Deep Learning based Failure Prognostic Model for Aerospace Propulsion Systems

Jose Nava, Alberto Ochoa

Universidad Autónoma de Ciudad Juárez,
Mexico

jose.nava@3pillarglobal.com, alberto.ochoa@uacj.mx

Abstract. The continuous growth of complexity in aerospace systems is making it increasingly difficult to analyze and process information from sensors in propulsion systems in real-time to obtain a good model for predicting failures in aviation turbofan units. In this work, a model of Neural Networks with deep learning to predict the Remaining Useful Life (RUL) in aerospace propulsion systems is proposed. Specifically, a Turbofan (jet or turbojet engine) and its set of sensors are analyzed to represent and mathematically predict the evolution of the failure state and engine degradation. A data set from NASA (National Aeronautics and Space Administration) has been used, found in the NASA Prognostics Center of Excellence (PCoE) repository at the Ames Research Center. This data set has been generated with the Commercial Modular Aero-Propulsion System Simulation (C-MAPSS) software, which is a dynamic model created by this institution. In addition, the TensorFlow platform has been used to program pre-processing, analysis, and create a Deep Learning model. The importance of this model goes beyond its usefulness in the aerospace industry, since the main functions of a gas turbine such as air compression, combustion, and energy recovery, allow these systems to be used for multiple purposes; therefore, this model can be replicated in gas turbines in other industries such as thermoelectric plants, petrochemical plants, in propulsion in the military and maritime industries, and in cargo ships.

Keywords: Remaining Useful Life (RUL), Long Short-Term Memory (LSTM), flight envelope.

1 Introduction

1.1 Background

In Artificial Intelligence, machine learning and deep learning can automatically learn hierarchical representations of data on a large scale, which makes these, effective tools for failure prediction applications within predictive maintenance, especially in the presence of industrial multidimensional and high-volume data. Traditional data-driven mathematical strategies require manual feature extraction and specific feature selection

5

processes, which is highly dependent on programmers' experience and knowledge of signal processing in electronics [12]. Furthermore, aviation systems currently send information directly from fuselage sensors to Electronic Flight Instruments (EFI), and these instruments only evaluate the information from the sensors within a predetermined range, that is, they do not analyze correlation information or patterns between the different dimensions to implement multivariate analysis in real time and predict effectively.

Conventional frameworks cannot be updated in real time and require a lot of work when dealing with large-scale data sets. In comparison, a deep learning algorithm makes it possible to integrate tasks such as feature extraction, feature selection, and regression into a dynamic, hybrid architecture, making automation of such complex predictions ideal for these systems. There are ongoing Aerospace Hardware product development efforts in the industry where it is possible to implement the model proposed in this work.

### 1.2 Aerospace Propulsion Systems

Propulsion by gas exhaust (or jet propulsion) can be defined as the force that is generated in the opposite direction to the expulsion of the gases. In a turbine engine, the intake, compression, combustion, and exhaust functions take place in the same combustion chamber.

Consequently, each of these functions must have an exclusive occupation of the chamber during its respective part of the combustion cycle [9]. An important feature of the gas turbine engine is that, in its design, separate sections are dedicated to each function, and all functions are performed simultaneously without interruption, hence the need for intensive sensors monitoring [4].

2   NASA C-MAPSS Dataset Description

As mentioned previously in this work, a dataset created by NASA has been used, generated with the dynamic model of Simulation of Commercial Modular Aeropropulsion Systems (C-MAPSS). The damage propagation modeling used to generate this synthetic dataset builds on the modeling strategy from previous work and incorporates two new levels of fidelity. First, considering the actual flight conditions recorded on board of a commercial aircraft. Secondly, it extends the degradation model by relating the degradation process to operation history [1].

The CMAPSS dynamic model is a high-fidelity computer model for the simulation of a large and realistic commercial turbofan engine. A schematic representation of the engine is shown in Figure 1, where many of the turbine sensors to be processed can be observed. In addition to the thermodynamic model of the engine, the package includes an atmospheric model capable of operating at altitudes from sea level to 40,000 feet, Mach numbers from 0 to 0.90 and sea level temperatures from -60 to 103 degrees F. The CMAPSS software also includes a power management system that allows the engine to run in a wide range of thrust levels across the full range of flight conditions.

Fig. 1. Schematic representation of C-MAPSS Model from NASA.

