Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

E-Pilots: A Machine Learning System Deployable in the Cockpit for Predicting Early Hard Landings in Commercial Aviation

Gunn Bansal[1], Tushar Gupta[2], Tanya Bansal[3], Sneha Dhyani[4], Dr. Satendra Kumar[5]

[1,2,3,4,5] Dept. of Computer Science and Engineering, Meerut Institute of Engineering and Technology Meerut, Uttar Pradesh, India

Abstract
Hard landings are a major and ongoing danger in commercial flights. They can cause damage to the aircraft's structure, raise maintenance costs, and even injure passengers or crew. Currently, safety regulations mainly use data collected after a flight to detect hard landings, which means problems are only identified after they occur, with no opportunity to prevent them in real time. This study introduces "E-Pilots," a new system designed for use in the cockpit. It is engineered to determine the likelihood of a hard landing while the aircraft is on approach. E-Pilots uses machine learning to analyse real-time flight parameters such as vertical speed, descent angle, and deviation from the correct flight path. The system is trained on a large volume of historical flight data, enabling it to identify subtle indicators of a potentially hazardous landing. Early results demonstrate that E-Pilots can predict hard landings with a high degree of accuracy, providing pilots with timely warnings. This would enable flight crews to modify their landing plan or execute a go-around, thereby improving flight safety. This paper explains the system architecture, the data requirements, and the methods employed. It concludes that a system such as E-Pilots could transform landing safety management and significantly reduce the frequency of hard landings.

Keywords: Hard landings, Aviation safety, Real-time prediction, Machine learning, Flight data analysis, Cockpit deployable system, Predictive modeling.

1. Introduction

A large percentage of fatal incidents worldwide occur during the approach and landing stages of commercial flights, which are statistically the most hazardous phases. In particular, the final approach and landing phase accounted for 49% of fatal commercial aircraft accidents between 2008 and 2017. Runway excursions, which are frequently preceded by a number of contributing factors, have been identified as a significant problem by safety bodies, notably the European Union Aviation Safety Agency (EASA).

An unstable approach, in which the aircraft fails to meet stabilisation criteria by a specific altitude, is frequently a precursor to accidents and incidents. Importantly, Boeing reported that although only 3% of approaches were unstable, 97% of them proceeded to land instead of executing a go-around. According to studies, a prompt go-around decision may have prevented up to 83% of runway excursion incidents.

Consequently, deciding quickly to perform a go-around manoeuvre is essential to reducing the overall accident rate in the aviation sector.

There is a clear and acknowledged need for onboard, real-time performance monitoring and warning systems to assist flight crew in this critical, high-workload decision. The large volumes of data gathered by aviation systems, combined with advances in machine learning, create an opportunity to develop such systems.

To address this need, this paper presents E-Pilots, a machine learning system deployable in the cockpit, designed to predict a Hard Landing (HL) event during the approach phase of commercial flights and to recommend a go-around manoeuvre. When the vertical acceleration at touchdown exceeds a defined threshold (for example, greater than 2G for Airbus aircraft), the event is classified as a hard landing and necessitates mandatory maintenance inspection.

Although hard landing prediction has been studied previously, most prior work focuses on UAVs (unmanned aerial vehicles) or employs prediction windows that are typically within 9 to 2 metres before touchdown. These windows are unsuitable for commercial aircraft, because the critical Decision Height (DH) required for a go-around is approximately 100 feet (30.5 metres). Furthermore, current deep learning methods have been criticised for being tested on a single aircraft type and for training on balanced datasets, which distorts the natural frequency of HL occurrences (approximately 3 to 4% of flights) and may produce misleadingly optimistic accuracy estimates.

This study proposes a hybrid technique for early HL prediction based on features describing the temporal interdependence of aircraft state variables, under the assumption that a hard landing event is preceded by observable antecedents. The technique combines an optimised neural network architecture with input features derived from the variability of aircraft parameters over specific altitude windows, including the Decision Height. The deployability of this methodology is evaluated through rigorous testing on a large dataset of 58,177 commercial flights spanning three distinct aircraft types: the A319, A320, and A321.

