Real-World Indoor Location Assessment with unmodified RFID antennas

Pedro Sobral[1,2][0009−0009−8683−8157], Rui Santos[1][0000−0003−2157−7602], Ricardo Alexandre[2][0000−0001−8836−0139], Pedro Marques[1][0000−0001−5656−1817], Mário Antunes[1,2][0000−0002−6504−9441], João Paulo Barraca[1,2][0000−0002−5029−6191], João Silva[3], and Nuno Ferreira[3]

[1] Departamento de Eletrónica, Telecomunicações e Informática
Universidade de Aveiro, Aveiro, Portugal
{sobral,ruipedro99,pedromm,mario.antunes,jpbarraca}@ua.pt
[2] Instituto de Telecomunicações, Universidade de Aveiro, Aveiro, Portugal
{sobral,rjfae,mario.antunes,jpbarraca}@av.it.pt
[3] Think Digital, Aveiro, Portugal
{jsilva,nferreira}@thinkdigital.pt

Abstract. The management of health systems has been one of the main challenges in several countries, especially where the aging population is increasing. This led to the adoption of smarter technologies as means to automate, and optimize processes within hospitals. One of the technologies adopted is active location tracking, which allows the staff within the hospital to quickly locate any sort of entity, from key persons to patients or equipment. In this work, we focus on exploring ML models to develop a reliable method for active indoor location tracking based on off the shelf RFID antennas with UHF passive tags. The presented work describes the full development of the solution, from the initial development made within a controlled environment, to the final evaluation made on a real health clinic. The proposed solution was able to achieved 0.47 meters on average on a complex medical environment, with unmodified hardware.

Keywords: Indoor Location · Machine Learning · Passive RFID tag · Regression Models

1 Introduction

A fundamental governance concern is the administration of healthcare systems and their economic viability, particularly in nations dealing with an increasingly sizable ageing population. Hospitals and clinical centres are actively looking for novel approaches to improve service efficiency, cut costs, and improve patient happiness due to the soaring patient numbers and the difficulties in matching annual budgets with this expansion. Some hospitals are starting to use active localization technology like Radio Frequency Identification (RFID) Radio Frequency Identification (RFID) [6,11] to expedite these procedures. Utilising Radio Frequency (RF) technologies makes it possible to track the location and use of medical equipment, keep tabs on the inventory and distribution of patient prescriptions, and even keep track of patient mobility across hospital departments.

The main goal of this study was to look into various approaches for indoor resource localization using Ultra High Frequency (UHF) passive tags while sticking to certain restrictions. Rapid response times, pinpoint location accuracy, minimal data requirements for each prediction (to reduce power consumption by components of the RFID radios), and, most importantly, compatibility with readily available, unaltered off-the-shelf hardware were all requirements that the models developed in this research had to meet. This study is a component of a larger initiative to create an inexpensive indoor positioning system using widely available RFID antennas and tags.

In this work, we studied how Machine Learning (ML) models can be used to predict asset locations by utilising passive tags. Our study expands on earlier study that examined related hardware in a controlled environment and produced encouraging results [8]. The main innovation in this study focuses on the evaluation of the previous approach using actual data acquired from a medical clinic. In addition, we have added improvements to the system's functionality that have been verified through practical testing in a clinical setting.

The remaining document is organized as follows. Section 2 described the current state of the art for active indoor location. The following section (Section 3) describes the hardware used in the execution of this study. Section 4 presents the previous proposed solution, that serves as the basis for this work. In Section 5 we describe the real-word scenario where the antennas where deployed. The new models (and the main contribution of this work) are presented in Section 6. Finally, the conclusion can be found on Section 7.

2 State of the Art

Due to their wide range of technological capabilities and the significant value they provide to numerous business sectors, Indoor Positioning System (IPS) have seen a rise in popularity over the previous two decades.

There are multiple techniques used in location systems, such as multilateration [1], angulation [7], fingerprinting [10] and others. These techniques require some information provided by the antennas and tags used, like the Time of Arrival (ToA) [9], Angle of Arrival (AoA) [15] and Received Signal Strength Indication (RSSI) [4]. The tags used in these systems can be active or passive, but they often require unique design and are not widely available in products from the main stream market. They are mostly present in specialised products and research projects. We will go into more detail about the most pertinent works that offer remedies for situations that are similar in the parts that follow.