Table 1. Descriptor Variables of Flight Envelope.

Index	Symbol	Description	Units
1	alt	Altitude	ft
2	Mach	Flight Mach number	-
3	TRA	Throttle-resolver angle	%

The NASA dataset provides synthetic run-to-failure degradation trajectories from optimal turbine condition until failure for a small fleet, comprising nine turbofan engines with unknown and different initial health conditions.

The actual flight conditions as they would be recorded on board a commercial aircraft (also known as the operating mode) were taken as input for the C-MAPSS model. Table 1 presents the scenario descriptor variables that describe the flight conditions, which is called Flight Envelope.

Within the NASA dataset, according to their structural and simulation analysis, contemplates that the datasets consist of multiple multivariate time series. Each dataset is in turn divided into training and testing subsets. Each time series comes from a different engine, that is, the data can be considered to come from a fleet of engines of the same type. Each engine starts with different degrees of initial degradation and manufacturing variations unknown to the user. These degradations and variations are considered normal, that is, it is not considered a fault condition. There are three operating settings/modes that have a substantial effect on engine performance. These adjustments are also included in the data. The data is contaminated with sensor noise.

The engine runs normally at the beginning of each time series and begins to degrade sometime during the series. In the training set, the degradation grows in magnitude until a predefined threshold is reached beyond which it is not preferable to operate the motor. In the test set, the time series ends some time before complete degradation. The objective is to predict the number of remaining operating cycles earlier in the test set, that is, the number of operating cycles after the last cycle that the engine will continue to operate normally.

Table 2. Dataset variables description.

No	Variable
1	Turbofan unit number
2	Time, in cycles
3	Operational setting 1
4	Operational setting 2
5	Operational setting 3
6	Sensor measurement 1
7	Sensor measurement 2
8	Sensor measurement 3
…	…
26	Sensor measurement 21

Fig. 2. Charts for all sensors of all turbofan units from cycle 1, until failure.

The data is provided by NASA as a compressed zip text file with 26 columns of numbers, separated by spaces.

Table 2 represents the structure of the columns in the dataset.

The NASA datasets used were PHM08 and the Turbofan Engine Degradation Simulation Data Set.

The main characteristics of PHM08 dataset are presented below:

−   Consists of 45,918 records

−   Training dataset trajectories: 218 different turbofans where each one fails at a different point.

− Test dataset trajectories: 218 different turbofans where each one fails at a different point.

As explained in the NASA repository instructions, the three operational settings stated previously in this work are also included in the data.

To work with this data, the following Python libraries were used:

− TensorFlow, used for numerical calculation, for deep learning models.

− SKLearn, used for machine learning.

− Pandas, used for numerical data analysis, whose function is to create a data frame structure to be able to work with it.

− Numpy, used for linear algebra.

− Matplotlib, used for graphics.

− mpl_toolkits, used for graphics.

Sensor's selection process is critical to improve aircraft engine diagnostics [10]. During the first analysis of the dataset, it has been observed that the last two columns are null data, so those columns have been removed from the data frame as a first strategy.

After this, an array was created with the names of the columns of the dataset, to be able to add them to the structure of the data frame, since the dataset comes as raw data, without structure. Next, descriptive statistics was applied in an exploratory analysis, and the most relevant sensors were plotted against the RUL variable (the number of cycles remaining for the turbofan to reach a failure state), where it can be observed that it starts at 357, which is the greater number of cycles of a turbofan in the dataset, up to where RUL is 0. Something interesting to observe is the density of the graph at different times in cycles.

It can be observed that the X axis shows the RUL variable from the highest, which is the maximum number of cycles (357), where there is less density in the graph due to the existence of fewer sensors with many cycles, until failure (RUL = 0). After looking at the behavior of all 21 sensors, 6 sensors behave with constant values during all cycles, so the following sensors ["T2", "P2", "epr", "farB", "Nf_dmd", "PCNfR_dmd"] were then extracted from the dataset.

StandardScaler class from SkLearn Machine Learning library has been used to standardize the characteristics by eliminating the mean and scaling to the variance unity.

The standard result of a sample x is calculated as follows:

$$z = (x - u) / s, \tag{1}$$

where ***u*** is the mean of the training samples or zero if ***with_mean = False***, and *s* is the standard deviation of the training samples or one if ***with_std = False***.