The following is a summary of this paper's primary contributions:

(i) Hybrid Model with Optimised Design: To overcome the complexity and convergence issues of models such as LSTM, a hybrid model is proposed. It uses the standard deviations of aircraft parameters across an isolated range of altitude windows as inputs to an optimised neural network architecture.

(ii) Superior Performance on a Large, Multi-Aircraft Dataset: The proposed method outperforms current LSTM techniques, achieving an average sensitivity (recall) of 85% and an average specificity of 75% at the go-around point. It is evaluated on an extensive dataset of 58,177 flights encompassing the A319, A320, and A321 aircraft types.

(iii) Analysis of Classifier vs. Regressor: It is demonstrated that classification models are substantially better suited for early identification and go-around recommendations for HL events than regression models.

(iv) Go-Around Recommendation Capability: The system's feasibility as a cockpit-deployed recommendation engine is demonstrated by its ability to identify HL events prior to the Decision Height.

The remainder of this paper is structured as follows. Related research, including existing classifier and regressor techniques, is reviewed in Section 2. The methodology, including the dataset, variable selection, and network models, is described in Section 3. Experimental findings regarding predictive power and deployability are presented in Section 4. Section 5 discusses the results, compares the approach against existing methods, and proposes future improvements. The paper is concluded in Section 6.

2. Related Work

In recent years there has been growing research on data-driven prediction of hazardous landing events, with several studies examining anomalies such as abnormal approach profiles, long landings, and runway excursions [9]–[12]. Nevertheless, the specific challenge of hard landing (HL) prediction in commercial aviation remains comparatively underexplored. The majority of prior research has focused on unmanned aerial vehicles (UAVs), whose flight dynamics, operating protocols, and aerodynamic characteristics differ significantly from those of commercial aircraft; consequently, these methods are not directly applicable to airline operations.

According to a survey of earlier studies, the majority of research on landing-anomaly prediction appears in conference papers, with very few journal articles, as illustrated in Figure 1. This disparity suggests that the field is still developing and that many approaches remain experimental rather than fully validated.

Figure 1: Category-Wise Distribution of Papers

$$Xstd(ta) = \frac{1}{2w}\int_{-w}^{w} (X(ta-t) - \overline{X})^2 \, dt \qquad (1)$$

Deep learning techniques employing Long Short-Term Memory (LSTM) architectures have been applied to capture temporal relationships in landing dynamics [20], [22]. Despite the encouraging recall rates (~70%) reported in these studies, there are significant drawbacks to their experimental designs:

(i) Balanced test sets, even though only approximately 3 to 4% of flights in actual operations are HL events.

(ii) Practical application is hampered by prediction windows below the Decision Height (10 to 2 seconds prior to impact).

(iii) Datasets that are exclusive to a particular aircraft type, making it impossible to generalise across aircraft families.

These limitations call into question the usefulness of such approaches for real-time decision support.

Regression-based approaches use deep LSTM models in conjunction with feature-selection strategies to predict the normal acceleration at touchdown [24], [25]. These models routinely underestimate high-acceleration events, which leads to poor HL detection ability when converted into binary predictors, despite reporting low numerical error (MSE = 10^{-3}). Additionally, because these models only predict one step forward ($t \rightarrow t+1$), their practical relevance is further diminished by uncertainty regarding their ability to produce meaningful early predictions.

There are also several broader gaps in the literature beyond the limitations noted in prior work. Despite the substantial influence of gusts, turbulent crosswinds, and low-visibility conditions on landing outcomes, meteorological factors are rarely incorporated in existing models. Similarly, due to data availability constraints, aircraft mass and centre-of-gravity position—both of which significantly affect landing dynamics—are frequently omitted. Cross-aircraft generalisation is another notable gap; despite aerodynamic and control-law differences among models such as the A319, A320, and A321, the majority of studies evaluate a single aircraft type. Finally, operational deployability—including the need for explainable machine learning (XAI) to support crew trust and safety investigations, real-time latency constraints, and pilot–automation interaction—is seldom addressed.