SpotON [4] is a location system that uses RSSI to locate active RFID tags in a three-dimensional space. LANDMARC [5] is a system that also reflects the relationship between RSSI and power levels, and makes use of reference tags and the K-NN algorithm to estimate positions. Results show an accuracy of 2 m and a location delay of 7.5 s. In [10] Dwi *et al.* propose a fingerprinting based positioning system using a Random Forest (RF) algorithm and RSSI data, which achieved an error of 0.5 m, which is 18% lower than the compared Euclidean distance method. In [2], Lummanee et al. compare the performance of a Gradient Boosting algorithm to a typical Decision Tree (DT) applied in a positioning system. The experiment was based on a 324 m^2 area divided in 9 zones. The DT based Gradient Boosting algorithm achieved an estimation error of 0.754 m for 19 reference radio signals at 50 samples per zone, 17.8% more accurate than the typical DT. In [3] Jae *et al.* developed a passive RFID based localization system which uses RSSI information and reference tags to predict one-dimensional position of the asset. It achieves an error of 0.2089 m using the K-NN technique in a 3 m space.

These methods are mainly based on RSSI, which has the disadvantage of suffering greatly from attenuation due to internal obstacles and dynamic environments. Unlike SpotON and LANDMARC, the approach in [14] by Wilson *et al.* does not depend on RSSI, however, is based on the same RSSI principles. This research work is based on passive RSSI technology. Two scenarios of stationary and mobile RSSI tags are considered. The method gives tag count percentages for various signal attenuation conditions. The tags are located by recording characteristic curves of readings under different attenuation values at multiple locations in an environment. Similarly, Vorst et al. [12] use passive RFID tags and an onboard reader to locate mobile objects. Particle Filter (PF) technique is exploited to estimate the location from a prior learned probabilistic model, achieving a precision of 0.20–0.26 m.

By carefully designing and validating ML models, we in our work take into consideration these earlier research initiatives and increase the precision of our earlier work.

3 RFID Antenna Description

Given the environment where this work has developed, the hardware was preselected (as seen in Figure 1). The antenna part is composed of two units (processing + radio), that communicate with each other through a physical bus (RS232, RS485 or Ethernet). The local processing unit was designed to communicate through Long Term Evolution (LTE) and Ethernet, exposing the radio frontend to the backend systems. It is powered through a 230 V Alternating Current (AC) power supply and uses an Advanced RISC Machine (ARM) processor running a GNU/Linux operating system (for low power consumption).

The antenna model used in our system is pretty typical and is used exactly as the manufacturer intended it to be, without any special firmware alterations. This decision lowers the cost of the system while also improving usability and accessibility. But there is a catch: the antenna's processing power or the data it transmits might not be at their best. The vendor claims that the firmware level automatically compensates for gain in order to aid in the detection of passive tags. However, this compensation has a substantial impact on how well it fits our planned scenario. Our work's main goal is to offer value by providing an effective indoor positioning solution, even when using unmodified hardware. It is important to keep in mind that, in the hypothetical situation, more reliable alternatives utilising active tags quickly exceed the average cost per patient, making them unworkable.

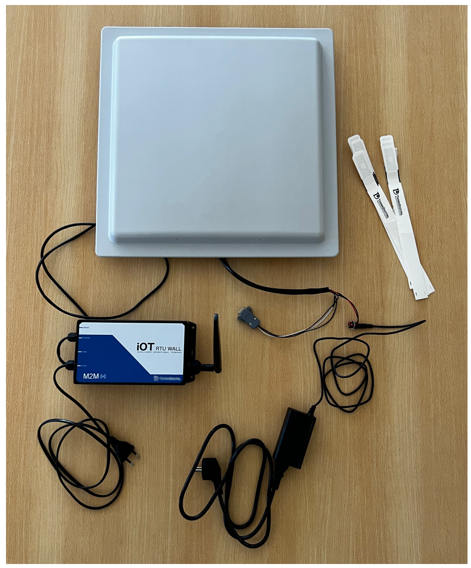

Fig. 1: Smart Antenna and passive RFID tags used for data gathering [8].