Jose Nava, Alberto Ochoa

3   Experiments and Creation of the Neural Network Model

In this work two hypotheses have been proposed that will be later supported with experiments: (H1) An LSTM architecture is more powerful than a traditional multilayer neural network when applied to multivariate time series prediction tasks, and when is included in an automated prediction framework such as the one in the present work, it can outperform multivariate time series predictions on RUL prediction reference data sets. The alternate hypothesis (H2) states that a traditional multilayer neural network outperforms an LSTM architecture in multivariate time series predictions on RUL prediction reference data sets.

For this type of problem, a deep recurrent neural network (RNN) is used to learn the multivariate time series regression function. It is important to state three complex conditions of the problem, which are the multiple operating conditions, the different operating behavior between sensors in terms of the range of parameter values, and the lack of knowledge of the exact starting point of failure or degradation.

In recent years, deep recurrent neural networks (RNN) based on gated units such as Long Short Term Memory [3] have been used successfully to model sequential data. RNNs have been shown to model the temporal (sequential) aspect of sensor data, as well as capture inter-sensor dependencies. The RNNs have been used to model the behavior of motors as a function of time series of multiple sensors with applications for the detection of anomalies and faults [5].

An LSTM unit maintains a cell state using an entry gate, a forget gate, and an exit gate: at a given time step, the entry gate decides what should be added to the cell state, the forget gate decides what should be removed from the state cell, and the output gate decides what part of the cell state should be the output from the LSTM unit [2].

In the equations below, column vectors are denoted with lowercase in bold and matrices with uppercase in bold. For a hidden layer with h LSTM units, the values for the input gate i_t, the forget gate f_t, the output gate o_t, the hidden state z_t, and the state of cell c_t at time t are calculated using the input current x_t, previous hidden state z_{t-1}, and cell state c_{t-1}, where i_t, f_t, o_t, z_t, and c_t are h-dimensional vectors with real values so that:

$$zt = f(xt,\ zt\text{-}1,\ ct\text{-}1), \tag{1}$$

as indicated in the following equation [11]:

$$\begin{pmatrix} i_t^l \\ f_t^l \\ o_t^l \\ g_t^l \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ tanh \end{pmatrix} W_{2h,4h} \begin{pmatrix} D(z_t^{l-1}) \\ z_t^{l-1} \end{pmatrix}, \tag{1}$$

where the time series goes through the following transformations iteratively in the *l-th* hidden layer for $t = 1$ to T, where T is the length of the time series.

LSTM has the powerful ability to remove or add information to the state of the cell, carefully regulated by these structures called gates [7]. Gates are a way to optionally pass information. They are composed of a sigmoid neural network layer and a point multiplication operation. The sigmoid layer generates numbers between zero and one, which describes how much of each component should be allowed to pass. A value of

```python
model = keras.Sequential([
    keras.layers.Dense(24, activation=tf.nn.relu,
                       input_shape=(train_data.shape[1],)),
    keras.layers.Dense(24, activation=tf.nn.relu),
    keras.layers.Dense(24, activation=tf.nn.relu),
    keras.layers.Dense(1)
])
```

Fig. 3. First Neural model with a traditional Sequential Layer Dense layer.

```python
model = tf.keras.Sequential()

normaliza =tf.keras.layers.experimental.preprocessing.Normalization()
normaliza.adapt(train_x)
model.add(normaliza)

model.add(tf.keras.layers.LSTM(32, return_sequences=True))
model.add(tf.keras.layers.LSTM(32))
model.add(tf.keras.layers.Dense(1))

model.add(tf.keras.layers.Lambda(lambda x: x * 206))
lr_schedule = tf.keras.callbacks.LearningRateScheduler(lambda epoch: 1e-8 * 10**(epoch / 20))

optimizer = tf.keras.optimizers.SGD(learning_rate=1e+8, momentum=0.9)
model.compile(loss=tf.keras.losses.Huber(),optimizer=optimizer,metrics=[tf.keras.metrics.RootMeanSquaredError()])
```

Fig. 4. Second Neural model with 2 LSTM layers of 32 nodes and a Dense layer.

zero means "let nothing go", while a value of one means "let everything go." An LSTM has three of these doors to protect and control the state of the cell [8].