Overall, the literature indicates that, while significant progress has been made in data-driven analysis of landing behaviour, no method currently in use satisfies the combined requirements of early prediction, cross-aircraft applicability, class-imbalance robustness, and cockpit deployability.

Research on landing-anomaly prediction has steadily increased, as shown by a review of publication trends from 2018 to 2025 (Figure 2). There were very few publications prior to 2020, but a noticeable increase around 2021 and 2024 suggests growing interest in data-driven approaches for approach-phase safety. This trend indicates that hard landing prediction has become an increasingly important research area in recent years.

Figure 2: Number of Retrieved Papers by Year

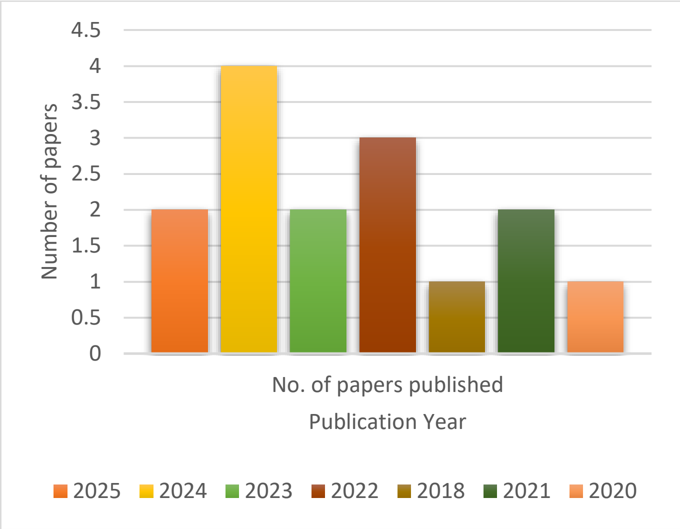

3. Methods

3.1 Dataset Description

The study is based on a large operational dataset obtained from the Flight Management System (FMS) of a European airline operating Airbus A319, A320, and A321 aircraft. The source dataset comprises 377,446 flights with 370 recorded parameters, covering several airports. To ensure consistency and comparability, the dataset was restricted to London Heathrow Airport (LHR), where aircraft follow a uniform straight-in approach corridor. This setting provides a controlled environment for training predictive models and eliminates variability introduced by differing approach geometries.

Following the removal of records with data quality issues—including signal dropouts, noise, and quantisation errors—a total of 58,177 flights remained in the usable dataset. Each flight record covers the period from the Final Approach Fix (FAF), which occurs approximately three minutes before touchdown, to 20 seconds after touchdown, in order to capture the maximum vertical acceleration (maxG).

Statistically, hard landings (HL) were defined as those exceeding 1.4037 g, calculated as the mean plus two standard deviations of the normal acceleration at touchdown. This threshold produced 2,673 HL instances (approximately 4.6%), consistent with the frequency reported by the industry. The input parameters were grouped into four categories:

1. Physical variables: Aircraft aerodynamics (e.g., airspeed, pitch, roll, accelerations, and glideslope deviation).
2. Pilot inputs: Thrust lever angles and sidestick pitch/roll commands.
3. Actuator states: Elevator, rudder, flap, slat, and aileron positions.
4. Automation factors: Flight director engagement, autopilot status, and autothrust.

3.2 Impact of Automation Factors

The correlations between maximum G and automation-related parameters (autopilot, autothrust, flight director, landing gear, and speed brake) were examined to determine whether automation systems influence hard landing events. Boxplot analysis comparing automation engagement for HL and non-HL flights showed no statistically significant differences, indicating that automation factors do not materially affect HL prediction.

3.3 Prediction Models

Two complementary machine learning approaches were considered:

(1) Regression models: Used to predict the continuous parameter maxG at touchdown.

(2) Classification models: Used to distinguish between HL and non-HL flights using a threshold of 1.4037 g.

