optimization.

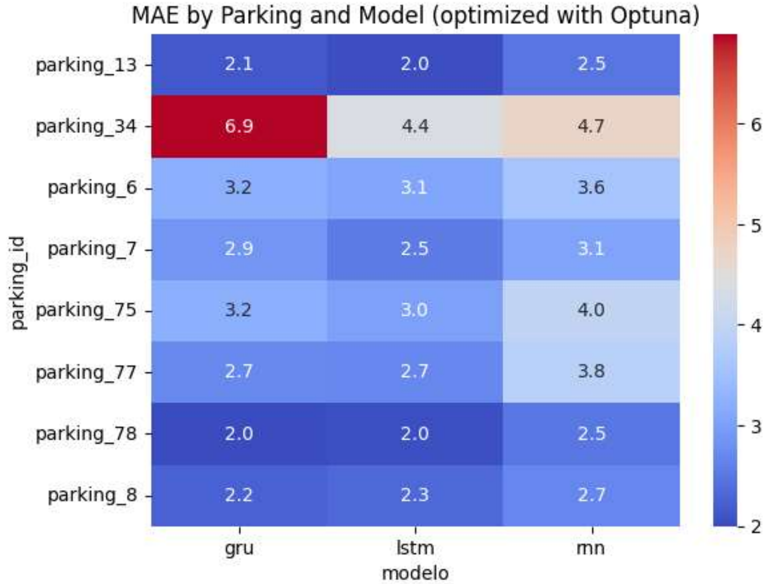

Source: Own elaboration, 2025.

4.4. Dynamics and Prediction Fidelity

Figure 8 underscored the ability of RNN, LSTM, and GRU models—both vanilla and optimized—to capture the strong periodicity in parking occupancy at a fine temporal resolution.

Even without optimization, all models closely tracked the cyclical patterns, indicating that the dominant temporal structure was readily learnable from historical data.

Optimized configurations obtained via Bayesian search tended to better fit peak and trough transitions, particularly during abrupt changes in occupancy (e.g., around hours 30–40 and 120–140), where vanilla models exhibited slight lag or underestimation.

Discrepancies between predictions and actual values were most visible during sharp, short-lived deviations from the main cycle, likely caused by atypical events or demand shocks not fully captured by the training data.

Given the small performance gap, the choice between vanilla and optimized models appeared to depend less on accuracy and more on computational constraints, latency requirements, and ease of retraining in production environments.

Figure 8: Predictions on the 1st week of the test set for car park ID = 6

Source: Own elaboration, 2025.

Across all models and car parks, hyperparameter optimization via Bayesian search yielded consistent MAE reductions compared to the vanilla configurations. However, improvements were modest, reflecting the relatively regular and predictable nature of the patterns. In some cases, the computational overhead of more complex models may outweigh their marginal accuracy gains, particularly for real-time deployment scenarios.

4.5. Implications for Deployment and Scalability

The results confirm the feasibility of applying recurrent neural networks to urban parking occupancy forecasting. While hyperparameter optimization enhances accuracy in most cases, the modest improvements—coupled with increased computational costs—suggest that vanilla models may already suffice for operational contexts, particularly where computational resources are limited.

This methodology, validated on the city of Valencia, is currently being integrated into the City Council's real-time information systems, enabling predictive occupancy data for public use. Beyond this specific case, the approach could be adapted to other cities with similar data availability, offering potential benefits in traffic management, urban mobility planning, and sustainable transport policy.

4.6. Comparison with previous works

Our findings align with and extend prior research on parking availability prediction within the broader smart-city agenda (Batty et al., 2012). Consistent with evidence that urban mobility time series are amenable to data-driven forecasting (Calabrese et al., 2011; Zheng et al., 2015), we observe highly learnable occupancy patterns across facilities, reflected in rapidly convergent loss curves and low validation errors.

Prior studies have reported that relatively simple machine learning models can, in certain contexts, outperform more complex neural network architectures (Awan et al., 2020; Inam et al., 2022). Our findings refine this perspective: baseline recurrent models (RNN, LSTM, and GRU) already achieve low MAEs—approximately 2–3% for most car parks—with Bayesian hyperparameter optimization yielding only marginal further improvements. This outcome aligns with the broader literature, which suggests that when signals exhibit strong regularity and periodicity, as is typical in parking demand, increases in architectural complexity offer diminishing returns unless complemented by richer and more diverse covariates.