To compile a model in keras with TensorFlow, an optimizer (keras.optimizers) has to be chosen, and for this model, Stochastic Gradient Descent (SGD, or gradient descent) has been selected, whose implementation has been widely used in the past with great success in many other projects. This is a first-order iterative optimization algorithm for finding a local minimum of a differentiable function. The goal is to take small steps repetitively in the direction opposite to the gradient of the function at the current point, because this is assumed to be the direction of the fastest and deepest descent.

This is done until a local minimum is found.

For the loss function, Huber Loss has been chosen, which is the combination of the Gaussian loss function and the Laplace loss function, but it has better performance than the Gaussian loss function. The Huber Loss function is less sensitive to outliers in the data than the squared error loss function. It is also differentiable at 0. It is basically an absolute error, which becomes quadratic when the error is small. How small that error must be to be quadratic depends on a hyperparameter, δ (delta), which can be adjusted in code. The Huber function gets closer to MSE when $\delta \sim 0$ and MAE when $\delta \sim \infty$, with large numbers [6].

In this first Neural model, a Sequential traditional Neural Network was implemented, which is a linear stack of layers in the network. This network configuration (as shown in figure 3) has 1,873 parameters with 4 Dense layers. The result of this simple neural network was very far from optimal, with a result of a loss of 66.45 and an RMSE of 73.21. These results will be kept to work with our hypothesis H1 and H2.

The next step was to build the second network, an LTSM Recurrent Neural Network. RMSE is in terms of the same units as the dependent variable, in this case, our RUL [14]. This means that there is no absolute good or bad threshold, however it can be defined based on its dependent variable. In this way, theoretical statements about

Fig. 5. Plot of RUL of the training data Vs the prediction of the model (left) and Plot of the result of the test data (data that the model has never seen) Vs the prediction of the model (right).

RMSE levels can be made if what is expected of the independent variable in the field of research is already known.

As can be noticed in the code, the model consists of 5 layers, of which 2 are LSTM, a normal layer (Dense), and the output layer, where the output by the mean RUL which is 206 (the mean of the descriptive statistics of the train_FD001 dataset cycles) are multiplied through a lambda function.

TensorFlow's powerful callbacks tool has been used, which provides a great degree of control in many aspects of the model, i.e., which has the advantage of the ability to stop training with this functionality when an expected efficiency is reached, and thus, generating an optimal strategy for the neural model is possible.

By calling to the *summary()* function, it can be observed that the total number of parameters is 14,398 divided into the 5 layers of the neural network.

The main strategy in this model is to use the ***LearningRateScheduler()*** function (from tf.keras.callbacks) in the first training of this model to move the learning rate in each epoch, adding a callback to move this learning rate. This will be called in the callback at the end of each epoch. What this does is change the learning rate to a value based on the epoch number. So, in epoch 1, it's 1 times 10 to the -8, times 10 to the 1 in 20. And when the epoch 100 is reached, it will be 1 times 10 to the -8, times 10 to the 5, and so on, and that is 100 over 20. This will happen in every callback because it is set in the callback parameter of the function. An advantage of this strategy is that at the end of the execution, a chart of loss Vs epochs and another chart of loss Vs learning rate can be created. These charts will provide extremely important information about what happens in each epoch and in each step of the training. The ***fit()*** function (***model.fit***) is called, where the strategy is training by 100 epochs in order to later visually analyze the behavior of the learning rate through the epochs. Callbacks is also implemented, to call ***lr_schedule*** as explained above.

Even though the result of the first training ended with a final loss or loss of 42.4795 and an RMSE of 53.2980, it can be observed an incremental trend in the previous graph, although there are also anomalous data that that can be attributed to the type of problem, where there is an infinite number of factors that can produce a great number of anomalous data in the sensors of a turbofan.

Fig. 6. Charts for Loss Vs Epochs (left) and Loss Vs Learning Rate (right).

Fig. 7. Plotting two red straight lines on our Loss Vs Learning Rate graph (left) and Plot of RUL of our training data Vs the prediction of our model after updating the learning rate and using *myCallback* and *on_epoch_end* (right).

To know how much the model fits, the Coefficient of Determination (also called R2, or R squared) is now calculated, which is the proportion of variance (%) in the dependent variable that can be explained by the independent variable [13].

This shows how close the data is to the fitted regression line. The model output with the test data (which the model has never seen) was -0.602 of R2, which states that the model needs to be improved.

This result when calculating RMSE was 52.60, this means that on average there could be a deviation of the RUL result of 52.60 units up or down. Now observing what the model did in each epoch is needed, and how the loss changed during the training.