The communication between the antenna and the tags is made through a carrier wave in the 865–868 MHz (UHF) frequency range as defined in the EN 302 208 v3.2.0 [4] directive for the European region, and cannot exceed 2 W emission power. In this way, the antenna controller allows the RF emission power adjustment 0–300 mW, allowing readings up to 25 m and writings up to 6 m according to the manufacturer. The antenna polarization is circular with a gain of 12 dBi. The controller uses the Impinj R2000 chipset supporting the EPC C1 GEN2 protocol [5], ISO18000-6C [6] (see Table 1). This setup should be one of the most commonly used, as the hardware and chipset are commonly used for similar tasks. We see this as a major contribution from our work, as the output can be applied to a wide set of existing or future, deployments.

The real-time communication with the smart antenna is achieved using an MQTT broker, over which we implemented a key set of control functions: a) Definition of emission power of the antenna; b) Tag reading request over a time

[4] `https://www.etsi.org/deliver/etsi_en/302200_302299/302208/03.02.00_20/en_302208v030200a.pdf`
[5] `https://www.gs1.org/standards/rfid/uhf-air-interface-protocol`
[6] `https://www.iso.org/standard/59644.html`

Table 1: Specification of the RFID UHF reader and writer [8].

Product Parameter	Parameter Description
Model	ACM818A UHF (20M)
Tag Protocol	EPC C1 GEN2 \ ISO 18000-6C
Output Power	Step interval 1.0dB, maximum +30dBm
RF Power Output	0.1W - 1W
Built-in Antenna Type	12dbi linear polarization antenna
Communication Ports	1) RS-232 2) RS-485 3) Wiegand 26 \ 32 bits
Communication Rates	115200bps
Reading/Writing	20m
Multi-tags Reading	200tags/s
Working Voltage	DC +12V

window (burst); c) Return data obtained at the end of the reading; d) Direct interaction with an antenna to manage it.

4 Previous Solution

This works expands on a previous one [8], which explored the same hardware but on a controlled environment. In this section we present a brief summary of the results achieved from our previous solution, which we expanded in this new work. Both works share the same environment, requirements, and the approach of using ML to provide usefulness to off the shelf hardware.

To achieve the previous solution we run three experiments to determine the best input features and ML model for location prediction.

Regarding the input features, we considered two different ones, apart from the typical RSSI, the number of tag activation's and the average time between activation's were also considered. Our results showed that these input features were robust and improved greatly the precision of the location solution.

Furthermore, we explored the usage of anchor points (static tags with well known location) for model evaluation and continuously training in a dynamic environment.

Finally, we were able to establish the baseline values to the models precision on a controlled environment. Our results were aligned with the state of the art, the solution obtained an error of 0.00 m within a range of 5 m and an error of 0.55 m within a range of 10 m, resulting in an average error of 0.275 m.

5 Scenario

As previously mentioned, the previous solution was developed based on three experiments conducted in a controlled environment. This works is aligned with a National Project, as such we had access to a real health center were five antennas, and eighteen anchor were deployed. Figure 2) depicts the location of the antennas and anchors in our real-world scenario.

