

Article

A Machine Learning and Feature Engineering Approach for the Prediction of the Uncontrolled Re-Entry of Space Objects

Francesco Salmaso, Mirko Trisolini *  and Camilla Colombo 

Department of Space Science and Technology, Politecnico di Milano, Via La Masa 34, 10156 Milan, Italy; camilla.colombo@polimi.it (C.C.)

* Correspondence: mirko.trisolini@polimi.it

Abstract: The continuously growing number of objects orbiting around the Earth is expected to be accompanied by an increasing frequency of objects re-entering the Earth's atmosphere. Many of these re-entries will be uncontrolled, making their prediction challenging and subject to several uncertainties. Traditionally, re-entry predictions are based on the propagation of the object's dynamics using state-of-the-art modelling techniques for the forces acting on the object. However, modelling errors, particularly related to the prediction of atmospheric drag, may result in poor prediction accuracies. In this context, we explored the possibility of performing a paradigm shift, from a physics-based approach to a data-driven approach. To this aim, we present the development of a deep learning model for the re-entry prediction of uncontrolled objects in Low Earth Orbit (LEO). The model is based on a modified version of the Sequence-to-Sequence architecture and is trained on the average altitude profile as derived from a set of Two-Line Element (TLE) data of over 400 bodies. The novelty of the work consists in introducing in the deep learning model, alongside the average altitude, and three new input features: a drag-like coefficient (B^*), the average solar index, and the area-to-mass ratio of the object. The developed model was tested on a set of objects studied in the Inter-Agency Space Debris Coordination Committee (IADC) campaigns. The results show that the best performances are obtained on bodies characterised by the same drag-like coefficient and eccentricity distribution as the training set.



Citation: Salmaso, F.; Trisolini, M.; Colombo, C. A Machine Learning and Feature Engineering Approach for the Prediction of the Uncontrolled Re-Entry of Space Objects. *Aerospace* **2023**, *10*, 297. <https://doi.org/10.3390/aerospace10030297>

Academic Editor: Vladimir S. Aslanov

Received: 28 February 2023

Revised: 13 March 2023

Accepted: 14 March 2023

Published: 17 March 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: re-entry predictions; machine learning; deep learning; features engineering; uncontrolled re-entry

1. Introduction

From the launch of the first artificial satellite on the 4 October 1957, the population of bodies orbiting around the Earth has been continuously growing with an exponential trend [1]. Furthermore, in the last years, this growth has been exacerbated by the creation of large constellations for telecommunication, navigation and global internet coverage. Post-mission disposal guidelines [2,3] have been introduced in order to preserve the Low Earth Orbit (LEO) region and incentivise the early disposal of satellites so that they would not affect the already polluted orbital environment. Focusing on the re-entry problem, as of March 2022, a total of 26,135 bodies have re-entered the atmosphere (www.space-track.org, accessed on 13 March 2022). According to Pardini et al. [4], since April 2013, only one out of five re-entries of intact objects was controlled, corresponding to a total mass of roughly 11,000 metric tonnes, mainly related to rocket bodies and spacecraft. Therefore, the continuously growing number of orbiting objects should be correlated with an expected increase in the re-entries of both controlled and uncontrolled objects. However, predicting the re-entry epoch and the location of an uncontrolled body remains critical because it is affected by several uncertainties, which could translate into risks of casualties on the ground [4].

In the typical dynamics-based approach to the re-entry prediction problem, the trajectory of the spacecraft is propagated until it reaches the altitude when the break-up occurs. This approach consists of determining the initial position of the object and accurately modelling the forces acting on it in order to predict the evolution of its trajectory. However, in general, it is difficult to accurately model the forces acting on the spacecraft and to precisely know its physical characteristics and state. In fact, a typically recommended relative uncertainty value of $\pm 20\%$ is recommended for re-entry predictions [5,6]. Ref. [4] provides an in-depth review of the different sources of uncertainties in re-entry predictions, the most relevant being the modelling of the drag acceleration. Other forces affect the re-entry dynamics such as solar radiation pressure and lunisolar attraction. However, in this work, we focused on the latest stage of the re-entry prediction of low eccentricity orbits; therefore, the dominant force is atmospheric drag and the main uncertainties are related to the atmospheric density. It is modelled through empirical models which describe the density variations on spatial and temporal scales [7–11]. However, they are affected by two types of uncertainties [10]: the simplified physical modelling on which they are based, and the uncertainties and complexities in predicting the space weather, in particular the solar index. Another important variable for the predictions of low eccentricity orbits is the ballistic coefficient of the object of interest. The ballistic coefficient depends on the object's shape, orientation, and attitude, which can be difficult to accurately know for uncontrolled objects.

In this circumstance, it becomes interesting to adopt a paradigm shift to model the problem, from a physical to a data-driven point of view. In this context, machine learning represents a new approach, where, instead of accurately modelling the space environment, its forces and how the spacecraft interacts with them, it autonomously builds its knowledge based on data, autonomously dealing with the complexities and uncertainties of the phenomenon. Jung et al. [12] proposed a machine learning approach to predict the re-entry epoch of an uncontrolled satellite based on Recurrent Neural Networks (RNNs). In their work, the RNN was trained on a set of Two-Line Elements (TLEs) sets using as the only feature the evolution of the average altitude of the object. TLE sets provide, among other information, the time evolution of the position of catalogued objects in orbit. However, as also addressed by Lidtke et al. [13], TLEs suffer from some inaccuracies that must be taken into account. Most notably, TLEs of an object are not homogeneous—they can be provided at irregular time steps and contain incorrect observations or observations belonging to other objects (outliers). In addition, TLEs do not directly provide information on the physical characteristics of the object (e.g., the ballistic coefficient), rather, they combine all the uncertainties and modelling errors in a single drag-like coefficient, referred to as B^* [8,9]. Finally, TLEs do not contain accuracy information on the measure they provide; that is, they do not provide covariance information. However, TLE sets are currently the only public source of data that can be used for re-entry predictions. Therefore, despite their limitations and inaccuracies, they were used in this work to assess the capabilities of a machine-learning approach using multiple input features. As shown in [14], initial covariance information can be estimated using pseudo-observations, and the accuracy of TLE data can be improved as shown [15,16]. While in this work we focus on the application of a machine learning technique to raw TLE data, future work can benefit from the inclusion of such information to increase the accuracy of the re-entry predictions.

Starting from the work of Jung et al. [12], we expand on the machine learning model by introducing additional features, namely the drag-like coefficient, B^* , and the solar index. The first feature will allow the machine learning model to account for the physical characteristics of the objects, while the second feature will give a temporal dimension to the model, accounting for the different solar activity levels and how they affect the lifetime of an object. The developed model was trained on hundreds of TLE sets belonging to uncontrolled objects such as spent payloads and rocket bodies, and debris. Consequently, the same model can be tested on different objects, in order to obtain the output sequence and assess its accuracy.

The paper is structured as follows. Section 2 describes the deep learning model architecture adopted; Section 3 contains the procedure used for the pre-processing and filtering of the data required by the deep learning model; Section 4 contains an analysis of the characteristics of the input features of the model; and Section 5 shows the results obtained by the model on selected test cases, together with the optimisation procedure of the model's hyperparameters.

2. Deep Learning Model

From a machine learning perspective, the task of predicting the re-entry of an object is a time-series problem, where the aim of the model consists of predicting the evolution of the trajectory based on an initial set of conditions so that the re-entry epoch can be determined. In other words, it is a regression problem. Time series are discrete-time data and can be considered as a particular case of sequential data, where each value is associated with a time stamp. Therefore, the order of the series is imposed by the time dimension and it must be preserved. Let us denote an input sequence of length T_x as [17]:

$$\mathbf{x} = \{\mathbf{x}^{<1>}, \mathbf{x}^{<2>}, \dots, \mathbf{x}^{<T_x>}\}, \quad (1)$$

where the vector $\mathbf{x}^{<t>} \in \mathbb{R}^p$ contains the input features at time step t . Furthermore, let us denote a target sequence of length T_y as [17]:

$$\mathbf{y} = \{y^{<1>}, y^{<2>}, \dots, y^{<T_y>}\}, \quad (2)$$

where $y^{<t>} \in \mathbb{R}$ is the target feature. We define the predicted sequence as [17]:

$$\hat{\mathbf{y}} = \{\hat{y}^{<1>}, \hat{y}^{<2>}, \dots, \hat{y}^{<T_y>}\}. \quad (3)$$

The aim of the deep learning problem consists of predicting the target sequence, minimising a loss function, which in this work was selected as the Mean Squared Error (MSE) [17]:

$$\mathcal{L}_{MSE}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{T_y} \sum_{t=1}^{T_y} (y^{<t>} - \hat{y}^{<t>})^2. \quad (4)$$

The choice of the MSE is natural in this context due to its similarity with the Euclidean distance. Indeed, the MSE can be intuitively considered as the squared distance between the true position and the predicted one, averaged by the number of observations in the output sequence.

Sequence-to-Sequence Architecture

In this work, we selected a Sequence-to-Sequence (Seq2Seq) [18] architecture to model the re-entry problem, due to its ability to deal with input and output sequences of different lengths. As shown in Figure 1, a Seq2Seq architecture is composed of two neural networks, the encoder and the decoder, and the context vector. The encoder receives as input a sequence of arbitrary length, T_x , and generates a fixed-length vector, called the context vector. This vector summarises the information contained in the input sequence. The outputs of the decoder are simply discarded. At each time step, the cell produces the hidden state considering the input at the same time step and the hidden state at the previous time step [19]:

$$\mathbf{h}_{Enc}^{<t>} = f(\mathbf{x}^{<t>}, \mathbf{h}_{Enc}^{<t-1>}, \boldsymbol{\theta}), \quad (5)$$

where $\mathbf{h}_{Enc}^{<t>}$ is the hidden state of the encoder at time t , $\mathbf{h}_{Enc}^{<t-1>}$ is the hidden state at the previous time step, $\mathbf{x}^{<t>}$ is the input vector at the current time, and $\boldsymbol{\theta}$ is the weights vector. The function f represents a neural network, which can be a simple Recurrent Neural Network (RNN) [20,21] or a more complex Gate Recurrent Unit (GRU) [19,20,22] or a

Long-Short Term Memory (LSTM) [19,20] network. At the final time step, T_x , the hidden state of the encoder is defined as the context vector as follows:

$$c_{context} = \mathbf{h}_{Enc}^{<T_x>}. \quad (6)$$

The decoder is initialised with the final state of the encoder, $\mathbf{h}_{Dec}^{<t'=0>} = c_{context}$. The aim of the decoder is to generate the output sequence based on this encoded vector. At each time step $t' \in [1, T_y]$, given the hidden state and the output at the previous time step $t' - 1$, the decoder hidden state can be defined as follows [19]:

$$\mathbf{h}_{Dec}^{<t'>} = g(\hat{y}^{<t'-1>}, \mathbf{h}_{Dec}^{<t'-1>}, \theta'), \quad (7)$$

where again $g(\cdot)$ can be a simple RNN, a GRU or an LSTM. Subsequently, it is possible to obtain the output $\hat{y}^{<t'>}$ by passing the corresponding output of the decoder to an output layer. Note that the hidden state at time $t' = T_y$ is simply discarded.