For doing this, Loss Vs epochs is plotted, and with this a notion of what was happening with the loss will be revealed, if it was going down and if it behaved erratically at some point, in order to adjust our variables at the right moment. The history variable is used, which is created by assigning the history of the call to the *model.fit*() function, thanks to history, access to all the information about what was happening at each moment in our training is possible.

Having this information on how loss has behaved during the epochs, it can be observed that it was going well during the first 50 epochs, when the loss variable began to rise, which was not efficient at all. Next step is looking at the chart of the right, Loss Vs Learning Rate, to get more information about the behavior of the learning rate and see if making a decision with this information is possible.

By visually observing in the graph at which moment the loss variable begins to have an erratic increasing behavior, the learning rate of the exact moment where this changes took place can be observed, to use it as the initial learning rate in Gradient Descent and run the whole model again. To do this, two **red** straight lines are drawn from the moment the learning rate began to behave erratically.

As can be seen in the previous chart, the learning rate $7 * 10^{-6}$ has been chosen to use it in the second training and see how the model behaves, doing the same tests and thus the same graphs after the training can be observed to be able to make other decisions.

To prepare the second training, the learning rate scheduler is no longer needed, since it has fulfilled its objective in the first training. In this training SGD will do its job, but now with an optimal initial learning rate.

A class **myCallback** is now created, in which the function **on_epoch_end()** will be defined, which will be in charge of stopping the training when RMSE reaches a target value, in this case 0.5. This is done to avoid that the training passes through optimal values and behaves erratically, returning to inefficient large values.

The call to the fit function (**model.fit**) can be executed next, but now, training for 500 epochs, as intensive training, and also calling **callbacks** function is needed (to **myCallback** class), to stop the training when reaching an RMSE less than 0.5.

When the model was executed again with **model.fit** using the new strategy, the result of the training with 500 epochs was a loss of 4.3183, which is quite good. But the expectation was to see if a better value could be given, so another execution of 500 epochs was performed, this time with **callbacks** to stop when RMSE is less than 3.5 and loss is approximately 2.

This strategy worked pretty well, **callbacks** function helped so that the behavior of the model did not raise the loss again, obtaining a final loss of 2.1686, which is much better than the previous one.

The right image in figure 7 shows a perfect line in graphing the training RUL and the result of the model predictions.

Next step was to observe how the model behaved with data that it had never seen, after knowing that the behavior of the model with the training data has been optimal.

The behavior with the test data that the model had never seen was observed, and with this, an R2 of 0.701 and an RMSE of 20.419 were reached. This result is quite good for this type of neural models with time series.

As it was shown in previous runs, it was possible to observe what the model did in each epoch and how the loss changed during training. And for this a graph of Loss Vs epochs was created, and with this, a good insight of what was happening with the loss was provided, if it was going down and if it behaved erratically at some point, to adjust our variables at the right time.

Analyzing Loss Vs Epochs chart can help to observe behaviors through the epochs in training, the algorithm arrived at a Loss that tends close to zero, which is ideal for the model, and analyzing Loss Vs Learning Rate graph provides information to conclude that a model with which a model can reliably calculate RUL in turbofan units in the aerospace industry was achieved.

Fig. 8. Charts for Loss Vs Epochs (left) and Loss Vs Learning Rate (right).

4 Results of Experiments

H1 has been accepted when creating the models in figure 3 and figure 4 and iterating through different strategies in each experiment. The traditional model in figure 3 threw results very far from an optimal model and using LTSM neural network layers reached an optimal model. H1 was accepted because when a traditional and sequential Neural Network was used the loss result was 66.45 and RMSE was 73.21, and when using LTSM the loss was 2.1686 and an RMSE of 3.5. H2 has been rejected when analyzing these results.

5 Conclusions and Future Research

With the model and adjustments made in this work, optimal results for a model of this type were achieved, where the main characteristic of the model is that the estimation of the remaining useful life from data such as those of multiple sensors in a turbine, and where this engine is it degrades through time, it can be considered as the learning of a regression function, which maps a multivariate time series to a real value number. Challenges in the supervised learning-based approach have been highlighted, such as missing data in datasets, learning to estimate RUL values, sensor noise, anomalies due to lack of knowledge of the beginning of engine degradation, and implementation of neural models with very large multivariate time series, dealing with multiple operating conditions, etc. A solution strategy has been executed in the context of these challenges.