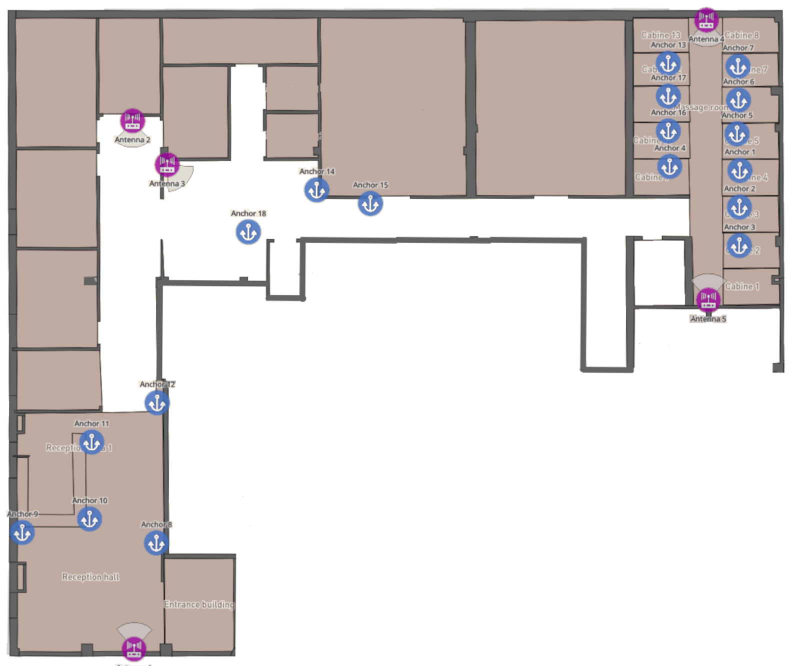

Fig. 2: The location of the antennas and anchor tags.

With this placement, the next step was to choose the best powers to generate our dataset. We want to maximize the model precision, but also wanted to increase the maximum number of tags detected. Due to latency and processing constrains we are limited to query with two power levels. After emitting a signal at the specified power level, we wait three seconds for tags responses. Querying more than two power levels would increase the latency of the system, which would impact the remaining components. Moreover, the connection system will need to store all tag activations detected during the scanning period, which presents a processing, memory and communication burden.

To achieve this analysis we collected data from all power levels for three days straight. We created a dataset with each single power level, and all the combinations of two power levels. Using those datasets we trained the best model from the previous work (Random Forest), and computed the Mean Absolute Error (MAE) for each dataset. Since this is a regression problem, we want to minimize the MAE, and at the same time maximize the number of tags detected. As such, the metrics score that we considered in this analysis is the subtraction the normalized MAE from normalized count of tags.

Figure 3 depicts the results of our analysis. In the diagonal, we can see the values of single power levels, while the remaining cells contains the metric value for each pair of power levels. It is possible to observe, that the pair $[210, 300]$mW offers the lowest error, while it captures the largest number of tags. Effectively, 300mW will get all tags possible (max power), while 210mW will provide distance discrimination.

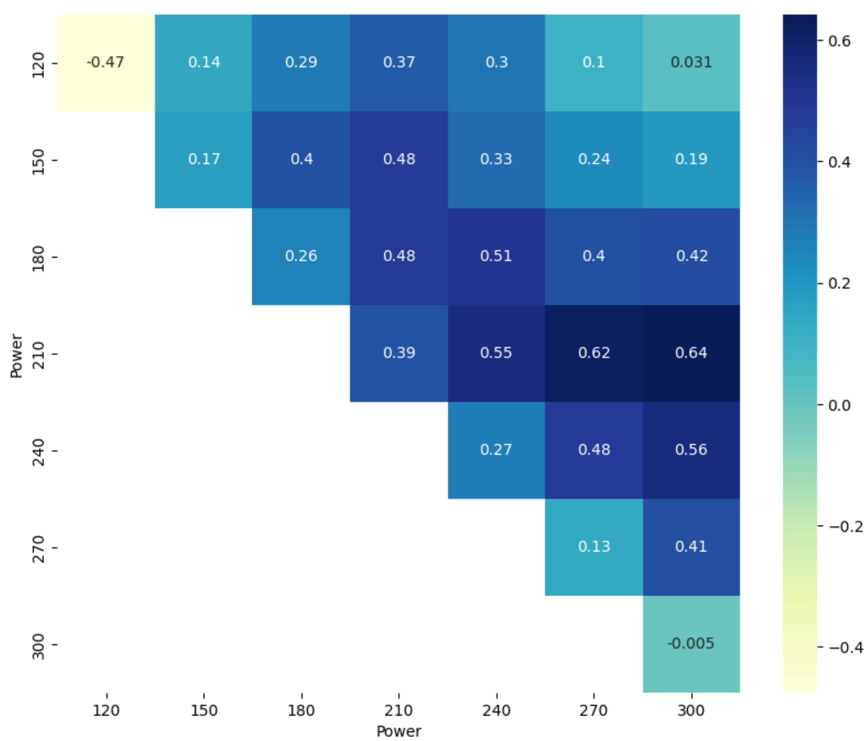

Fig. 3: Result of the power levels analysis impact on performance

The dataset that will be used to train the ML models are distributed in the following way (as depicted in Figure 4), and had 68368 rows, each row representing a tag that was read by some antenna.