During training, the decoder receives its own previous hidden state and the true input, or the predicted output at the previous time step, to generate a new state and output sequentially. When only the “true input” is provided to the decoder, the training method is referred to as Teacher Forcing [20]. Therefore, in this method, the predicted output of the model is substituted with our reference data, which is considered the ground truth. However, this method introduces a discrepancy between the training phase and the inference phase in which the network is used in open-loop, which can introduce errors in the prediction as the network is not able to learn from its own mistakes. To mitigate this problem, Bengio et al. [23] proposed an approach called Curriculum Learning. It consists of randomly selecting either the previous true output or the model estimate, with a probability that is proportional to the training epoch. Intuitively, at the initial epochs of training, the model should be trained with teacher forcing since it is not still well trained and errors can be made. Instead, at the final epochs, the model is expected to be well trained so that it is possible to use the model output, as during inference. Specifically, for every output to predict $y^{<t'>}$ of every batch at the j -th epoch, a coin is flipped to decide to feed either the previous true output, associated with a probability ϵ_j , or the model estimate, with a probability $(1 - \epsilon_j)$. Given the fact that, in this context, the size of the dataset is quite small and the training epoch number is large, the probability is set as a function of the training epoch, instead of the iteration [23], as:

$$\epsilon_j = k^j, \quad (8)$$

where $k < 1$ is a constant that depends on the speed of convergence, which is the decay rate of the scheduled sampling.

The representation of the Seq2Seq model during training is provided in Figure 1, where $x^{<t'>}$ denotes each time component of the input sequence; $\hat{y}^{<t'>}$ the predicted output; and $y^{<t'>}$ the ground truth.

Gate Recurrent Unit

In this work, both the encoder and the decoder were based on Gate Recurrent Units (GRUs) [22], given their capability of dealing with the vanishing gradient problem that characterises simple RNN when trained with long sequences. It can be seen as a simplified version of the LSTM and it has the advantage of reducing the computational cost, providing comparable performances [24]. At each time step, the GRU computes:

$$z^{<t>} = \sigma(W_z x^{<t>} + V_z c^{<t-1>} + b_z) \tag{9}$$

$$r^{<t>} = \sigma(W_r x^{<t>} + V_r c^{<t-1>} + b_r) \tag{10}$$

$$\tilde{c}^{<t>} = \tanh(W_c x_t + V_c (r^{<t>} \odot c^{<t-1>} + b_c)) \tag{11}$$

$$c^{<t>} = (1 - z^{<t>}) \odot \tilde{c}^{<t>} + z^{<t>} \odot c^{<t-1>}, \tag{12}$$

where $W_z, W_r, W_c \in \mathbb{R}^{k \times p}$ denote the input weight matrices; $V_z, V_r, V_c \in \mathbb{R}^{k \times k}$ are the hidden state weights; $b_z, b_r, b_c \in \mathbb{R}^k$ are the bias vectors; \odot the element-wise product; $\sigma(\cdot)$ and $\tanh(\cdot)$ are the sigmoid function and the hyperbolic tangent function, respectively. The variables $z^{<t>} \in \mathbb{R}^k$ and $r^{<t>} \in \mathbb{R}^k$ represent the update and reset gate respectively, which helps to tackle the vanishing gradient problem. Their aim consists in controlling the quantity of information of the past time step that needs to be carried forward. In particular, the update gate is responsible for updating the old state by controlling the similarity between the new state and the past one. The reset gate is responsible for setting the amount of information that needs to be forgotten. For these reasons, the values of the gates are designed to be bounded between 0 and 1, by using the sigmoid function. The output hidden state is computed in two steps. First, the candidate memory cell $\tilde{c} \in \mathbb{R}^k$ is calculated according to Equation (11). It takes as input the previous hidden state, which is multiplied by the reset gate, and the current input. Depending on the entries of the reset gate, two limit cases can be identified. When its values are close to unity, the information of the previous time step is carried forward in time, as with a simple RNN. Instead, when the values are close to 0, the information from the previous time step is ignored. After having derived the candidate memory cell, the variable $c \in \mathbb{R}^k$, called the memory cell, can be computed according to Equation (12). The memory cell receives as input the current candidate and the previous memory cell and it is governed by the reset gate only. In this case, when the entries of the reset gate are close to 0, the new memory cell is updated with the candidate, meaning that the information of the previous time step is dropped. Instead, when the values of the gate approach to unity, the candidate is simply discarded and the memory cell maintains its state. Finally, the hidden state and the output of the GRU can be defined as, respectively:

$$h^{<t>} = c_{context}^{<t>} \tag{13}$$

$$\hat{o}^{<t>} = c_{context}^{<t>}. \tag{14}$$

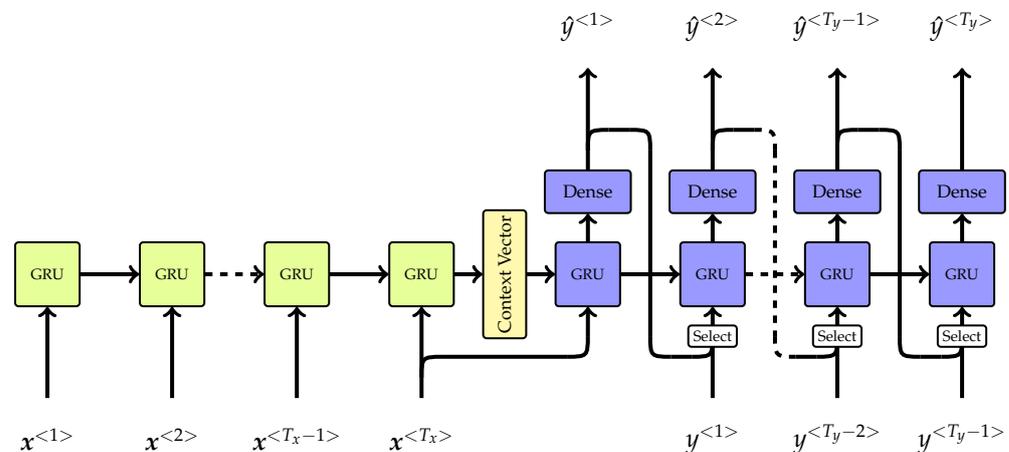


Figure 1. Representation of Seq2Seq architecture during training.

threshold ((optimal tuning parameters for the filter can be found in Tables 4 and 5 of [13]);

5. Filter out TLEs that present negative values of B^* because they can be associated with unmodelled dynamics or manoeuvres.

For a complete description of the filters and their performance, as well as the optimal tuning of the filter parameters, the reader is referred to the work of Lidtke et al. [13]. It is worth noting that some outliers may not be identified by such filters; however, the percentage of outliers is considered to be sufficiently small to ensure the validity of the data. After the pruning process, the mean elements contained in the TLEs were converted to the corresponding osculating variables using the SGP4 propagator [9]. It is also important to consider other characteristics of the dataset. Specifically, we wanted to focus on re-entries that had a sufficient number of TLE points, and whose re-entry prediction was sufficiently accurate. The minimum number of TLE points was selected via a trade-off between the number of pruned TLEs and the accuracy of the fitting procedure (see Section 4.1) [27]. In addition, the orbit should not be too eccentric as the re-entry mechanism for these types of orbits changes significantly. The following criteria were used to filter the dataset:

- Maximum re-entry uncertainty: 20 min;
- Maximum average altitude of the initial TLE: 200 km;
- Minimum average altitude of the final TLE: 180 km;
- Maximum eccentricity: 0.1;
- Minimum number of points: 4 TLE.

The described criteria adopted the average altitude, h , which was computed according to:

$$h = a - R_{\oplus}, \quad (15)$$

where a is the osculating semi major axis; and R_{\oplus} the Earth radius. The result of this pruning and filtering procedure is a dataset with a total of 417 bodies, which is then split in 80% training and 20% validation (see Appendix A. Note that the objects in the validation set must resemble the characteristics of the training set. For example, the average, μ , and the relative standard deviation, σ , of the B^* coefficient and the eccentricity for the training and validation sets should be comparable (Table 1).

Table 1. B^* and eccentricity characteristics for the training and validation sets.

Variable		Training Set	Validation Set
Eccentricity [-]	μ	7.2842×10^{-3}	5.8856×10^{-3}
	σ	9.3963×10^{-3}	7.1684×10^{-3}
B^* [1/ER]	μ	6.3362×10^{-4}	5.9503×10^{-4}
	σ	5.4924×10^{-4}	3.8341×10^{-4}

Moreover, the re-entry epochs are distributed quite uniformly from roughly January 2005 to April 2021 for both the sets, therefore the different objects experience different solar flux conditions and, therefore, similar values of drag acceleration (Figure 3). Finally, it should be highlighted that the majority of the bodies in the training and validation sets are characterised by uncertainty in the re-entry epoch of less than 2 min.

The objects for testing the prediction capabilities of the model were selected among the bodies analysed during IADC campaigns, as listed in Table 2. Note that the re-entry epochs of 23,853 and 19,045 were manually changed to, respectively, 16:25 UTC and 15:08 UTC [6].

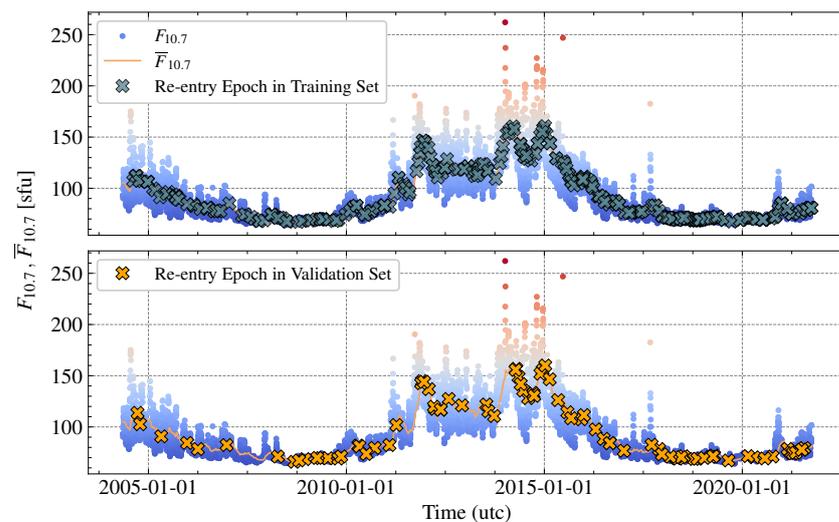


Figure 3. Re-entry epochs in training (Top) and validation (Bottom) sets and solar flux indexes.

Table 2. List of test objects derived from TIP.

NORAD	Name	Re-Entry Epoch [UTC]
23853	COSMOS 2332	2005-01-28T18:05 *
26873	CORONAS F	2005-12-06T17:24
10973	COSMOS 1025	2007-03-10T12:56
31599	DELTA 2 R/B	2007-08-16T09:23
31928	EAS	2008-11-03T04:51
20813	MOLNIYA 3-39	2009-07-08T22:42
11601	SL-3 R/B	2010-04-30T16:44
21701	UARS	2011-09-24T04:00
20638	ROSAT	2011-10-23T01:50
37872	PHOBOS-GRUNT	2012-01-15T17:46
19045	COSMOS 1939	2014-10-29T15:32 *
40138	CZ-2D R/B	2015-06-13T23:58
39000	CZ-2C R/B	2016-06-27T19:04
38086	AVUM	2016-11-02T04:43
37820	TIANGONG 1	2018-04-02T00:16

Note: the epochs denoted with * were manually changed according to the indications of the IADC campaigns.