Input features were derived from the variability of aircraft parameters at the following altitude levels:

A = {1500, 1000, 500, 400, 300, 200, 150, 100, 50, 40, 30} ft

Variability at each altitude ta was computed using the standard deviation over a sliding window of half-width w:

$$Xstd(ta) = \frac{1}{2w}\int_{-w}^{w}(X(ta-t) - \bar{X}\ )^2 dt$$

Equation (1)

where the mean $\bar{X}$ over the window is defined as:

$$\bar{X}\ = \frac{1}{2w}\int_{-w}^{w} X(ta-t)\, dt$$

Equation (2)

This approach captures temporal variability without requiring recurrent network architectures, thereby avoiding the convergence difficulties inherent in LSTM-based models.

3.4 Altitude Ranges for Training

Three altitude ranges were used for model training:

(1) AP2TD (Approach to Touchdown): Covers all sampled altitudes from the start of the approach phase to touchdown. Models trained on this range establish the system's maximum achievable accuracy.

(2) AP2DH (Approach to Decision Height): Covers altitudes {1500, 1000, 500, 400, 300, 200, 150, 100} ft. Models trained on this range determine the system's effectiveness for go-around recommendations and early HL identification.

(3) DH2TD (Decision Height to Touchdown): Covers altitudes {50, 40, 30} ft. Models trained on this range assess the upper limit of HL prediction accuracy when data below the Decision Height is used.

4. Experiments

4.1 Experimental Design

The experimental evaluation aimed to assess the predictive capability of the proposed models for hard landing (HL) detection. Both classification and regression approaches were examined. The performance of classification models was measured using sensitivity and specificity, defined as follows:

$$\text{Sensitivity} = \frac{TP}{(TP+FN)} \qquad \text{Specificity} = \frac{TN}{(TN+FP)}$$

where TP, FP, TN, and FN denote true positives, false positives, true negatives, and false negatives, respectively.

The efficacy of regression models was assessed using the mean squared error (MSE) between the actual and predicted maximum normal acceleration at touchdown:

$$MSE = \frac{1}{N}\sum_{i}(maxG_i - \widehat{maxG_i})^2$$

Equation (3)

Regression predictions were also converted to binary HL classifications at the 1.4037 g threshold for comparison with classification models.

All neural network architectures were optimised using scaled conjugate gradient backpropagation with default settings ($\sigma = 5\times10^{-1}$, $\lambda = 5\times10^{-1}$). Early stopping was applied when the gradient remained below 10^{-2} for six consecutive epochs. Each training run used the full training subset as a single batch.

A 15-fold cross-validation strategy was employed to ensure statistically robust comparisons. Class imbalance was addressed by randomly downsampling the non-HL (NHL) class in each training fold to match the size of the HL class, while the test set was kept unchanged to reflect the actual HL occurrence rate (~4 to 5%). Each test fold contained approximately 3,280 samples with around 100 HL events, while each training fold comprised 6,066 balanced samples.

The experimental analysis comprised two components:

(1) Assessment of Predictive Power: Training performance was used to select the best-performing model architectures. Sensitivity, specificity, and MSE were then evaluated on the test data for the top-performing classifier and regressor.

(2) Evaluation of Cockpit Deployability: Experiments examined performance differences across: distinct variable categories (physical, actuator, pilot, and combined); multiple altitude ranges (AP2TD, AP2DH, and DH2TD) to determine the minimum data requirements; the capacity for early HL prediction; and the feasibility of enabling go-around recommendations.

4.2 Results

4.2.1 Predictive Power of the Models

Classifier sensitivity and specificity distributions show largely consistent performance across most network configurations. Config5 and Config7, which exhibited reduced sensitivity in certain settings, were notable outliers; ANOVA confirmed that these configurations had significantly lower sensitivity across all altitude ranges.

Specificity analysis indicated that Config1, Config3, Config4, and Config6 performed considerably worse when trained on physical or combined variable sets. ANOVA confirmed these differences. Regarding regression MSE, Config6 and Config4 consistently underperformed across most settings, whereas Config2, Config3, and Config7 achieved the lowest MSE values. ANOVA results confirmed that Config6 had substantially higher MSE across all variable categories, and Config4 showed elevated MSE across altitude ranges.

