A Functional Analysis Approach for Prototyping a Prognostic Model

Darin Dunlap, Jonathan Lawrence, Brandon Osborn, Taylor Spoo, and Kyle Turner*

BAE Systems, Fort Worth, Texas, USA

darin.dunlap@baesystems.us
jonathan.lawrence@baesystems.us
brandon.osborn@baesystems.us
taylor.spoo@baesystems.us
kyle.turner@baesystems.us*

ABSTRACT

Large datasets in prognostics and health management (PHM) can often present challenges for degradation prediction models because of size and computational cost. In this work, a methodology is introduced that employs functional data analysis (FDA) to compress high dimensional sensor data using a B-spline basis. This method was able to reduce the dimensionality of the data while preserving the essential characteristics of the underlying trend. The reduced dimensional representation serves as input to a k-Nearest Neighbors (KNN) algorithm for remaining useful life (RUL) estimation of a jet engine nacelle fan. Experimental results demonstrate that the functional k-Nearest Neighbors (fKNN) framework can at times predict part failure within 50 flight hours, achieving sufficiently reliable RUL forecasts on the test data set. The optimal hyperparameters for this dataset as validated by the test set used 25 nearest neighbors, a prediction count of 3, and a prediction minimum threshold of 0.44. Therefore, the combined use of FDA for dimensionality reduction and KNN for RUL prediction offers a scalable solution for handling big data (~20 GB) PHM challenges while maintaining predictive performance.

1. INTRODUCTION

Man-made objects degrade over time, and if appropriate precautions are not taken, they will degrade until they can no longer provide their intended benefit (Stanton, Munir, Ikram, & El-Bakry, 2022). Aviation is not excluded, and in many cases, these failures are not only costly but also pose a safety risk. Modern aircraft are comprised of complex integrated systems which contain a level of uncertainty in their behavior due to operating in a multitude of harsh environments and under varying operating conditions (Kordestani, Orchard, Khorasani & Saif, 2023). Due to this requirement, identifying, implementing, and deploying an algorithm focused on predicting degradation and ultimate failure of a given part is not straightforward, and no one predictive model will work for all parts. Many of the classical "rule-based" health monitoring strategies miss subtle patterns which are often non-linear and develop over the lifespan of the part (Kabashkin, 2025). A process is needed to capture and mature an idea into an accurate model that provides the end user the benefit of concepts like Just-In-Time supply chain, which seek to have materials ready as needed, eliminating the need for extra inventory holdings. These benefits range from reduced operating costs, lowered inventory, shortened lead time, and increased responsiveness (Yang et al., 2021). From the algorithm development standpoint, the process should be geared to highlight the effectiveness of a proposed methodology by considering the algorithm's performance around metrics like missed detections and false alarms. For this project, the precision metric was identified as the most critical factor. This metric was chosen based on customer feedback and a cost-benefit analysis for this project, which prioritized making sure the predictions were true positives with misdetections becoming less consequential.

This current paper provides a snapshot of the work required to take an idea for a prognostic algorithm to a working prototype that can produce a warning that a part will likely fail soon. To that end, the following section highlights literature concerning some of the data science techniques applied in the cleaning and analyzing of the signal data to produce the prototype algorithm to monitor the life of the nacelle fan. Then followed by the methodology and data

*Contact this author for questions

cleaning approach that was taken for building the model, followed by the data analysis section where the model specifications are provided. Afterward, a summary of results is given, and lastly, a discussion about the future work and limitations of the model is presented.

2. Literature Review

The following sections provide a summary of literature over PHM regarding the F-35 Joint Strike Fighter and an overview of common data cleaning and data analytic techniques.

2.1. F-35 PHM

Prognostics and health management (PHM) has been studied and used in aerospace for the last several decades for the F-35 specifically (Brown, McCollom, Moore, & Hess. 2006). In their report of using data from signals on the jet, they highlighted a process to incrementally develop prognostic algorithms to better streamline part acquisition and maintenance planning to reduce cost and downtime for the fleet. The primary methodology for this development of a prognostic algorithm for a particular subsystem of the jet was to extract data from health report codes (HRCs), the raw signal data, and event data which could include fluid flow rates, air vehicle motion data, and other relevant data. Brown et al. (2006) then explains the high-level understanding of how this data is used to schedule maintenance and how PHM would be an invaluable tool for the F-35 fleet. Brown et al. (2006) provides a large-scale picture of the PHM process, which provides the context for the development of prognostic algorithms to detect failures for this paper.