4. Features Engineering

In this section, we outline the processes to select and pre-process the most important input features required to predict the re-entry epoch by using a machine learning model. However, this is not a straightforward task because it is necessary to find input variables that are correlated with the output of the deep learning model. Therefore, it is essential to analyse each feature from a physical point of view, so that this knowledge can be then applied to constructing the deep learning model, which is then going to benefit from each input variable. This process is typically known as feature engineering. In this optic, the deep learning model is considered a physical model, where, given the input variables, the output is obtained by means of physical relations. Instead, in machine learning, the model is treated as a black box, which is autonomously capable of learning the input-output relations, based on the true values and the input features. Therefore, each feature can be selected and generated based on the knowledge of the physical principles

that govern the re-entry phenomenon, leaving the deep learning model to cope with the learning and the uncertainties of the physical problem. By analysing each feature, the relations between the considered variables can be highlighted so that we can make a more informed decision on whether the considered feature can be beneficial to the deep learning model [27,28].

In addition, the raw data, provided by the TLEs, require a pre-processing and regularisation step in order to be used with an RNN. In fact, TLE data are characterised by a non-uniform generation frequency, which differs from object to object; they also lack high accuracy and may present outliers. All these aspects must be taken into account before feeding the data to the RNN. The following sections analyse the contribution of the features considered in this work, which are the average spacecraft altitude, the ballistic coefficient, the solar index, and the B^* coefficient.

4.1. Average Altitude

The average altitude feature is of fundamental importance for deep-learning-driven re-entry prediction. In fact, it is related to the definition of the loss function used to train the RNN Equation (4). This means that it is used both as input and as output for the deep learning model, such that given a part of the average altitude profile as input, the remaining part is predicted. The pre-processing of the average altitude profile follows the work of Jung et al. [12]. In particular, each trajectory is generated by fitting the available TLEs, below 240 km, according to the following equation:

$$f(t) = a_1 + a_2(t_{ref} - t)^{\frac{1}{2}} + a_3(t_{ref} - t)^{\frac{1}{3}} + a_4(t_{ref} - t)^{\frac{1}{4}}, \quad (16)$$

where t_{ref} is the re-entry time, and a_1 , a_2 , a_3 and a_4 , are the parameter obtained through curve fitting using the Levenberg–Marquardt algorithm, which is available through the Python library, SciPy [29]. The first coefficient, a_1 , is equal to 80 km, which is the reference altitude at which TIP messages are provided [6]. This assumption avoids the fitting algorithm changing the re-entry epoch from the TIP message and, therefore, ensures consistency with the TLE data. Subsequently, the epoch corresponding to an altitude of 200 km is set as the initial epoch so that the time span from the starting altitude of 200 km to the re-entry at 80 km, identifies the residual time.

As each object will have different residual times, it is convenient to interpolate the residual time as a function of the average altitude. To do so, we subdivided the considered altitude range (between 200 km and 80 km) into a uniform grid of 25 points. Consequently, each trajectory was identified by a series of observations $\{x_i, t_i\}$, with $i = 1, 2, \dots, 25$, where x_i represents the average altitude and t_i the relative epoch. In this way, it is possible to use only t_i as the output of the deep learning model. Figure 4 shows a result of such a procedure for the object identified by NORAD 27392. The resulting trajectories are represented for the training, validation, and test sets are represented in Figure 5. It can be seen that there is a concentration of trajectories characterised by residual lifetimes between 0.5 days and 3 days. Meanwhile, only three trajectories with residual lifetimes higher than 6 days are present, and only two of them were used for training. This is important because this distribution incorporates the ballistic coefficient information and it can be directly related to the prediction capability of the model. We also note that, in the test set, the leftmost trajectory has a residual time of roughly 14 min and it is associated with the Molniya 3-39 spacecraft. In this case, the trajectory is highly elliptical and does not satisfy our selection criterion defined in Section 3; however, it was included in the test set to assess the capabilities of the model with such trajectories.

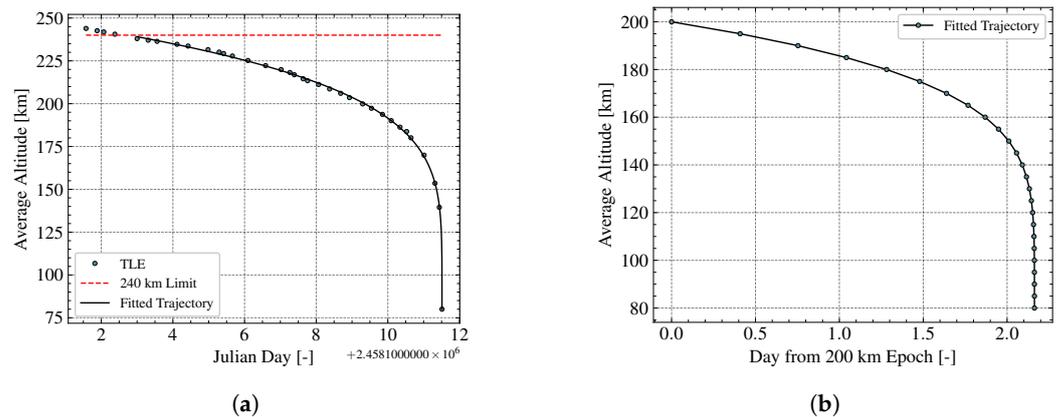


Figure 4. Example of fitted (a) and sampled (b) re-entry trajectory for the object with NORAD number 27392.

In a real case scenario where the entire trajectory is not known a priori, Equation (16) can be used with t_{ref} as the epoch of the last available TLE. However, the resulting curve will be different from the one obtained with all the available TLEs. We thus expect to observe some differences in performance between the training and testing sets. Further analysis should be carried out on this aspect because there may be variations in the predicted trajectory that depend on the adopted procedure and on the number of available TLEs. However, this difference should be minimised when a high number of TLE points is available, which is the case, for example, with the objects analysed during IADC campaigns.

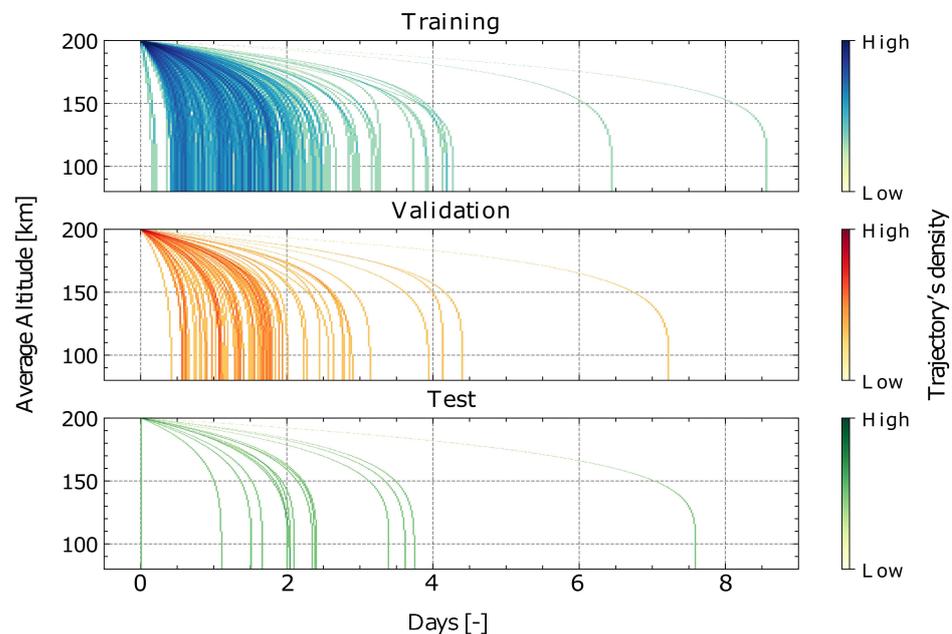


Figure 5. Trajectories generated with the fitting and sampling procedure for the training (top), validation (middle), and test (bottom) sets. Darker shadings are associated to higher concentrations of trajectories for a given re-entry time.

4.2. Ballistic Coefficient and B^*

The ballistic coefficient, B , has a key importance for the re-entry of low eccentricity orbits because the dynamic is dominated by the atmospheric drag perturbation. The ballistic coefficient, as defined in Equation (17) [30], describes the aerodynamic interaction of the spacecraft with the atmosphere via its shape, mass, and drag coefficient.

$$B = C_D \frac{A}{m}. \quad (17)$$

However, as we used TLE data for the re-entry prediction, we used the drag-like coefficient, B^* , as a proxy of the ballistic coefficient. In fact, B^* is the coefficient introduced in the SGP4 propagator, which in turn was used for the generation of TLE data. The expression for the B^* is as follows [9]:

$$B^* = \frac{1}{2} \frac{C_D A}{m} \rho_0 R_{\oplus}, \quad (18)$$

where ρ_0 is the atmospheric density at the perigee of the orbit, assumed as $2.461 \times 10^{-5} \text{ kg/m}^2/\text{ER}$, with ER = 6375.135 km; and R_{\oplus} the Earth radius of 6378.135 km. In general, the B^* coefficient of the SGP4 formulation soaks up different modelling errors. Despite the fact that the ballistic coefficient can be retrieved from B^* , they are not identical. However, in this context, we assumed that B^* can be used as a proxy of B as we were considering objects with similar orbital parameters that will be subject to similar modelling uncertainties. Similarly to the average altitude, it is essential to pre-process the B^* data in order to feed them to the RNN. The first step of the processing consists of defining a sliding window with a fixed size to create a series of average values, in which each term is given by the average of all the previous points. The second step consists in applying a linear piecewise interpolation scheme, in which, between two consecutive values, the preceding one is kept constant. Subsequently, the interpolated curve was sampled at the epochs defined through the average altitude processing. An example of the aforementioned procedure is represented in Figure 6 for NORAD 27392.

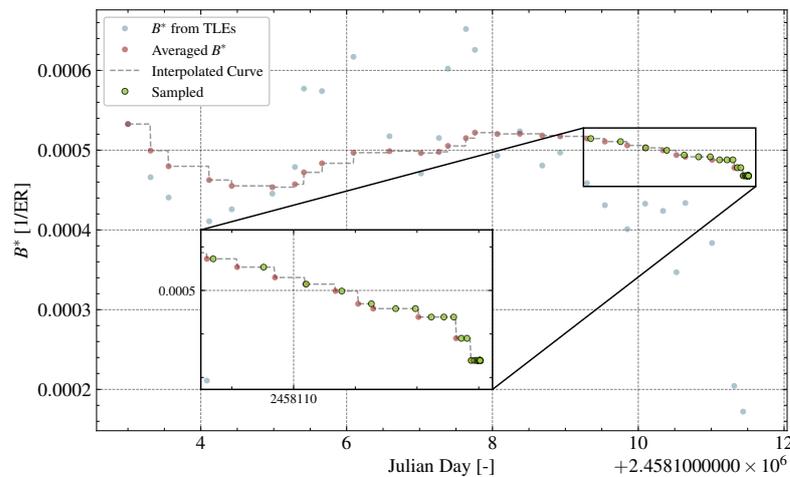


Figure 6. Representation of the pre-processing of the B^* feature generation: step-wise linear interpolation with moving window average.

Given the features generated so far, it is possible to apply the PCA to observe if the expected relations between the B^* , the residual time in days and average altitude exist. The results of the analysis are summarised in Figure 7. Focusing on the loading plot and the PC1, the feature related to the time has a strong positive loading; meanwhile, the others have negative loadings. Indeed, the time feature can be seen as a synonym of the residual lifetime, therefore as the B^* increases, the lifetime of an object decreases due to the drag acceleration growth. Furthermore, the altitude-time relation is associated with the altitude decrease, as time proceeds.