Regarding LSTMs and GRUS, prior work shows they are well suited to long-range temporal dependencies in parking and traffic time series (Barraco et al., 2021; Cho et al., 2014; Hochreiter & Schmidhuber, 1997; Zhao & Zhang, 2024). Our cross-model comparison corroborates this: LSTM and GRU perform similarly and slightly better than a simpler multilayer network baseline, with stable generalization. This stability is consistent with good practice reported in the literature—bounded inputs and early stopping—which mitigates exploding/vanishing gradients and overfitting in recurrent models.

Studies leveraging contextual features (e.g., day type, time of day, weather) report that neural networks can outperform classical time-series baselines such as ARIMA/SARIMA (Sebatli, 2023), and that architectures capturing temporal—and sometimes spatial—structure (e.g., temporal convolutional networks) can further improve accuracy (Chen et al., 2023; Zhang et al., 2024). Our results, obtained with recurrent networks and the available feature set, reach low error rates without heavy architectural machinery. This suggests that, for many facilities, periodic structure dominates the signal. Nevertheless, the persistent "hard" cases (e.g., Parking 34) echo prior findings that context variability can erode performance (Yang et al., 2019; Zheng et al., 2014). In such settings, incorporating exogenous variables and/or spatial coupling (e.g., neighboring facilities, events, transit supply) is likely to close the gap noted in the literature.

Bayesian hyperparameter search brings small but consistent MAE reductions for most facilities, yet with notable sensitivity in outliers (e.g., GRU on Parking 34). This aligns with reports that improvements from advanced tuning or deeper architectures may be incremental compared to their computational cost, especially for real-time deployment—supporting the pragmatic choice of well-regularized vanilla RNN/LSTM/GRU when latency and maintainability matter.

5. AI explainability

When applying Artificial Intelligence (AI) and Machine Learning (ML) models, interpretability is crucial to understand and validate the results. Explainable AI (XAI) refers to techniques and methods that make the decision-making processes of AI systems understandable to humans.

In this work, two complementary explainability approaches were used: SHapley Additive exPlanations (SHAP) and surrogate models.

5.1. SHAP Analysis

SHAP, based on game theory, estimates the marginal contribution of each input variable to a specific prediction. It does so by evaluating all possible combinations of features and measuring the change in the model's output when each variable is included or excluded.

In our case, the model input consists of the previous 24 time instants (hours) of occupancy data. SHAP computes the contribution of all combinations of these instants to the prediction at the current time step.

Figure 9 shows the SHAP values for the LSTM vanilla model applied to car park ID 6. The results indicate that the most recent 10–12 instants have the highest influence on the prediction, in contrast to the exploratory data analysis (EDA) findings, which suggested that all 24 previous instants were relevant. This suggests that the input time window could be reduced by half without significantly degrading accuracy, thus lowering computational cost.

Figure 9: SHAP values for car park ID 6, LSTM vanilla model

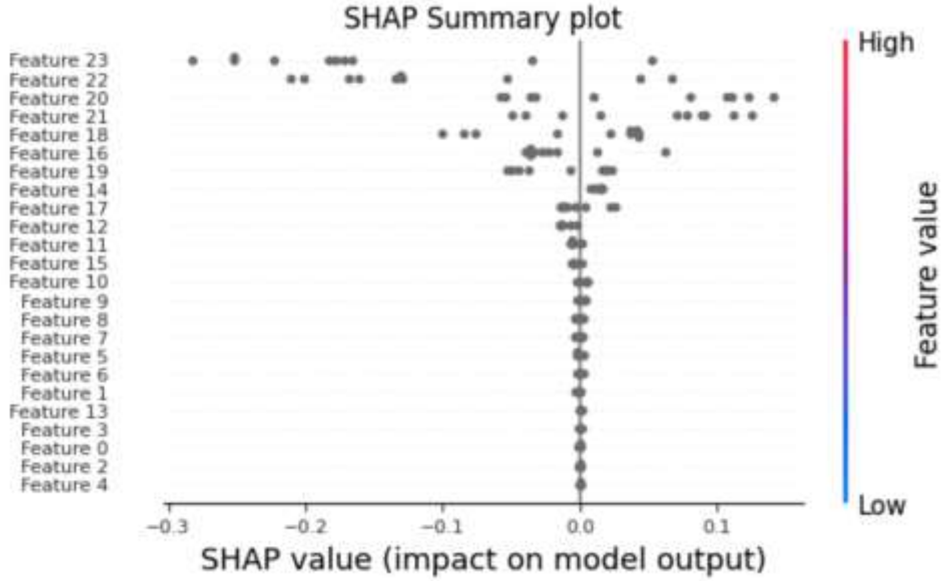

Source: Own elaboration, 2025.