As shown in Figure 4, the distribution of samples is not uniform for all distances, which is not the ideal, as it introduces a bias within the models. However, the major of the distances have a very good distribution (between 6.5%

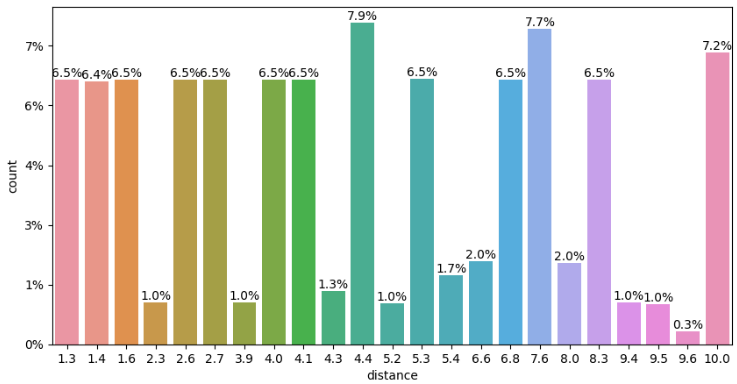

Fig. 4: Dataset Distribution

and 7.9%). To recap, we placed eighteen anchors that are able to produce twenty three different distances. This is possible as there are some overlaps between antennas, increasing the number of effective distances for a given set of anchor tags.

6 Evolved Solution

We started by applying the ML models from the previous work, in the real-world scenario, as it showed capabilities without our requirements. These models would serve as a baseline performance for the system in a real environment. Given that there was a considerable increase in the models error, we applied two different pre-processing methods to clean the dataset and improve the model capabilities.

The first pre-processing method was simple applying Interquartile Range (IQR) to filter out the outliers. IQR is a measure of statistical dispersion, it is an example of a trimmed estimator, which improves the quality of the of dataset by dropping lower contribution samples.

The second pre-processing method implies the usage of an AutoEncoder (AE) to compress the original dataset into a smaller one, by reducing the number of input features per sample. The main advantage of this approach is that it does not removes any samples from the dataset (opposite to the outlier removal method from the previous approach). The main disadvantage is the necessity of training the AE.

These pre-processing approaches were taken into account due to the significant level of noise present in the dataset, which was not present in the controlled environment. The noise level in the real-world dataset can be explained by two factors. First, is the clinical environment itself. As a working environment, it contains equipment, its own Wi-Fi network, and other components that can interfere with the propagation of the signal emitted by the localization antennas.

Second, is related to the utilization of tags by the staff and users of the space. Although instructions were provided on how to use the tags, there are no guarantees that the tags will always be in the optimal positions for their reading. Actually, they will almost never be in such situation, as they reside in patient and staff members pockets, purses or backpacks.

In order to mitigate this noise, we implemented these additional approaches with the aim of refining the data quality, while minimizing the impact of external influences.

Three metrics were used to evaluate the performance of the models: Mean Squared Error (MSE), MAE, and Coefficient of Determination (R2). The first two are dissimilarity metrics, meaning that the lower the value, the better the result. The later is a similarity metric, meaning that the closer it is to 1, the better the result is.

6.1 Baseline

The first step was to evaluate the models previously developed. This means, using the models without applying any additional process to handle the new dataset. The results of the models can be found in Table 2.

Table 2: Baseline performance metrics

Model	MSE	MAE	R2
LinearRegression	5.62	1.86	0.23
DecisionTree	1.64	0.57	0.78
RandomForest	2.23	1.03	0.70
KNeighbors	1.87	0.59	0.75
XGBoost	1.50	0.58	0.80
Voting Model	2.01	0.98	0.73

As presented in Table 2 the best possible value is 0.57 (MAE) achieved by the DT model. This represents a performance decrease compared to the best model from the previous work by an average of 0.02 m.