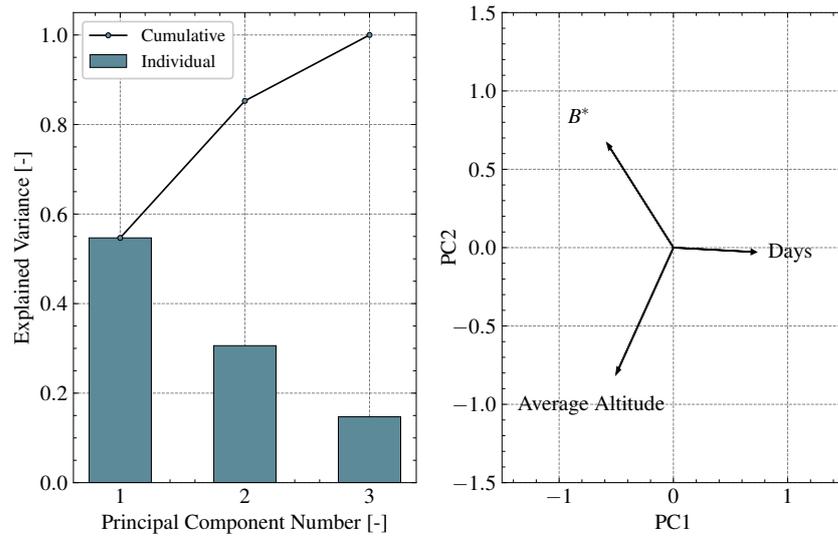


Figure 7. Variance plot (Left) and loading plot (Right) with B^* , residual time and altitude.

4.3. Solar Index

The solar index, $F_{10.7}$, is a parameter that describes the intensity of solar radiation. It is important for predicting the variations of the atmospheric density [9]. The data from TLEs do not include the atmospheric density; however, the B^* accounts also for the solar index variations. To highlight this dependence, Figure 8 represents the drag-like coefficient as a function of the last 81-day average of the solar index $\bar{F}_{10.7}$ for different bodies. It can be observed that, as $\bar{F}_{10.7}$ increases, the B^* grows with a non-linear behaviour, showing the existence of a relation between the variables. In addition, this is also highlighted by the same colours of B^* groups for the same values of $\bar{F}_{10.7}$, and by the terms of the curve fitting associated with $\bar{F}_{10.7}^2$. Therefore, the $\bar{F}_{10.7}$ index is introduced as a feature of the RNN, along with the area-to-mass ratio. This decision stems from the fact that the model must deal with objects with similar B^* but completely different solar indexes. From an implementation point of view, it was found that including only the $\bar{F}_{10.7}$ index corresponding to the starting epoch of the prediction led to the highest performance for the RNN. Therefore, it is included and repeated in the input sequence as a constant value. The same procedure is applied to the area-to-mass ratio, which is kept fixed.

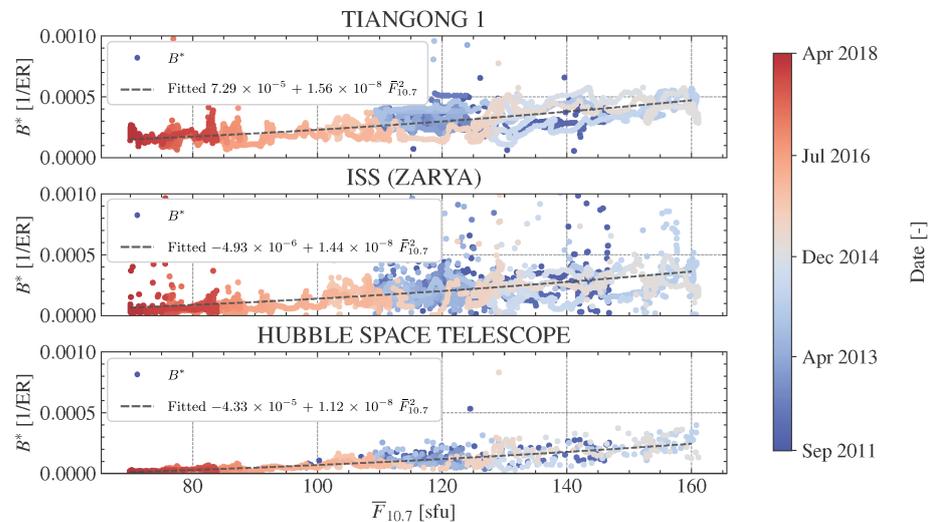


Figure 8. B^* as a function of the mean solar index for Tiangong 1, ISS and HST.

Similarly to Sections 4.1 and 4.2, it is possible to explore the features generated so far in order to find if the expected relations can be found. To this aim, the PCA can be applied considering the $\bar{F}_{10.7}$ index, the B^* and the area-to-mass ratio. The results are summarised in Figure 9. From the loading plot, it can be seen that, along PC1, all the variables are characterised by positive loadings. In particular, the relation between the A/m and the B^* is related to the drag-like coefficient definition. The relation between the B^* and the solar index is the same as discussed in Figure 8, where a positive correlation is present. Moreover, the relation between the area-to-mass ratio and the $\bar{F}_{10.7}$ can be explained through the B^* .

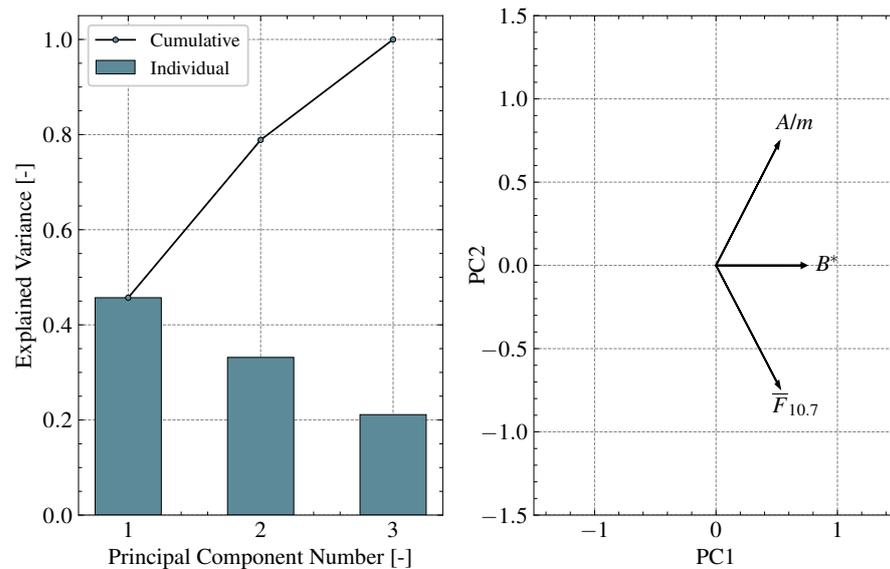


Figure 9. PCA characteristics of solar index, B^* and area-to-mass ratio.

4.4. Data Regularisation

The features generated so far include the average altitude profile, B^* , the $\bar{F}_{10.7}$ and the area-to-mass ratio. Due to their nature, the solar index and time-related feature are normalised according to [17]:

$$\bar{x} = \frac{x - \min(x)}{\max(x) - \min(x)}, \quad (19)$$

where \bar{x} represents the normalised feature; and x the related original input feature.

5. Results

The Seq2Seq architecture described in Section 2 was independently trained on four different cases, characterised by different starting altitudes, as shown in Table 3. In this way, the sensitivity of the prediction accuracy with respect to the starting altitude can be assessed.

Table 3. Characteristics of training cases.

Case	T_x [-]	Starting Altitude [km]	Training Epochs [-]
A	5	180	2900
B	9	160	3000
C	13	140	1200
D	17	120	1200

In each one of the cases in Table 3, the entire set of TLEs was split into two independent parts, the training and validation sets. Furthermore, both the inputs and outputs of the two sets have the same lengths. Therefore, the model was trained to predict the output sequence of length T_y , given an input sequence of length T_x , and the same strategy was also applied for validation. This means that the losses, \mathcal{L}_{MSE}^{Val} and $\mathcal{L}_{MSE}^{Train}$, were computed on the values associated with the same average altitudes. The only difference is represented by the objects analysed, which are different between the sets. The input dataset is mathematically represented by a tensor of rank 3, which, in TensorFlow's notation, has dimension [Objects \times Time \times Features]. Therefore, according to the problem definition, the model is faced only with a part of the input trajectory of the objects, together with all the related input features. The length of the time dimension can be chosen according to the desired needs, by varying the value of T_x , which is associated with the starting altitude of the prediction. Hence, the five different cases in Table 3 were identified by different lengths of the input sequence T_x . Regarding the output, the model has to predict the final part of the average trajectory altitude until the re-entry epoch. The sum of the input and output sequence is $T_x + T_y = 25$.

5.1. Hyperparameters Optimisation

The deep learning model described in Section 2 requires the definition of specific hyperparameters that are variables on which the learning algorithm relies for its optimisation and prediction processes. In general, hyperparameters may be different depending on the problem under exam; therefore, it is necessary to perform an informed selection of such parameters. This is usually achieved by optimising the hyperparameters in order to maximise the performances of the deep learning model. However, the hyperparameters optimisation process can be computationally intensive; therefore, for this work, it was decided to perform the optimisation only of Case A of Table 3. In fact, this case is characterised by the smallest input sequence and is therefore considered the most challenging of the four cases. The optimisation is designed using *HyperOpt* [31] module, that is, the Bayesian optimisation library, together with the Asynchronous Successive Halving (ASHA) algorithm [32] to combine random search with the early stopping of hyperparameters combinations that exhibit poor performances. These algorithms can be conveniently combined and parallelised using the Python package *Ray Tune* [33]. In this work, part of the hyperparameter space was kept constant and another part was optimised. Table 4 shows the fixed hyperparameters and their associated value. As shown in Equation (4), the loss function is the Mean Squared Error; the selected optimiser is Adam [34], which is a commonly used algorithm for deep learning model weights optimisation. It is a first-order gradient-based optimisation, based on adaptive estimates of first and second-order moments (i.e., mean and variance). As specified in Table 4, the decay rate of the first order moment, β_1 , and of the second order moment, β_2 , were set, together with the *Clipnorm* parameters, which forced a re-scaling of the gradient whenever its norm exceeds the specified value.

Table 4. Fixed hyperparameters.

	Parameter	Value [-]
	Loss	Mean Squared Error
	Optimiser	Adam
Fixed	β_1	0.999
	β_2	0.999
	Clipnorm	0.1

The remaining hyperparameters that were optimised are summarised in Table 5, together with the relevant search spaces and sampling mode. The parameter related to the layer is defined as the number of stacked GRU units of the encoder and decoder,

respectively. This means that if the number of layers is two, the model is going to be composed of four units, where two are associated with the encoder and the remaining with the decoder. Specifically, the final hidden state of each layer of the encoder is passed to their respective layer of the decoder. For example, the first GRU unit of the decoder is initialised with the hidden state of the first unit of the encoder.

Table 5. Optimised hyperparameters and their relevant search space in *HyperOpt*.