Additionally, the embedding distributions for categorical variables (Figure 10) reveal that months with similar behaviours cluster together in the vector space, corroborating patterns identified in the EDA.

Figure 10: Embedding distributions for categorical variables

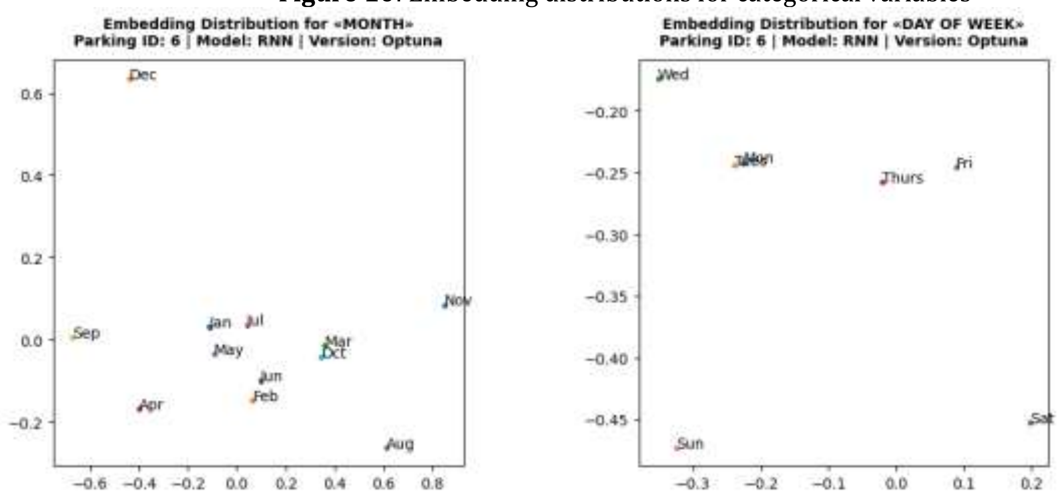

Source: Own elaboration, 2025

5.2. Surrogate Model Analysis

A surrogate model was employed to approximate the behaviour of the “black box” neural network using a simpler, interpretable model—in this case, a regression decision tree.

Let $X \in \mathbb{R}^{n \times d}$ be the input features and $y_{\text{true}} \in \mathbb{R}^n$ the true labels. The LSTM predictions are:

$$\hat{y}_{\text{LSTM}} = \text{LSTM}(X, y_{\text{true}})$$

With residuals:

$$e_{\text{LSTM}} = y_{\text{true}} - \hat{y}_{\text{LSTM}}$$

A regresion tree is trained on $(X, \hat{y}_{\text{LSTM}})$ to mimic the LSTM's predictions:

$$\hat{y}_{\text{TREE}} = TREE(X, \hat{y}_{\text{LSTM}})$$

The resulting tree (Figure 11) highlights the most influential features. The analysis confirms that the most recent time instants dominate the predictions, except during peak traffic hours (e.g., start and end of the workday), when occupancy changes are more abrupt and slightly longer temporal dependencies become relevant. These findings further support the potential benefit of reducing the temporal input window to improve computational efficiency without sacrificing accuracy.