The best-performing classifier (Config6) achieved average HL detection rates of approximately 85% sensitivity and 72% specificity across the AP2TD and AP2DH ranges when trained on physical or all variable categories. The best regressor (Config2) achieved the lowest MSE (approximately 2.6×10^{-3}) over the DH2TD range using physical variables. Despite this low numerical error, regression-based HL classification produced poor sensitivity (maximum approximately 46%), demonstrating that regression methods are inadequate for reliable HL detection.

4.2.2 Cockpit Deployability Potential

Models trained using physical variables achieved the highest sensitivity, with only a minor reduction in specificity compared to models trained on all variable categories. Regression and classification results were substantially worse for models trained solely on actuator states or pilot inputs. Physical variables alone did not significantly outperform the combined variable set.

Altitude-range analysis showed that AP2TD and AP2DH performance levels are statistically equivalent. However, the DH2TD range exhibits a notable decline, with sensitivity dropping by approximately 10 to

13 percentage points and specificity by approximately 7 to 9 percentage points relative to the AP2TD range. This behaviour reflects the highly dynamic nature of the touchdown flare and the effect of sudden environmental disturbances close to the runway surface.

The unsuitability of regression models for early HL diagnosis and cockpit deployment was further confirmed by the observation that they only produced acceptable MSE values when data recorded close to touchdown was included.

5. Discussion

5.1 Comparison with Existing Models

The proposed approach was compared against an LSTM model, a Support Vector Machine (SVM), and a Logistic Regression (LR) classifier reported in prior research. All three baseline methods were reimplemented from scratch using the same dataset, parameters, and evaluation metrics as the proposed method, to ensure a fair comparison.

The results demonstrate that the proposed approach outperforms the best-performing LSTM model currently in use. Specifically, the proposed method achieved 5% greater sensitivity than the LSTM in both the AP2TD and AP2DH altitude ranges. Average specificity for these ranges was 7.7% greater than that of the LSTM. With an MSE of approximately 2.6×10^{-3} for maxG prediction, the proposed regression model is comparable to state-of-the-art LSTM regression approaches. However, the regression model achieved only 46% sensitivity in identifying HL events, despite this low MSE. This discrepancy is explained by the regression loss function's declining gradient with respect to maxG, which results in systematic underestimation of maxG within the HL class. Regression predictions are therefore considered inadequate for HL detection.

The sensitivity and specificity results for different model configurations are illustrated in Figures 3 and 4. While certain configurations attain very high specificity with few false positives, only a subset maintain the sensitivity required to reliably detect hard landings. The findings highlight the importance of selecting architectures that effectively balance both metrics.

Figure 3: Accuracy Results of Different Models (Specificity)

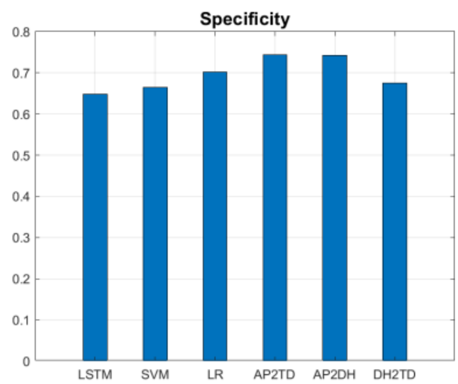

Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Figure 4: Accuracy Results of Different Models (Sensitivity)

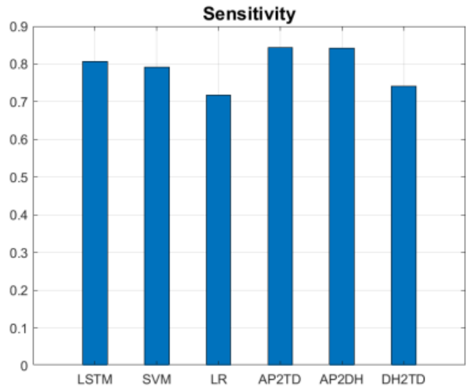

5.2 Improvements and Future Work

The current models rely on parameters recorded during the final approach phase (below 2,000 feet). However, the probability of an unstable approach may be substantially influenced by external conditions such as wind direction, turbulence, and gusts. While the variability of physical parameters captures some of these effects, icing conditions and certain low-visibility scenarios were not explicitly accounted for.