To better understand the model perform at different ranges, we computed the MAE for the best model (DT) at each possible distance and plotted them in the following histogram (see Figure 5).

Through an analysis of the figures, we could observe that the model tends to make more errors in distances where there is less data, for distances greater than approximately 4 m. Comparing this with the distribution of the dataset (Figure 4), it is quite noticeable.

6.2 First pre-processing: IQR

The first approach to reduce the natural noise in the dataset was to apply the IQR outlier filtering method. This method calculates all the values between

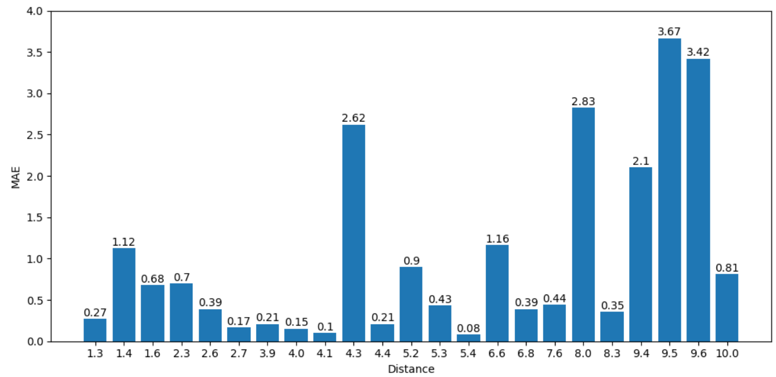

Fig. 5: MAE histogram at different ranges from the best baseline models (DT).

the first quartile (Q1) and the third quartile (Q3), giving more importance to the inner 50% of the data and eliminating the majority of outliers. The constant value used was 1.5, which is used to establish a threshold that determines how far the data can be considered "far" from the median. This value (1.5) is based on Tukey's fence, allowing for sensitive identification of potential outliers while remaining relatively conservative. The idea behind exploring the removal of outliers (data points with significantly different characteristics from what is considered normal) is to obtain a cleaner dataset.

With this approach pre-processing, we obtain the results presented in Table 3.

Table 3: IQR performance metrics

Model	MSE	MAE	R2
LinearRegression	4.19	1.60	0.44
DecisionTree	1.27	0.47	0.83
RandomForest	1.88	0.94	0.75
KNeighbors	1.60	0.51	0.79
XGBoost	1.17	0.48	0.84
Voting Model	1.61	0.83	0.79

As we can see in table 3 the best possible value is 0.47 (MAE) achieved by the DT model. This represents a performance increase compared to the best model from the initial training by an average of 0.1 meters.

To better understand the model with the highest perform in this approach (DT) at different ranges, we computed the MAE for each possible distance and plotted them in the following histogram (see Figure 6).

As in the previously approach (Baseline), it is evident that the model exhibits a higher error rate in distances with few data, particularly for distances exceeding

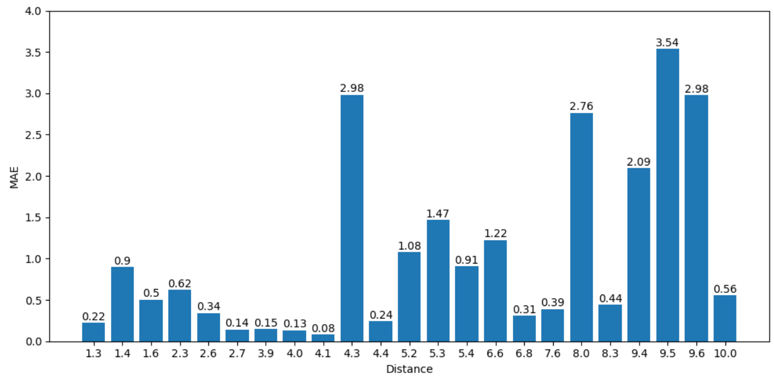

Fig. 6: MAE histogram at different ranges from the IQR best models (DT).

approximately 4 m. This is aligned with the distribution of samples per range of the acquired dataset (as seen in Figure 4).

6.3 Second pre-processing: AutoEncoder

We also implemented an AutoEncoder (AE) [13] for feature reduction. An AutoEncoder is a neural network that compresses the input data into a lower-dimensional representation, and then reconstructs it from that lower dimension.