Figure 11: Surrogate decision tree for the optimized RNN model, car park ID 6

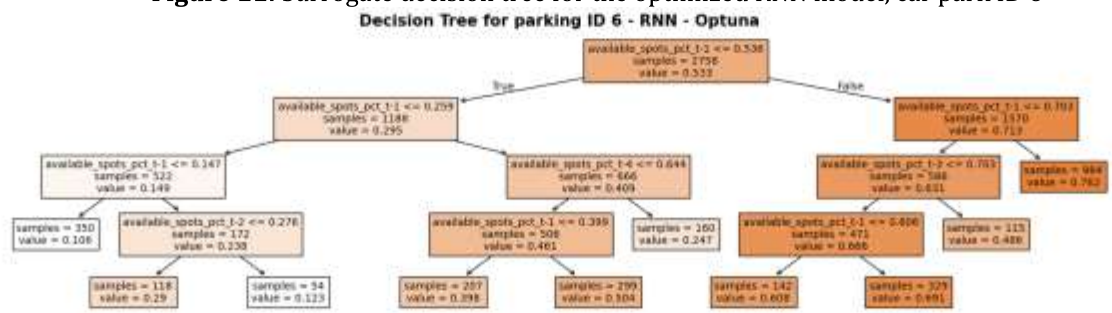

Source: Own elaboration, 2025

5.3. Limitations

The explainability analysis also highlights certain limitations of the current approach. Contextual factors such as special events, cultural activities, or road closures are not included in the model, yet they may significantly affect parking demand. Moreover, the available data are not uniformly distributed across all areas of the city, which could affect the generalizability of the explainability findings.

6. Conclusions and future work

This study demonstrates the feasibility of using machine learning techniques to predict the availability of parking spaces in urban environments, achieving a level of precision that supports integration into smart mobility applications. The developed system, applied to public car parks in the city of Valencia, is based on recurrent neural network architectures and incorporates both optimized and baseline configurations.

An exhaustive process of data acquisition and preprocessing was carried out, identifying and addressing multiple limitations that could have compromised prediction quality. This preliminary work proved essential to ensuring the robustness and reliability of the final system.

The evaluation of three recurrent architectures—RNN, LSTM, and GRU—showed that all are capable of capturing cyclical occupancy patterns with high accuracy. While hyperparameter tuning via Bayesian optimization yielded modest improvements, even the vanilla configurations performed satisfactorily, striking a balance between predictive accuracy and computational cost.

Our results are congruent with the previous works showing that recurrent models are a strong default for parking occupancy time series and that simple, well-regularized baseline can be sufficient for most facilities.

Explainability techniques, including SHAP analysis and surrogate models, provided valuable insights into the decision-making process of the models, clarifying the influence of temporal and categorical variables on the predictions. These analyses also revealed opportunities to reduce the temporal input window without significantly degrading performance, potentially lowering computational requirements.

Despite the low prediction errors achieved, model performance depends heavily on the quality and temporal resolution of the input data. Incomplete or irregularly updated datasets, as well as the absence of contextual variables such as weather conditions, special events, or traffic incidents, may reduce accuracy.

The Valencia case study contributes empirical evidence from a European, open-data setting where smart-mobility infrastructure is being progressively deployed.

Future work will explore hybrid modelling strategies that integrate recurrent neural networks with external data sources to capture a wider range of influencing factors. Moreover, transfer learning approaches could also facilitate deployment in cities with limited historical data. Additionally, real-world evaluations of the system's impact on driver behaviour and urban mobility will be conducted to refine both the models and their practical applications.

In the medium term, the system could evolve into an automated MLOps pipeline capable of integrating new data in real time, continuously optimizing model parameters, and streamlining deployment in production environments. The techniques and methodologies developed here are readily extensible to other domains of mobility and urban services, offering a scalable foundation for future smart city initiatives.

7. Acknowledgements

The present work has been developed within the framework of the Chair funded by the Valencian City Council: "Cátedra de Inteligencia Artificial y Visualización de Información para el Gobierno Abierto en la Ciudad de Valencia"

References

Awan, F. M., Saleem, Y., Minerva, R., & Crespi, N. (2020). A comparative analysis of machine/deep learning models for parking space availability prediction. *Sensors*, *20*(1), 322. https://doi.org/10.3390/s20010322

Barraco, M., Bicocchi, N., Mamei, M., & Zambonelli, F. (2021). Forecasting parking lots availability: Analysis from a real-world deployment. *IEEE International Conference on Pervasive Computing and Communications Workshops* (pp. 299–304). IEEE. https://doi.org/10.1109/PERCOMWORKSHOPS51409.2021.9430942

Batty, M., Axhausen, K. W., Giannotti, F., Pozdnoukhov, A., Bazzani, A., Wachowicz, M., Ouzounis, G., & Portugali, Y. (2012). Smart cities of the future. *The European Physical Journal Special Topics, 214*(1), 481–518. https://doi.org/10.1140/epjst/e2012-01703-3

Calabrese, F., Colonna, M., Lovisolo, P., Parata, D., & Ratti, C. (2011). Real-time urban monitoring using cell phones: A case study in Rome*. IEEE Transactions on Intelligent Transportation Systems*, *12*(1), 141–151. https://doi.org/10.1109/TITS.2010.2074196