It is important to highlight that having such powerful tools as TensorFlow, Keras, Numpy and Pandas provide the programmer with the ability to create strategies with neural networks as powerful as observing the learning rate and loss during the epochs in a visual way, generating great positive impacts on the artificial intelligence industry, such as accelerating strategies to reach exponentially better results faster.

In the United States, in collaboration with NASA through multiple aviation contracts, the company Ampaire, Inc. is validating the capabilities of electric aircrafts. Such systems are starting to use RTOS operating system have the advantage of being really "in real time" (which, so far, do not exist similar systems in aviation). Advanced

aviation systems such as Ampaire's, could be implementing their sensor strategy with a microcontroller working with RTOS, such as the product of the company Hover, Inc., which in partnership with the company Pinnacle Aerospace, Inc., have created the first Hardware under RTOS for aviation systems, certified by the FAA. This type of systems could be adapted to work with turbofan units and could take advantage of high availability architectures for decision making in monitoring systems, where, besides the great potential provided by Machine Learning solutions, they can implement a reliable approach in the sensor reading strategy that combines algorithms for reading sensors, pre-processing data with inferential statistics, and producing a high decision-making capacity, with less computing resources.

References

1. Chao, M. C., Kulkarni, Goebel, K., Fink, O.: Aircraft Engine Run-to-Failure Dataset under real flight conditions, NASA Ames Prognostics Data Repository (http://ti.arc.nasa.gov/project/prognostic-data-repository), NASA Ames Research Center, CA, USA (2020)
2. Eide, I., Westad, F.:Automated multivariate analysis of multisensor data submitted online: Real-time environmental monitoring. PLoS ONE 13(1): e0189443 (2018)
3. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation, 9(8), 1735–1780 (1997)
4. Jansohn, P.: Overview of gas turbine types and applications. In: Modern Gas Turbine Systems. High Efficiency, Low Emission, Fuel Flexible Power Generation. Paul Scherrer Institute, Switzerland (2013)
5. Malhotra, P., Vig, L., Shroff, G., Agarwal, P.: Long Short Term Memory Networks for Anomaly Detection in Time Series. In: ESANN, 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, pp. 89–94 (2015)
6. Mayer, G. P.: An Alternative Probabilistic Interpretation of the Huber Loss, Uber Advanced Technologies Group. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. USA (2019)
7. Nguyen, K.T.P., Medjaher, K.: A new dynamic predictive maintenance framework using deep learning for failure prognostics. Reliability Engineering and System Safety, 188. 251−262 (2019)
8. Olah, Ch.: Understanding LSTM Networks, Recurrent Neural Networks, CA, USA (2015, Aug 27) http://colah.github.io/posts/2015-08-Understanding-LSTMs/
9. Petersen, R., Khan, O.: Aviation Maintenance Technician Certification Series, Propulsion, Module 14 for B Certification. In: EASA, European Aviation Safety Agency. Aircraft Technical Book Company, CO, USA (2019)
10. Shane, T.: Systematic Sensor Selection Strategy for Turbofan Engine Diagnostics. In: GRC Propulsion Control and Diagnostics Workshop, NASA, QinetiQ North America, USA (2009)
11. TV, V., Gupta, P., Malhotra, P., Vig, L., Shroff, G.: Recurrent Neural Networks for Online Remaining Useful Life Estimation in Ion Mill Etching System. In: Annual Conference of the PHM Society, 10(1) (2018)
12. Mendez, R., Ochoa, A., Zayas, B., Perez, M., Quintero, O.: Implementation of Big Data in Intelligent Analysis of Data from a Cluster of ROVs Associated with System of Prevention and Reparation of Hydrocarbon Leaks to Optimize their Distribution in Gulf of Mexico. Research in Computing Science, 149(7), pp. 73−85 (2020)

13. Reyes, V., Cossio, E., Pescador, A.: Demand Forecasting Applied to Radio Frequency Identification Technology. Research in Computing Science, 149(7), pp. 97−109 (2018)
14. Hernandez, R., Gabbasov, R., Suárez, J.: Improving Performance of Particle Tracking Velocimetry Analysis with Artificial Neural Networks and Graphics Processing Units. Research in Computing Science, 104, pp. 71−79 (2015)