This network is composed of two different blocks, the encoder and the decoder. The encoder is a initial part of the network that compresses the input into a latent space with a smaller dimension. The decoder tries to reconstruct the input signal from the constrained latent space. The network is trained on the input data, while trying to reproduce it using a constrained latent space. The loss function is the difference between the input features and the reconstructed features. After an initial training phase, we only use the encoder part of the network to compress the input features into a lower-dimensional called latent space. One consequence of compressing the input features into a lower-dimensional latent space is the reduction of noise in the input signals.

After training the AutoEncoder, we simple use the encoder part of the network to transform the dataset into a new one on the constrained latent space. As previously mentioned, with reduced noise levels. After we just apply the ML models naturally.

With this pre-processing method, we obtained the results presented in Table 4.

As we can see in Table 4 the best possible value is 0.64 (MAE) achieved by the K-Nearest Neighbors (K-NN) model. This represents a performance decrease compared to the best model from the initial training by an average of 0.07 m, and a performance decrease compared to the results in the IQR approach by 0.17 m.

Table 4: AutoEncoder performance metrics

Model	MSE	MAE	R2
LinearRegression	5.70	1.93	0.22
DecisionTree	2.71	1.02	0.64
RandomForest	3.15	1.27	0.57
KNeighbors	2.05	0.64	0.72
XGBoost	2.05	0.84	0.72
Voting Model	2.49	1.14	0.66

In order to better understand how the best model performs (K-NN model) in the existing ranges, we plotted it (see Figure 7), similarly to the previous approaches.

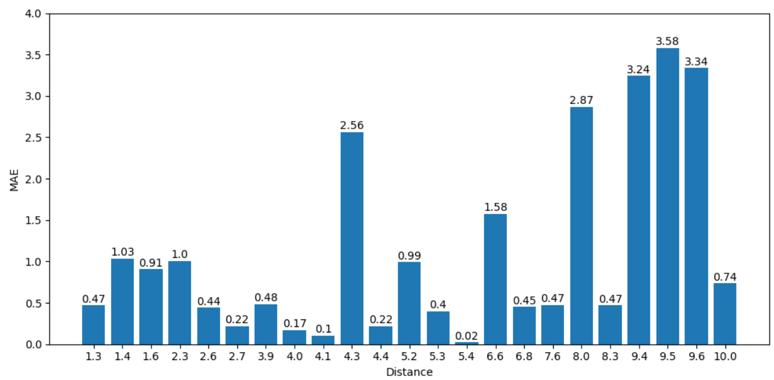

Fig. 7: MAE histogram at different ranges from the AE best model (K-NN).

Similarly to the previous two approaches (baseline and IQR) it is apparent that the model demonstrates increased error rates in ranges where training data is limited. Again, especially visible for distances greater than approximately 4 m, and aligned with the distribution of samples per range of the acquired datasets.

7 Conclusion

The management of health systems has been one of the main challenges in several European countries, especially where the ageing population is increasing. One of the technologies adopted is active location solutions, which allows the staff within the hospital to quickly find any sort of entity, from key persons to equipment.

In this work, we evaluated the usage of dedicated hardware (namely the RF antenna) for indoor location within a medical environment. From a previous work we devise two new feature (number of activation's and the average time of activation's for passive tags) that in conjunction with RSSI produce models that are more robust. We also established a performance based, based on data acquired on a controlled environment.

The current work is aligned with collaboration with medical entities, and as such we were able to acquired data in a real-world health clinic.

In this work, we present the details of the scenario, and how we selected the ideal power levels to explore with the antennas. It is a compromise between the amount of data acquired from the antennas and the latency of the overall system.

Following that, we evaluated the previous solution to compare the current results with the previous ones. Simply applying the previous models lead to a decrease of 0.02 meters in performance, when taking into account the best ML models.

Our analysis show an increase noise level within the dataset. That noise is a consequence of the environment itself, arising from moving people, sub-optimal tag location and orientation and othre interference. To deal with this we applied two different pre-processing approaches: outlier removal with IQR and noise reduction with AutoEncoder.