Chen, G., Zhang, S., Weng, W., & Yang, W. (2023). Residual spatial-temporal graph convolutional neural network for on-street parking availability prediction. *International Journal of Sensor Networks, 43*(4), 246–257. https://doi.org/10.1504/IJSNET.2023.135840

Cho, K., van Merrienboer, B., Bahdanau, D., & Bengio, Y. (2014). On the properties of neural machine translation: Encoder–decoder approaches. *Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8*) https://doi.org/10.48550/arXiv.1409.1259

Hasan, S. (2020, June 26). *Recurrent Neural Network and it's variants*. Medium. https://medium.com/@hasanshujaat4/recurrent-neural-network-and-its-variants-de75f9ee063

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Inam, S., Mahmood, A., Khatoon, S., Alshamari, M., & Nawaz, N. (2022). Multisource data integration and comparative analysis of machine learning models for on-street parking prediction. *Sustainability, 14*(12), 7317. https://doi.org/10.3390/su14127317

Junta de Andalucía & Telefónica. (2021, July). *Un sistema inteligente basado en Big Data permite conocer en tiempo real las plazas libres de aparcamiento en el PCT Cartuja*. APTE / PCT Cartuja. https://www.apte.org/sistema-inteligente-basado-big-data-permite-conocer-tiempo-real-plazas-libres-aparcamiento-pct-cartu/

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, 4765–4774. https://doi.org/10.48550/arXiv.1705.07874

Moreira-Matias, L., Gama, J., Ferreira, M., Mendes-Moreira, J., & Damas, L. (2013). Predicting taxi–passenger demand using streaming data*. IEEE Transactions on Intelligent Transportation Systems, 14*(3), 1393–1402. https://doi.org/10.1109/TITS.2013.2262376

Moreno, D. (2021). *Predicting the use of car parks in the province of Barcelona* [Bachelor's thesis, Universitat Pompeu Fabra]. UPF Repository. http://hdl.handle.net/10230/49219

Pozo, R.F., González, A. B. R., Wilby, M. R., Vinagre Díaz, J. J., & Matesanz, M. V. (2022) Prediction of on-Street parking level of service based on random undersampling decision trees. *IEEE Transactions on Intelligent Transportation Systems*, *23*(7), 8327-8336. https://doi.org/10.1109/TITS.2021.3077985

Sebatli, A., & Çavdur, F. (2023). Prediction of parking space availability using ARIMA and neural networks. *Endüstri Mühendisliği Dergisi, 34*(1), 86–108. https://doi.org/10.46465/endustrimuhendisligi.1241453

Shoup, D. C. (2005). *The high cost of free parking*. Planners Press, American Planning Association.

Transformative Mobility SUTP. (2024). *Gestión de estacionamientos: Una contribución hacia la movilidad sostenible*. GIZ / SUTP. https://transformative-mobility.org/wp-content/uploads/2024/01/GIZ_SUTP_SB2c_Parking-Management_ES-2.pdf

Vieta, G. L. (2024). *Modelos predictivos para ocupación de estacionamiento* [Master's thesis, Universidad de Sevilla]. idUS Repository. https://hdl.handle.net/11441/161136

Yang, S., Ma, W., Pi, X., & Qian, S. (2019). A deep learning approach to real-time parking occupancy prediction in transportation networks incorporating multiple spatio-temporal data sources. *Transportation Research Part C: Emerging Technologies, 107*, 248–265. https://doi.org/10.1016/j.trc.2019.08.010

Zhang, F., Shang, K., Yan, L., Nan, H., & Miao, Z. (2024). Prediction of parking space availability using improved MAT-LSTM network. *ISPRS International Journal of Geo-Information, 13*(5), 151. https://doi.org/10.3390/ijgi13050151

Zhao, X., & Zhang, M. (2024). Enhancing predictive models for on-street parking occupancy: Integrating adaptive GCN and GRU with household categories and POI factors. *Mathematics*, *12*(18), 2823. https://doi.org/10.3390/math12182823

Zheng, Y., Capra, L., Wolfson, O., & Yang, H. (2014). Urban computing: Concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology, 5*(3), 1–55. https://doi.org/10.1145/2629592

Zheng, Y., Rajasegarar, S., & Leckie, C. (2015). Parking availability prediction for sensor-enabled car parks in smart cities. In *IEEE Tenth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)* (pp. 1–6). IEEE. https://doi.org/10.1109/ISSNIP.2015.7106955