Aircraft mass (weight) data was excluded from the study due to reliability issues in the dataset, which represents a significant limitation. Aircraft weight directly affects inertia and the amount of potential energy that must be dissipated at touchdown, and is therefore a critical factor. Similarly, the centre of gravity (CG) is an important parameter for control margins and stability; incorporating both CG and mass could meaningfully improve prediction accuracy.

Future machine learning enhancements may include the use of Convolutional Neural Networks (CNNs) to capture deeper temporal relationships in the data. Further research is also required to quantify the influence of meteorological factors on landing aerodynamics. Regarding cockpit implementation, the proposed fully connected network architecture is compatible with low-resource microcontrollers, with predicted inference latency below 50 ms, ensuring that real-time operating requirements are met.

6. Conclusion

This paper has presented E-Pilots, a cockpit-deployable machine learning system for predicting hard landings (HL) during the approach phase of commercial flights. With a typical sensitivity of 85% and specificity of 74% at the Decision Height, the proposed hybrid model demonstrates strong predictive performance by leveraging a large dataset of more than 58,000 flights across multiple aircraft types (A319, A320, and A321). These results substantially exceed current state-of-the-art methods, particularly those based on LSTM architectures, which often operate with restricted altitude windows below the Decision Height and use unrealistically balanced test sets.

The approach uses variability measurements gathered at distinct altitudes to represent the temporal interdependence of aircraft state variables. This method maintains high accuracy and practical viability while avoiding the convergence difficulties of deep recurrent networks. The study also demonstrates that automation factors—such as autopilot and autothrust—do not materially affect HL outcomes, and were therefore excluded from model training.

A comparison of regression and classification models reveals that, although regression networks can predict vertical accelerations with low mean squared error (approximately 10^{-3}), their sensitivity for HL detection remains low (approximately 46%). This finding reinforces the advantage of classification-based approaches for real-time cockpit decision support. The system's capacity to identify HL precursors prior to the Decision Height enables timely go-around recommendations, which may reduce the frequency of runway excursions and improve overall flight safety. The model's generalisability across different aerodynamic profiles is validated by its evaluation on multiple aircraft types.

References

1. Boeing Commercial Airplanes, Statistical Summary of Commercial Jet Airplane Accidents – Worldwide Operations 1959–2017, Aviation Safety, Seattle, WA, USA, 2018.
2. D. Gill, "Prediction of Hard Landings through Flight Data Monitoring Using Machine Learning Techniques," in Proc. 10th Int. Conf. Aviation Technology, Integration, and Operations (ATIO), Denver, CO, USA, 2017, pp. 1–8.
3. Federal Aviation Administration, Advisory Circular AC No: 91-79A Mitigating the Risks of a Runway Overrun upon Landing, Technical Report, Federal Aviation Administration, 2016.
4. M. Coker, Why and When to Perform a Go-Around Manoeuvre, Boeing Edge, 2014, pp. 5–11.
5. T. Blajev, W. Curtis, Go-Around Decision Making and Execution Project: Final Report to Flight Safety Foundation, Flight Safety Foundation, March 2017.
6. Eurocontrol, European Action Plan for the Prevention of Runway Excursions, Report, Eurocontrol, 2013.
7. European Union Aviation Safety Agency, Artificial Intelligence Roadmap – A Human-Centric Approach to AI in Aviation, Technical Report, EASA, 2020.
8. European Union Aviation Safety Agency, The European Plan for Aviation Safety (EPAS 2020–2024), Technical Report, EASA, 2019.
9. L. Wang, C. Wu, R. Sun, Pilot Operating Characteristics Analysis of Long Landing Based on Flight QAR Data, International Conference on Engineering Psychology and Cognitive Ergonomics, Springer, 2013, pp. 157–166.
10. L. Li, R.J. Hansman, R. Palacios, R. Welsch, Anomaly Detection via a Gaussian Mixture Model for Flight Operation and Safety Monitoring, Transportation Research Part C: Emerging Technologies, 2016, 64, pp. 45–57.
11. M. Miwa, T. Tsuchiya, S. Yonezawa, N. Yokoyama, S. Suzuki, Real-Time Flight Trajectory Generation Applicable to Emergency Landing Approach, Transactions of the Japan Society for Aeronautical and Space Sciences, 2009, 52 (175), pp. 21–28.