Hyperparameter	Interval [min, max]	Sampling Mode ¹
Learning Rate	$[1 \times 10^{-6}, 1 \times 10^{-1}]$	Uniform Logarithmic
Number of layers	[1, 3]	QUniform
Hidden size GRU	[16, 256]	QUniform
Batch Size	[8, 64]	QUniform
Decay Scheduled Sampling	$[1 \times 10^{-1}, 9.9 \times 10^{-1}]$	Uniform Logarithmic

¹ refers to the particular distribution from which the hyperparameter is selected. See http://hyperopt.github.io/hyperopt/getting-started/search_spaces/, accessed on 27 February 2023, for further details.

Finally, it is also necessary to specify the parameters for the ASHA algorithm. Table 6 shows the aforementioned parameters and the selected values. Specifically, the total number of trials refers to the amount of hyperparameters combinations the algorithm attempts before stopping, the reduction factor represents the fraction of trials removed after each run and the grace period that corresponds to the epochs for which a trial must be trained before halving. These values have to be selected considering the computational resources of the hardware and the expected results. The most suitable parameters, which were found with a trade-off between results and computational time, are represented in Table 6. It has to be noted that a small grace period may bias ASHA towards those solutions that converge quite rapidly but do not lead to the best performances. Therefore, a good compromise is experimentally found to be at 400 epochs. Furthermore, the maximum value of the training epoch is set at 2100 because it was observed that the optimal value is typically above 2000.

Table 6. Input parameters of ASHA.

Input	Value
Number of Trials	100 [-]
Reduction Factor	4 [-]
Grace Period	400 [Epochs]
Maximum Training Epoch	2100 [Epochs]

Figure 10 shows the validation losses as a function of the value of each hyperparameter. We can observe that the curves associated with lower losses are approximately distributed at low values of the batch size, between 20 and 40, and hidden state size, in particular between 50 and 150. Furthermore, these distributions provide the qualitative importance of each parameter. Indeed, an important parameter can be defined by observing the magnitude changes of the resulting loss with respect to a variation of its value. This means that a change, even small, in such an input parameter, determines an important variation of the resulting loss. For example, by looking at the learning rate in the parallel coordinates plot, the distinction between red and blue curves can be seen, where the lower rates lead to high validation losses. Instead, values between 0.002 and 0.0075 seem to lead to better performances. Contrarily, the decay rate of the Scheduled Sampling mechanism appears to have a negligible effect on the validation loss. Indeed, it is not possible to identify a range of values that leads to higher performances. This is related to the fact that the decay rate in Equation (8) does not have a significant importance on the convergence because the threshold for selecting the ground truth or the model output decreases quickly for all the

values that the optimiser selected, forcing the model to learn from its predicted outputs. Therefore, this may suggest that the model performs better when it has to deal with its own errors in a way that is more similar to inference than training.

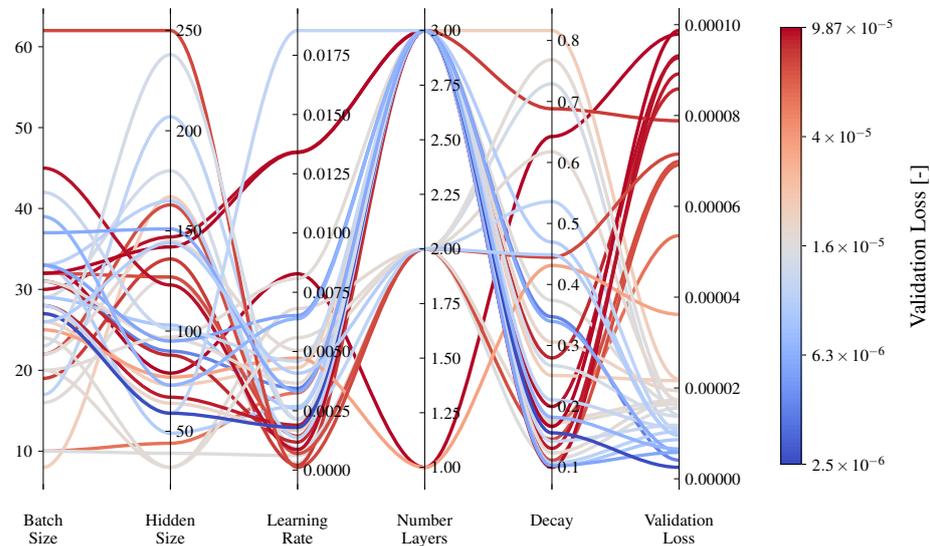


Figure 10. Representation of selected hyperparameter combinations as parallel lines that pass to the vertices corresponding to the hyperparameters associated with each trial, where each colour is related to the validation loss.

Finally, the optimised values of the hyperparameters are given in Table 7.

Table 7. Optimised hyperparameters.

	Parameter	Value [-]
Optimised	Learning Rate	0.001795
	Number of Layers	3
	Hidden Size GRU	59
	Batch Size	27
	Decay Scheduled Sampling	0.15665

5.2. Dataset Characterisation

To analyse and understand the performances of the model, the test objects have to be divided depending on their physical properties with respect to the training set. We defined two categories of objects as in Table 3 (Appendix A.3) in which the objects of the test set shared similar characteristics to the training set. Specifically, as shown in Figure 11, the two categories were subdivided with respect to the B^* coefficient: *Category 1* objects (red label) had a median B^* coefficient inside the 0.25 and 0.75 quantiles of the training set distribution, while the median B^* of objects in *Category 2* (black label) lay outside this range. Therefore, given the significant importance of B^* , it was expected that the deep learning model would perform better with those objects for which their distributions are similar to the training.

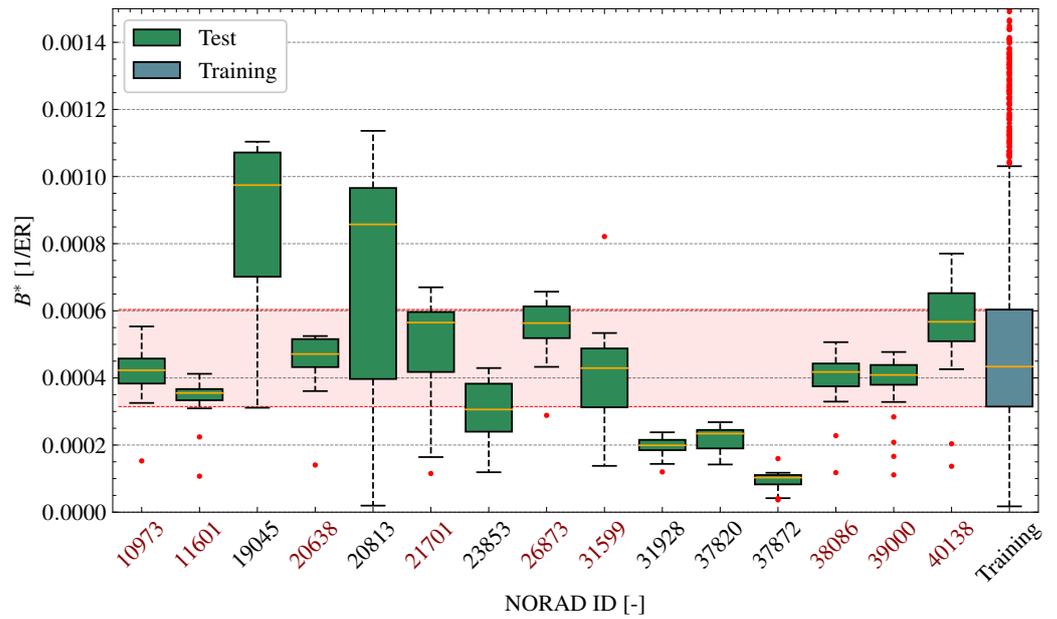


Figure 11. B^* distribution of objects in training and test sets, where red NORAD IDs identify test bodies with similar B^* distributions as the training set (Category 1).

To analyse and understand the performances of the model, the test objects have to be divided depending on their physical properties with respect to the training set. Hence, Figure 11 represents the B^* distributions of the objects in the test and training sets, where two categories of bodies can be identified, as highlighted by the different colours of the NORAD IDs. Another parameter to analyse is the eccentricity, the distributions of which for the training and test sets are represented in Figure 12. It can be seen that 20813 represents an anomaly with respect to all the other objects. Indeed, it is characterised by a median of roughly 0.7, which identifies it as a highly elliptical orbit, for which the nature of the re-entry is different. However, we decided to maintain this object to further test the capabilities and limitations of the developed deep learning model.

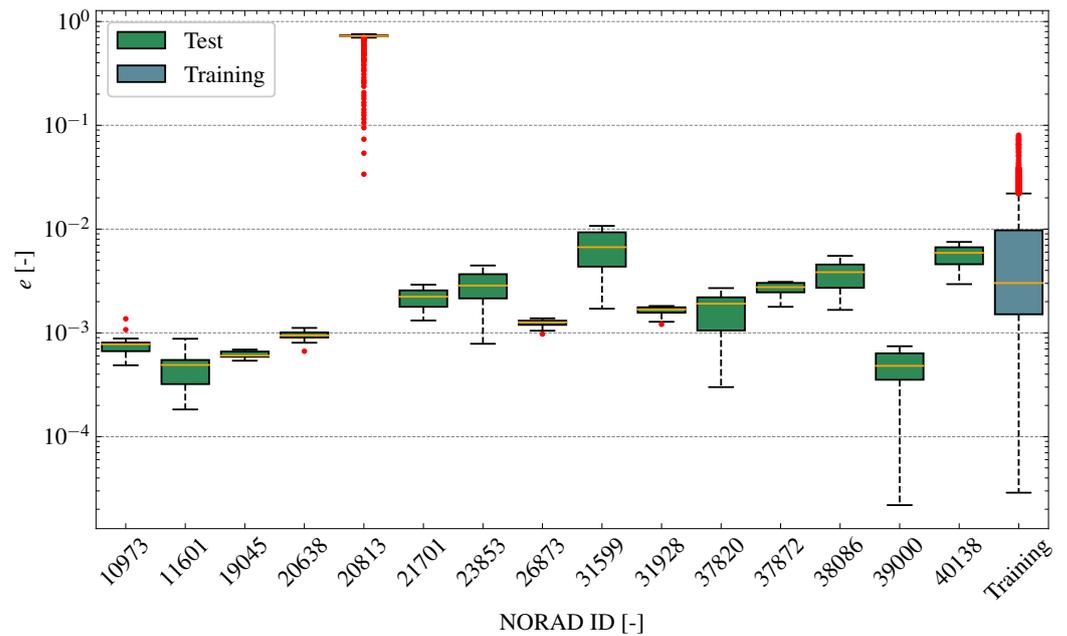


Figure 12. Eccentricity distribution of objects in training and test set.

To characterise the performances and compare the results, let the absolute error, ϵ_{abs} , and the relative error, ϵ_{rel} , be, respectively:

$$\epsilon_{abs} = |t_{Predicted} - t_{Actual}| \quad (20)$$

$$\epsilon_{rel} = \frac{|t_{Predicted} - t_{Actual}|}{t_{Actual} - t_{Initial}} \times 100, \quad (21)$$

where $t_{Predicted}$ is the predicted re-entry epoch; and t_{Actual} is the assessed re-entry epoch.