The approach based on IQR was able to improve the results, achiving an MAE of 0.47. While the approach based on AutoEncoder did not improve the results. It is worth mentioning that further research on this is required, as the IQR based method reduces the dataset size (rows), while the AutoEncoder transforms it into another dimensional space (reduces the number of columns). Is is possible, that with a larger dataset, we could see some gains on the approach based on AutoEncoder.

Finally, we want to mention that for all the proposed approaches, the MAE per range was correlated with the amount of data that we had for that specific range. This can lead us to increase the number of passive anchors (tags with well known location) to improve the continuous training of the models.

Acknowledgements This work is supported by the European Regional Development Fund (FEDER), through the Competitiveness and Internationalization Operational Programme (COMPETE 2020) of the Portugal 2020 framework [Project SDRT with Nr. 070192 (POCI-01-0247-FEDER-070192)]

References

1. Carotenuto, R., Merenda, M., Iero, D., Della Corte, F.G.: An indoor ultrasonic system for autonomous 3-d positioning. IEEE Transactions on Instrumentation and Measurement **68**(7), 2507–2518 (2019)
2. Chanama, L., Wongwirat, O.: A comparison of decision tree based techniques for indoor positioning system. In: 2018 International Conference on Information Networking (ICOIN). pp. 732–737 (2018)

3. Choi, J.S., Lee, H., Elmasri, R., Engels, D.W.: Localization systems using passive uhf rfid. In: 2009 Fifth International Joint Conference on INC, IMS and IDC. pp. 1727–1732 (2009)
4. Hightower, J., Vakili, C., Borriello, G., Want, R.: Design and calibration of the spoton ad-hoc location sensing system. unpublished, August **31** (2001)
5. Ni, L., Liu, Y., Lau, Y.C., Patil, A.: Landmarc: indoor location sensing using active rfid. In: Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, 2003. (PerCom 2003). pp. 407–415 (2003)
6. Paiva, S., Brito, D., Leiva-Marcon, L.: Real time location systems adoption in hospitals–a review and a case study for locating assets. Acta Scientific Medical Sciences **2**(7), 02–17 (2018)
7. Pomárico-Franquiz, J., Khan, S.H., Shmaliy, Y.S.: Combined extended fir/kalman filtering for indoor robot localization via triangulation. Measurement **50**, 236–243 (2014)
8. Santos., R., Alexandre., R., Marques., P., Antunes., M., Paulo Barraca., J., Silva., J., Ferreira., N.: Towards improved indoor location with unmodified rfid systems. In: Proceedings of the 12th International Conference on Pattern Recognition Applications and Methods - ICPRAM. pp. 156–163. INSTICC, SciTePress (2023)
9. Shen, H., Ding, Z., Dasgupta, S., Zhao, C.: Multiple source localization in wireless sensor networks based on time of arrival measurement. IEEE Transactions on Signal Processing **62**(8), 1938–1949 (2014)
10. Suroso, D.J., Rudianto, A.S., Arifin, M., Hawibowo, S.: Random forest and interpolation techniques for fingerprint-based indoor positioning system in un-ideal environment. International Journal of Computing and Digital Systems (2021)
11. Tegou, T., Kalamaras, I., Votis, K., Tzovaras, D.: A low-cost room-level indoor localization system with easy setup for medical applications. In: 2018 11th IFIP Wireless and Mobile Networking Conference (WMNC). pp. 1–7 (2018)
12. Vorst, P., Schneegans, S., Yang, B., Zell, A.: Self-localization with rfid snapshots in densely tagged environments. In: 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 1353–1358 (2008)
13. Wang, W., Huang, Y., Wang, Y., Wang, L.: Generalized autoencoder: A neural network framework for dimensionality reduction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2014)
14. Wilson, P., Prashanth, D., Aghajan, H.: Utilizing rfid signaling scheme for localization of stationary objects and speed estimation of mobile objects. In: 2007 IEEE International Conference on RFID. pp. 94–99 (2007)
15. Xiong, J., Jamieson, K.: ArrayTrack: A Fine-Grained indoor location system. In: 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13). pp. 71–84. USENIX Association, Lombard, IL (Apr 2013)