12. G. Holmes, P. Sartor, S. Reed, P. Southern, K. Worden, E. Cross, Prediction of Landing Gear Loads Using Machine Learning Techniques, Structural Health Monitoring, 2016, 15 (5), pp. 568–582.
13. D. Zhou, X. Zhuang, H. Zuo, H. Wang, H. Yan, Deep Learning-Based Approach for Civil Aircraft Hazard Identification and Prediction, IEEE Access, 2020, 8, pp. 103665–103683.
14. F.M. Vieira, C. de Oliveira Bizarria, C.L. Nascimento, K.T. Fitzgibbon, Health Monitoring Using Support Vector Classification on an Auxiliary Power Unit, 2009 IEEE Aerospace Conference, IEEE, 2009, pp. 1–7.
15. Y.L. Luo, X. Zhang, Aircraft Electronic Board Fault Detection Based on Infrared Thermal Imaging and Integrated SVM, Meas Control Technol, 2012, 12, pp. 012–012.
16. A.O. Alkhamisi, R. Mehmood, An Ensemble Machine and Deep Learning Model for Risk Prediction in Aviation Systems, 2020 6th Conference on Data Science and Machine Learning Applications (CDMA), IEEE, 2020, pp. 54–59.
17. S. Qian, S. Zhou, W. Chang, An Improved Aircraft Hard Landing Prediction Model Based on Panel Data Clustering, Proc. 29th Chinese Control and Decision Conference (CCDC 2017), 2017, pp. 438–443.
18. L. Zheng, J. Xie, S. Qian, Risk Prediction Method of Aircraft Hard Landing Based on Flight Data, Proc. 28th International European Safety and Reliability Conference (ESREL 2018), 2018, pp. 1827–1832.
19. C. Hu, S.-H. Zhou, Y. Xie, W.-B. Chang, The Study on Hard Landing Prediction Model with Optimized Parameter SVM Method, Proc. 35th Chinese Control Conference (CCC), IEEE, 2016, pp. 4283–4287.
20. S. Hochreiter, J. Schmidhuber, Long Short-Term Memory, Neural Computation, 1997, 9 (8), pp. 1735–1780.
21. J. Schmidhuber, Deep Learning in Neural Networks: An Overview, Neural Networks, 2015, 61, pp. 85–117.
22. H. Zhang, T. Zhu, Aircraft Hard Landing Prediction Using LSTM Neural Network, Proc. 2nd International Symposium on Computer Science and Intelligent Control (ISCSIC ’18), 2018, pp. 1–5.
23. Boeing, Why and When to Perform a Go-Around Manoeuvre, Boeing Aeromagazine, QTR 02, 2014.
24. C. Tong, X. Yin, J. Li, T. Zhu, R. Lv, L. Sun, J.J.P.C. Rodrigues, An Innovative Deep Architecture for Aircraft Hard Landing Prediction Based on Time-Series Sensor Data, Applied Soft Computing, 2018, 73, pp. 344–349.
25. C. Tong, X. Yin, S. Wang, Z. Zheng, A Novel Deep Learning Method for Aircraft Landing Speed Prediction Based on Cloud-Based Sensor Data, Future Generation Computer Systems, 2018, 88, pp. 552–558.
26. International Air Transport Association et al., Unstable Approaches, Risk Mitigation Policies, Procedures and Best Practices, Final Report, 2016.
27. C. Banbury et al., Micronets: Neural Network Architectures for Deploying TinyML Applications on Commodity Microcontrollers, Machine Learning and Systems, 2021.