5.3. Testing

Starting from the objects of the first category, Table 8 represents the predictions' errors for all the different starting altitudes defined in Table 3. Generally, it can be seen that the model is successfully capable of predicting the output trajectory and therefore the re-entry epoch. For example, 39000 shows an absolute error that is less than 2 min, roughly 24 h before the re-entry. Furthermore, this error stays confined below roughly 3 min, for all the predictions. This behaviour is related to the similarities of the B^* distributions of each object with respect to the one of the training set. Indeed, it should be highlighted that the highest absolute error is 14 min and it is associated with 40138. Moreover, observing the MSEs, all the relative errors lay between 10^{-5} and 10^{-7} , showing that the similarities of the B^* are associated with excellent performances. It is interesting to note that, comparing the different cases, the MSE decrease does not necessarily correspond to a lower absolute error, as for example in 11601 and 39000. This indicates that better results could also be obtained by selecting a different loss function.

Table 8. Results of objects in category 1.

Case	Metric	NORAD								
		10973	11601	20638	21701	26873	31599	38086	39000	40138
Case A	ϵ_{abs} [hrs]	0.0129	0.0012	0.0593	0.0287	0.0381	0.0835	0.0856	0.0263	0.2225
	ϵ_{rel} [%]	0.0582	0.0051	0.2930	0.1262	0.1800	0.5004	0.3469	0.1093	1.2464
	MSE [day ²]	4.4482×10^{-6}	1.1707×10^{-6}	1.0218×10^{-5}	1.0236×10^{-5}	1.7067×10^{-6}	5.2804×10^{-6}	2.1809×10^{-5}	1.0643×10^{-6}	1.9246×10^{-5}
Case B	ϵ_{abs} [hrs]	0.0915	0.0411	0.0417	0.0021	0.0376	0.0640	0.0187	0.0191	0.0191
	ϵ_{rel} [%]	0.9903	0.4878	0.6032	0.0265	0.4828	1.0072	0.2198	0.2134	0.2952
	MSE [day ²]	7.0168×10^{-6}	1.8629×10^{-6}	2.6473×10^{-6}	4.4563×10^{-6}	2.6635×10^{-7}	1.2836×10^{-6}	1.1987×10^{-6}	1.7427×10^{-6}	1.8259×10^{-6}
Case C	ϵ_{abs} [hrs]	0.0934	0.0043	0.0125	0.0268	0.0610	0.1025	0.0368	0.0065	0.0198
	ϵ_{rel} [%]	2.8694	0.1839	0.7061	1.3862	2.8083	5.6717	1.7670	0.2461	1.1732
	MSE [day ²]	1.8350×10^{-5}	2.1482×10^{-6}	9.5146×10^{-7}	4.1151×10^{-6}	1.5597×10^{-6}	5.9912×10^{-6}	4.7233×10^{-6}	1.5328×10^{-6}	6.7500×10^{-7}
Case D	ϵ_{abs} [hrs]	0.0670	0.0560	0.0244	0.0340	0.0581	0.0271	0.0580	0.0420	0.0093
	ϵ_{rel} [%]	8.1350	14.1151	9.1041	13.5330	15.8149	8.8874	20.3210	8.3022	3.7092
	MSE [day ²]	6.4009×10^{-6}	3.1455×10^{-6}	1.3601×10^{-6}	1.3073×10^{-6}	2.2189×10^{-6}	1.0388×10^{-6}	2.2950×10^{-6}	2.1382×10^{-6}	1.5021×10^{-6}

Regarding the objects in the second category, the errors for all the bodies are listed in Table 9. It can be generally seen that the performances of the model are inferior with respect to the previous category of objects, due to the low number of bodies characterised by B^* distributions similar to the one of the training set. This can be clearly observed in 37872, which can not be associated with any B^* distributions in the sets. Hence, the resulting outputs are characterised by the highest MSEs, due to the fact that the model tends to overestimate all the relative epochs. However, it can be demonstrated that the output follows the shape of trajectory altitude quite accurately. Furthermore, it can be noted that 20813 is characterised by the highest relative errors and equally significant MSEs. Hence, this indicates that the model is not capable of predicting the output trajectory and it completely misses all the predicted epochs. Indeed, by inspecting Figure 11, 20813 is characterised by a large distribution of B^* values, which maximum is similar to 19045. However, the model was shown to follow the output profile reasonably well for 19045 in all the cases, with a maximum relative error of roughly 18 min and MSE between 10^{-5} and

10^{-7} [days²]. However, object 20813 also has a large value of eccentricity (see Figure 12, which justifies the difficulties of the deep learning model in predicting its re-entry and indicates that, as expected, the contribution of the eccentricity cannot be ignored for orbits that are not quasi-circular.

Table 9. Results of objects in category 2.

Case	Metric	NORAD					
		19045	20813	23853	31928	37820	37872
Case A	ε_{abs} [hrs]	0.2994	1.2758	0.3083	0.6206	0.5829	1.3427
	ε_{rel} [%]	2.5476	964.0562	0.8570	1.7715	1.6664	1.6942
	MSE [day ²]	6.1088×10^{-5}	1.1331×10^{-3}	9.9711×10^{-5}	3.6330×10^{-4}	3.2982×10^{-4}	1.0313×10^{-2}
Case B	ε_{abs} [hrs]	0.0355	0.7159	0.0273	0.0854	0.0894	2.5015
	ε_{rel} [%]	0.8068	1309.4202	0.2152	0.7246	0.7297	9.3210
	MSE [day ²]	1.7582×10^{-6}	2.6453×10^{-4}	3.2465×10^{-6}	7.4105×10^{-6}	5.7928×10^{-6}	1.2797×10^{-2}
Case C	ε_{abs} [hrs]	0.0029	0.2194	0.0033	0.0080	0.0222	1.9536
	ε_{rel} [%]	0.2301	2254.3293	0.1022	0.2579	0.6580	31.1668
	MSE [day ²]	1.1096×10^{-7}	6.2342×10^{-5}	6.7083×10^{-6}	1.6303×10^{-5}	1.5640×10^{-5}	1.1795×10^{-2}
Case D	ε_{abs} [hrs]	0.0805	0.0098	0.1186	0.0847	0.0788	2.9645
	ε_{rel} [%]	35.9520	2532.4688	26.520	16.8331	13.4101	373.4046
	MSE [day ²]	9.3967×10^{-6}	2.5700×10^{-5}	2.4415×10^{-5}	2.3780×10^{-5}	1.9272×10^{-5}	1.5264×10^{-2}

Orbit Prediction—Case A

In this section, we show in detail the results for Case A. Specifically, we present the evolution of the training and validation losses and the results of the trajectory predictions of the developed deep learning model for selected objects in the test set. For the test case, both categories of objects (Section A.3) were included. Figure 13 shows the training and validation losses as a function of the training epoch, for a total of 2900 epochs. We observe how the rate of convergence of the validation loss is much lower than the training loss; nonetheless, the final loss is below 1×10^{-5} , which has been observed to produce accurate results.

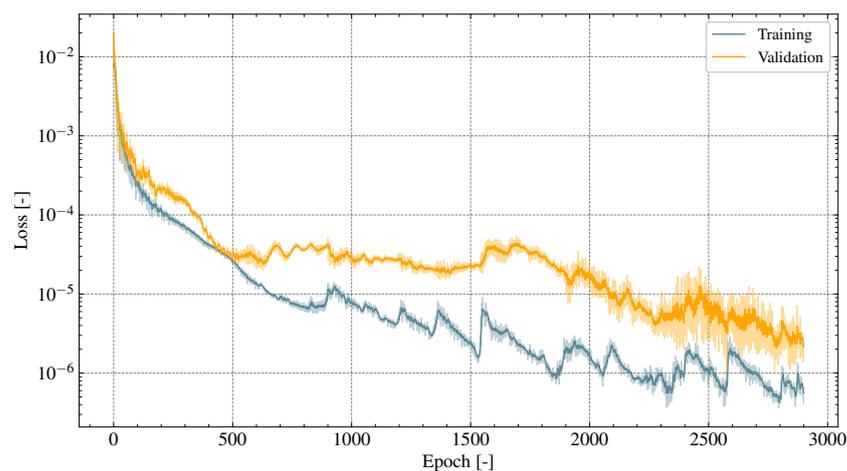


Figure 13. Representation of the training and validation loss curves in case A, where each curve is represented through its exponentially weighted average (coloured) and its original values (shaded).

Figure 14 shows the trajectory prediction for selected objects of Category 1 in the test set. The black curve represents the input average altitude (the initial 5 elements of the input sequence, T_x), the red line represents the predicted sequence, and the blue trajectory defines the true output (as obtained from TLE data).

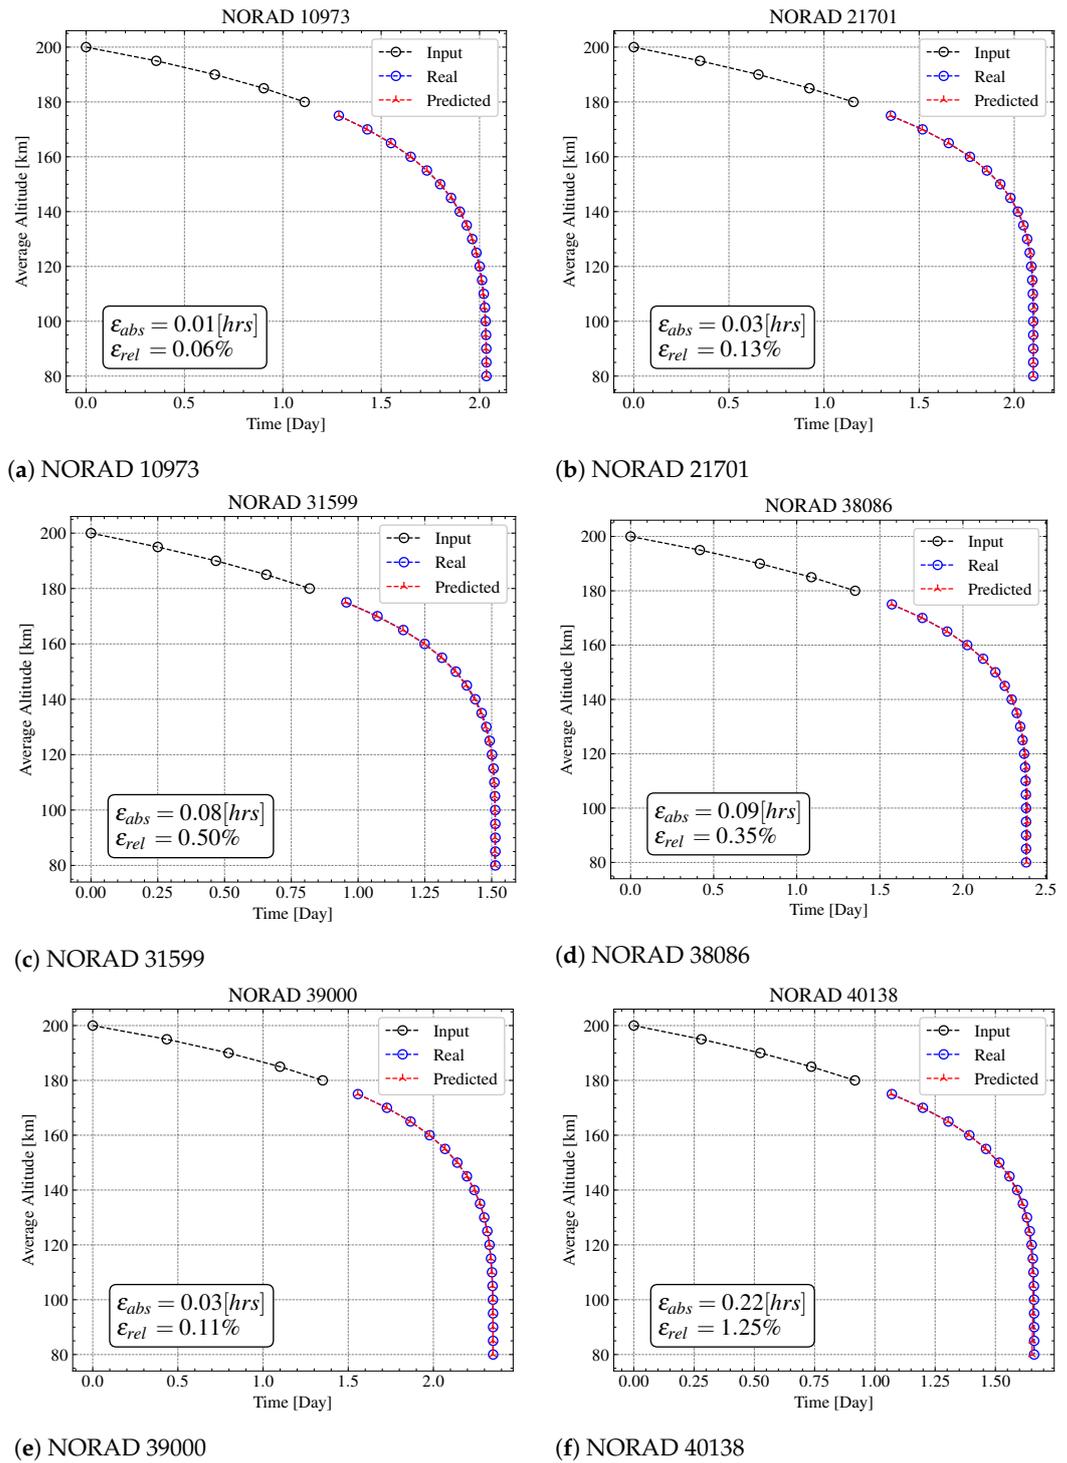


Figure 14. Predictions of selected objects in category 1 from 180 km altitude.

In general, it can be seen that the model is capable of predicting the re-entry epoch with excellent results. The model can successfully estimate the output trajectory of each object because the outputs follow the true paths quite accurately. Indeed, as also shown in Table 8, the highest MSE is associated with 38086; however, the relative errors on the re-entry epoch are small. Therefore, even if the MSE appears to be relatively high, the absolute and relative errors on the re-entry epoch are small. Moreover, the absolute errors are all lower than approximately 13 min, except for 40138, which is characterised by an error of 13.35 min. This can be explained by the higher B^* distribution of 40138 with respect to the

training set. Furthermore, considering objects 40138 and 10973, it appears that the model has correctly assimilated the knowledge related to the relation between the solar index and the drag-like coefficient.

Concerning the objects with B^* distributions different from the one of the training set (Category 2), the results are summarised in Figures 15 and 16.

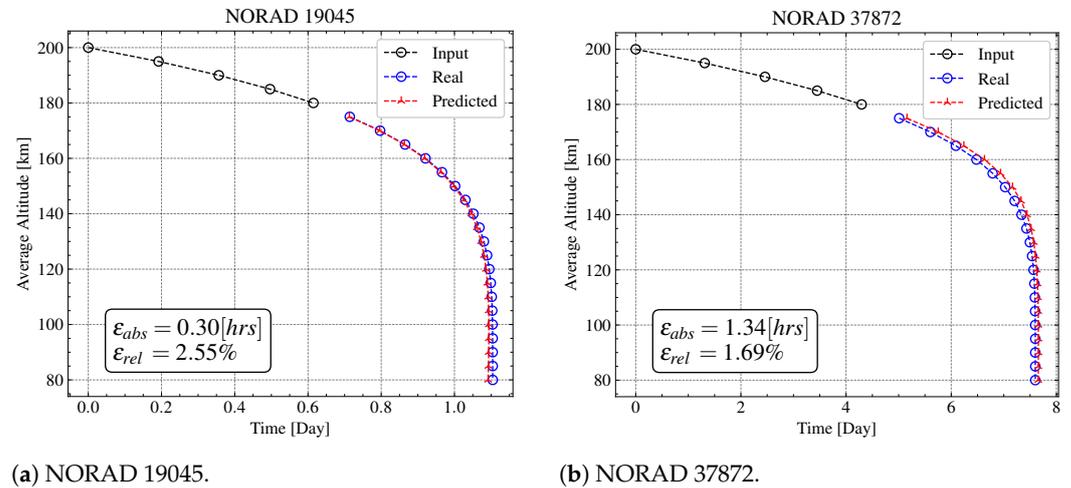


Figure 15. Predictions of 19045 and 37872 from 180 km altitude.

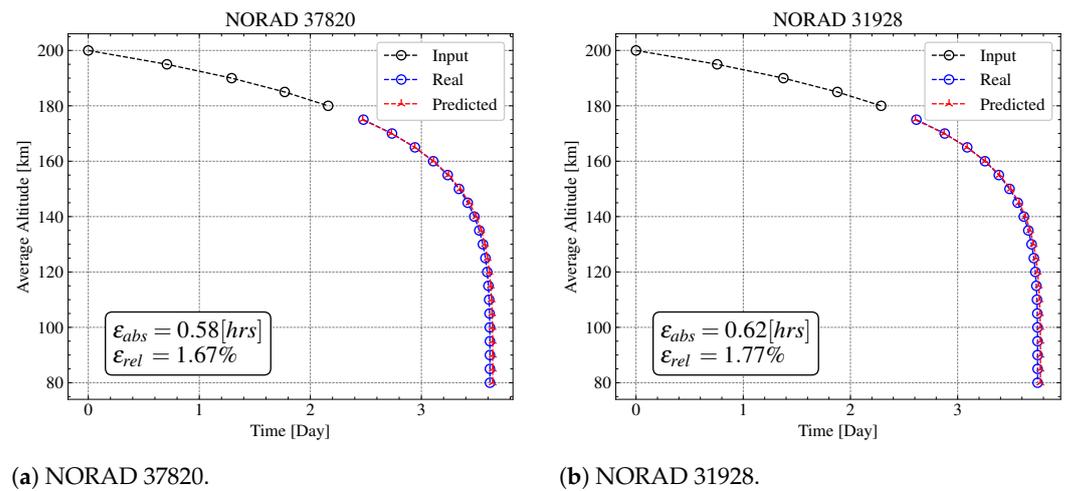


Figure 16. Predictions of 37820 and 31928 from 180 km altitude.

Generally, it can be observed that all errors are higher with respect to the previous case. This is related to the low number of objects characterised by a B^* distribution similar to the training, which makes the model perform less accurately. However, in four out of seven cases, the absolute errors are lower than approximately 40 min. Therefore, it can be concluded, in this case, that the model is capable of approximating the true output reasonably well. Nevertheless, object 20813 represents an exception in terms of errors, as can be seen in Figure 17, where the output trajectory is completely distorted.

By inspecting Figure 11, it can be seen that the B^* distribution of 20813 is characterised by a large range of values. However, its interquartile is lower than 19045, and the model has proved itself to provide acceptable performances on the latter, with an output that accurately follows the true one. Therefore, it appears that the drag-like coefficient cannot be directly linked to bad performance. Indeed, 20813 is characterised by high values of eccentricity, as highlighted in Figure 12. This makes the nature of the re-entry different, and the model is not capable of acquiring this knowledge because it is not trained for

it. Therefore, the prediction capabilities have to be related not only to the B^* but also to the eccentricity.

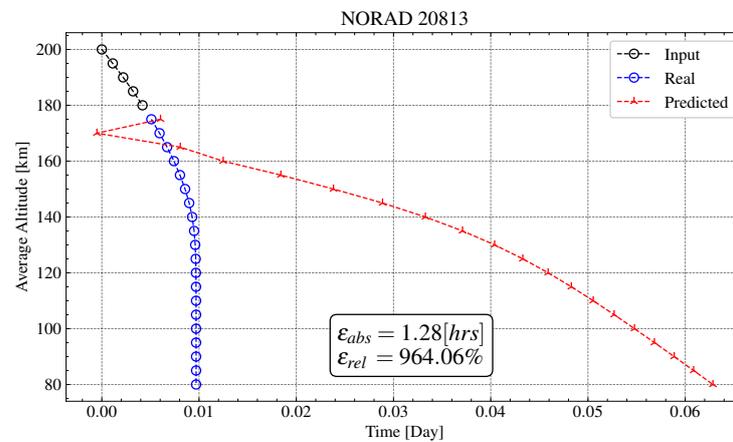


Figure 17. Prediction of NORAD 20813 from 180 km altitude.

6. Conclusions and Discussion

This work presents the development of a deep learning model that is based on a variant of the Seq2Seq architecture using curriculum learning for the re-entry prediction of uncontrolled objects in Low Earth Orbit based on TLE data. We, therefore, propose a paradigm shift from a physical to a data-driven approach. Based on the physical knowledge of the re-entry phenomenon, we introduced additional features alongside the average altitude profile, in an effort to better characterise and predict the re-entry epoch. Two additional key features were identified: the drag-like coefficient, B^* , and the solar index, $\bar{F}_{10.7}$. The drag-like coefficient is of key importance as it relates the re-entry time to the interaction of the object with the atmosphere. In our work, we present a methodology to introduce such a parameter as a feature into a deep learning model, via a stepwise moving window interpolation scheme. The solar index is equivalently relevant in the re-entry prediction as it affects the density of the atmosphere and, with it, the re-entry time. Analysing this parameter, the need to introduce an additional feature has arisen, in order to relate the solar index to the physical properties of the object. Therefore, the area-to-mass ratio was also added as a feature. Studying the behaviour of the deep learning model with respect to these features, we identified possible ways to introduce them in the model. The developed deep learning model was tested on a set of objects studied during the IADC campaigns. It was observed, in accordance with Jung et al. [12], that the B^* distribution has significant importance for the performances of the model. However, the poor performance related to object 20813 (Figure 17) suggests that eccentricity also has a relevant impact. Hence, future works may include the processing of highly eccentric orbits. In addition, it has been found that a decrease in the selected loss function (MSE) did not necessarily correspond to higher accuracy of the re-entry prediction. Consequently, future studies could also investigate the effect of different loss functions on the output of the deep learning model. Overall, the analyses showed promising results for re-entry prediction using machine learning; however, as can be expected, the prediction capabilities of the model strongly depend on the quality and quantity of the training data. It is therefore crucial, in order to expand the capabilities of this data-driven approach, to expand the size of the training data and, potentially, improve the quality of such data by considering additional sources other than TLEs.

Author Contributions: Conceptualization, M.T., F.S. and C.C.; methodology, F.S. and M.T.; software, F.S.; validation, F.S.; formal analysis, F.S.; investigation, F.S.; writing—original draft preparation, M.T. and F.S.; writing—review and editing, M.T. and C.C.; visualization, F.S.; supervision, M.T. and C.C.; funding acquisition, C.C. All authors have read and agreed to the published version of the manuscript.

Funding: The research presented in this thesis received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme as part of the COMPASS project (Grant agreement No 679086).

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

TLE	Two Line Element
TIP	Tracking and Impact Prediction
PCA	Principal Component Analysis
Seq2Seq	Sequence to Sequence
TPE	Tree-structured Parzen Estimator
ASHA	Asynchronous Successive Halving Algorithm
ANN	Artificial Neural Network
RNN	Recurrent Neural Network
GRU	Gated Recurrent Unit
MSE	Mean Squared Error
OMM	Orbital Mean-Elements Message

Appendix A

Appendix A.1. Norad IDs of Training Set

22277, 33594, 38223, 43538, 49130, 38976, 43033, 37727, 12389, 28883, 45942, 45729, 33440, 40459, 28824, 45601, 39573, 41813, 25431, 29504, 25577, 26481, 12072, 40714, 40700, 33448, 41640, 43703, 40120, 43586, 1377, 37346, 28701, 44500, 36604, 39262, 28642, 38349, 28842, 33387, 40085, 28827, 39561, 24968, 28400, 49045, 28822, 34670, 24840, 38048, 25288, 33592, 48853, 34906, 28486, 28424, 29480, 41478, 39259, 39683, 25274, 27451, 39147, 28641, 25291, 3835, 46930, 39507, 27375, 24837, 31602, 10861, 46390, 25285, 28655, 40949, 40452, 41313, 38075, 8063, 41628, 28504, 12071, 43705, 43475, 25173, 22781, 24872, 40589, 42718, 39145, 40897, 43007, 29058, 33507, 40096, 28896, 40728, 44210, 28625, 32267, 37674, 37397, 39033, 37874, 39566, 28777, 25064, 11745, 34872, 40957, 43243, 25342, 44228, 32007, 39409, 32377, 42701, 44070, 36835, 41865, 46397, 36096, 27372, 39178, 26897, 42734, 29247, 39116, 33456, 43166, 32261, 40745, 25647, 27474, 23769, 40421, 39560, 40430, 40740, 25275, 35642, 41821, 27450, 31596, 26874, 49065, 44551, 41819, 44385, 27374, 48804, 32485, 40723, 40453, 25290, 41100, 32269, 39193, 24945, 37883, 37821, 43532, 29053, 40314, 39580, 41559, 28933, 40668, 40945, 32270, 46614, 43494, 27642, 28823, 28362, 41480, 43757, 14207, 29157, 39571, 33054, 41776, 28425, 40363, 28878, 37197, 29109, 41909, 28519, 38848, 39503, 32757, 37172, 12054, 34603, 21694, 42723, 43863, 33435, 45112, 25468, 39527, 40127, 45596, 28099, 11474, 29667, 40702, 38462, 10582, 38074, 11849, 14694, 27392, 43691, 44372, 36130, 21931, 38249, 25276, 39623, 25040, 25106, 32477, 30778, 29080, 42757, 41486, 39519, 32262, 37860, 27841, 43090, 31394, 28867, 44844, 43239, 32002, 39126, 44707, 35694, 41596, 37680, 41769, 25171, 23020, 39514, 29229, 29263, 42731, 38862, 32751, 29253, 47619, 44438, 32059, 25287, 39393, 20299, 27700, 43212, 35818, 39031, 38037, 35867, 29403, 41574, 28363, 33457, 44206, 24794, 37255, 44834, 41777, 40727, 27550, 28907, 37217, 35947, 37942, 29653, 43667, 39569, 39513, 25777, 7338, 39062, 35011, 36362, 41125, 37634, 42703, 43244, 29394, 40898, 40126, 43637, 44826, 39220, 42972, 25263, 41027, 39171, 24839, 45604, 36506, 27706, 24905, 41487, 37758, 43027, 46669, 28403, 25530, 41572, 44799, 29386, 39180, 39149, 41597, 47255, 28130, 29246, 39386, 37878, 47348, 42730, 41387, 41178, 35941, 38672, 24966

Appendix A.2. Norad IDs of Validation Set

41868, 45177, 29488, 39682, 60, 32386, 38336, 33332, 38550, 24792, 48160, 36522, 32284, 39649, 35949, 43064, 41636, 28191, 40247, 6212, 25346, 37362, 11671, 28643, 44317, 42938, 24904, 43129, 11332, 28445, 27475, 41107, 40393, 44505, 40313, 36512, 48866, 36087, 37873,

37729, 42716, 39363, 48870, 44111, 40359, 13402, 39187, 33273, 47782, 26405, 37856, 28506, 41455, 39408, 39493, 41475, 39137, 24869, 34840, 42733, 41629, 48275, 23757, 46463, 40361, 40619, 28997, 42702, 42800, 28942, 44075, 36749, 32713, 33449, 38076, 37876, 41449, 28774, 31798, 41483, 45938, 39532, 37383, 40886

Appendix A.3. Norad IDs of Test Set

20813, 19045, 11601, 10973, 40138, 38086, 23853, 21701, 37872, 39000, 20638, 31928, 37820, 26873, 31599

Test Set Subdivision in Category 1 and Category 2

Table A1 shows the NORAD IDs of the test set objects subdivided into Category 1 and Category 2. Category 1 contains objects with a median B^* within the interquartile distance of the B^* distribution of the training set, while objects in Category 2 have a median B^* outside this interval.

Table A1. Subdivision of the bodies in the test set according to the B^* distributions.

Category	NORAD IDs									
1	10973	11601	20638	21701	26873	31599	38086	39000	40138	
2	19045	20813	23853	31928	37820	37872				

References

1. ESA Space Debris Office. *ESA'S Annual Space Environment Report*; ESA Space Debris Office: Darmstadt, Germany, 2021.
2. Alby, F.; Alwes, D.; Anselmo, L.; Baccini, H.; Bonnal, C.; Crowther, R.; Flury, W.; Jehn, R.; Klinkrad, H.; Portelli, C.; et al. The European Space Debris Safety and Mitigation Standard. *Adv. Space Res.* **2004**, *34*, 1260–1263. [\[CrossRef\]](#)
3. NASA. *Process for Limiting Orbital Debris, NASA-STD-8719.14B*; NASA: Washington, DC, USA, 2019.
4. Pardini, C.; Anselmo, L. Re-entry predictions for uncontrolled satellites: Results and challenges. In Proceedings of the 6th IAASS Conference "Safety Is Not an Option", Montreal, QC, Canada, 21–23 May 2013.
5. Pardini, C.; Anselmo, L. Performance evaluation of atmospheric density models for satellite reentry predictions with high solar activity levels. *Trans. Jpn. Soc. Aeronaut. Space Sci.* **2003**, *46*, 42–46. [\[CrossRef\]](#)
6. Pardini, C.; Anselmo, L. Assessing the risk and the uncertainty affecting the uncontrolled re-entry of manmade space objects. *J. Space Saf. Eng.* **2018**, *5*, 46–62. [\[CrossRef\]](#)
7. Braun, V.; Flegel, S.; Gelhaus, J.; Kebschull, C.; Moeckel, M.; Wiedemann, C.; Sánchez-Ortiz, N.; Krag, H.; Vörsmann, P. Impact of Solar Flux Modeling on Satellite Lifetime Predictions. In Proceedings of the 63rd International Astronautical Congress, Naples, Italy, 1–5 October 2012.
8. Vallado, D.A.; Finkleman, D. A critical assessment of satellite drag and atmospheric density modeling. *Acta Astronaut.* **2014**, *95*, 141–165. [\[CrossRef\]](#)
9. Vallado, D. *Fundamentals of Astrodynamics and Applications*, 2nd ed.; Springer: Dordrecht, The Netherlands, 2001.
10. Anselmo, L.; Pardini, C. Computational methods for reentry trajectories and risk assessment. *Adv. Space Res.* **2005**, *35*, 1343–1352. [\[CrossRef\]](#)
11. Frey, S.; Colombo, C.; Lemmens, S. Extension of the King-Hele orbit contraction method for accurate, semi-analytical propagation of non-circular orbits. *Adv. Space Res.* **2019**, *64*, 1–17. [\[CrossRef\]](#)
12. Jung, O.; Seong, J.; Jung, Y.; Bang, H. Recurrent neural network model to predict re-entry trajectories of uncontrolled space objects. *Adv. Space Res.* **2021**, *68*, 2515–2529. [\[CrossRef\]](#)
13. Lidtke, A.A.; Gondelach, D.J.; Armellin, R. Optimising filtering of two-line element sets to increase re-entry prediction accuracy for GTO objects. *Adv. Space Res.* **2019**, *63*, 1289–1317. [\[CrossRef\]](#)
14. Flohrer, T.; Krag, H.; Klinkrad, H. Assessment and categorization of TLE orbit errors for the US SSN catalogue. In Proceedings of the Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, USA, 17–19 September 2008.
15. Levit, C.; Marshall, W. Improved orbit predictions using two-line elements. *Adv. Space Res.* **2011**, *47*, 1107–1115. [\[CrossRef\]](#)
16. Aida, S.; Kirschner, M. Accuracy assessment of SGP4 orbit information conversion into osculating elements. In Proceedings of the 6th European Conference on Space Debris, Darmstadt, Germany, 22–25 April 2013.
17. Raschka, S. *Python Machine Learning*, 3rd ed.; Packt Publishing Ltd.: Birmingham, UK, 2019.
18. Sutskever, I.; Vinyals, O.; Le, Q.V. Sequence to Sequence Learning with Neural Networks. In Proceedings of the 27th International Conference on Neural Information Processing Systems, Montreal, Canada, 8–13 December 2014; MIT Press: Cambridge, MA, USA, 2014; Volume 2, pp. 3104–3112.
19. Zhang, A.; Lipton, Z.C.; Li, M.; Smola, A.J. Dive into Deep Learning. *arXiv* **2021**, arXiv:2106.11342.
20. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, UK, 2016.
21. Elman, J.L. Finding structure in time. *Cogn. Sci.* **1990**, *14*, 179–211. [\[CrossRef\]](#)

22. Cho, K.; van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; Bengio, Y. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; Association for Computational Linguistics: Doha, Qatar, 2014; pp. 1724–1734. [[CrossRef](#)]
23. Bengio, S.; Vinyals, O.; Jaitly, N.; Shazeer, N. Scheduled sampling for sequence prediction with recurrent neural networks. *Adv. Neural Inf. Process. Syst.* **2015**, *28*, 1171–1179.
24. Chung, J.; Gulcehre, C.; Cho, K.; Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv* **2014**, arXiv:1412.3555.
25. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems (Software). 2015. Available online: <https://www.tensorflow.org> (accessed on 10 January 2022).
26. Chollet, F. *Deep Learning with Python*; Simon and Schuster: New York, NY, USA, 2021.
27. Salmaso, F. Machine Learning Model for Uncontrolled Re-Entry Predictions of Space Objects and Feature Engineering. Ph.D. Thesis, Politecnico di Milano, Milan, Italy, 2022.
28. Dong, G.; Liu, H. *Feature Engineering for Machine Learning and Data Analytics*; CRC Press: Boca Raton, FL, USA, 2018.
29. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [[CrossRef](#)] [[PubMed](#)]
30. Curtis, H.D. *Orbital Mechanics for Engineering Students*, 3rd ed.; Butterworth-Heinemann: Oxford, UK, 2014.
31. Bergstra, J.; Yamins, D.; Cox, D. Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures. In Proceedings of the 30th International Conference on Machine Learning, Atlanta, GA, USA, 16–21 June 2013; Volume 28, pp. 115–123.
32. Li, L.; Jamieson, K.; Rostamizadeh, A.; Gonina, E.; Hardt, M.; Recht, B.; Talwalkar, A. A System for Massively Parallel Hyperparameter Tuning. *arXiv* **2020**, arXiv:1810.05934.
33. Liaw, R.; Liang, E.; Nishihara, R.; Moritz, P.; Gonzalez, J.E.; Stoica, I. Tune: A research platform for distributed model selection and training. *arXiv* **2018**, arXiv:1807.05118.
34. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.