2.2. Dimensionality Reduction Techniques

Both principal component analysis (PCA) and linear discriminant analysis (LDA) are commonly used techniques to reduce the dimensions of the data. PCA (as discussed in Bro & Smilde, 2014) is the linear combination of features where each feature is provided a weight to show its significance in a model in an unsupervised approach. LDA (as discussed in Tharwat, Gaber, Ibrahim, & Hassanien, 2017 and Qu & Pei, 2024) is a supervised approach to reduce dimensions which uses an orthogonal transformation to create a linear combination of features. Each of these methods primarily handles discrete data.

Functional data analysis (FDA) has been used in PHM frequently as a robust method for dimensional reduction by creating the coefficients of a polynomial function to represent the data. This has been used in aerospace (e.g. Zhang, Zheng, Huang, Feng, Zhou, & Zhang, 2020) and has general applications as expressed in Wang, Chiou, and Müller's (2016) thorough explanation of FDA. Wang, Zheng, Farahat, Serita, & Gupta (2019) present a methodology for estimating the remaining useful life (RUL) of any piece of equipment, assuming that sensor data are available over a suitable time interval.

2.3. Data Analysis Techniques

A frequently used technique to analyze data for prognostics is the k-Nearest Neighbors (KNN) algorithm, which is a supervised learning algorithm used in classification and regression problems. A detailed description of the uses and benefits of KNN is given in Taunk, De, Verma, & Swetapadma (2019), where they describe several different variants of KNN as well as advantages, such as its simple application and easily interpretable result. See additionally Yao and Ruzzo, 2006 for an example of a regression-based KNN for applications in biology. This article noted that although simple, "KNN methods are among the best performers in a large number of classification problems" (Yao & Ruzzo, 2006). This is largely due to KNN being highly flexible as it does not make assumptions about the data. This allows KNN the ability to optimize hyperparameters and distance metrics.

One variant of KNN is functional KNN (fKNN) which uses the same classification or regression approach as the standard KNN approach but uses function representation of the data as inputs. This method is explained by Thomas Laloë (2008) and Tianming Zhu & Jin-Ting Zhang (2010) to show how KNN can be adapted to work with functional data. Functional KNN is applicable to either the classifier or regression problem for KNN.

3. Methodology

For this project, the team used data from two data sources: F-35 signal database and a system for facilitating the collection of maintenance data. The F-35 signal database is a proprietary system that is used to collect, analyze, and report data on the F-35 Lightning II fleet's performance and operational activities. Our features used in our project come from this database. The maintenance database is a proprietary system that records F-35 maintenance events such as cause of maintenance and codes to characterize the repairs. This information is used to inform and label our signal data as health and unhealthy. The following sections will describe how data is routinely cleaned, and the data's dimensionality is reduced for our analysis.

3.1. Routine Data Cleaning

For this work, the team used the signals related to calibrated airspeed, the fan speed, fan voltage, and fan current. These signals were selected based on feedback from a subject matter expert familiar with this system, as they were identified as the primary controls governing the part's functionality and determining its failure. Data was considered good for our investigation if it met the following conditions:

- The flights had to occur in 2020 or later.
- The jet needed to have flown at least once during the power cycle to be considered in our sample. This was done to exclude ground runs where signals may not have been present or useful. A power cycle in this context is the time series data from power on of the aircraft to power off.
- The historical jet data had to have the targeted HRC asserted.

In addition, the team performed routine cleaning such as removing null values and any signals recorded outside their reported operating range as well as converting data types to match similar signals between tables for table merges. All these criteria were chosen to ensure high accuracy and quality of the data. Once the data was cleaned, all signals were normalized to be between zero and one. Then the signals were synchronized to a common normalized system time, enabling better matching of recorded signal occurrences.

After data collection and cleaning, the sample consisted of approximately 110 jets, each with up to four years of data, spanning the entire lifetime of the fan until the first recorded failure for that jet. Given the extensive range of this lifespan value, spanning more than 800 flight hours till failure, a bucketing strategy was employed to categorize the data into manageable intervals of 50-hour increments, initially ranging from 0-50 hours to 200 hours to failure. Upon examination of the data, the signal behavior was found to remain consistent across the broader range of 200-1000 flight hours till failure. As a result, the categorization was simplified by labeling all values above 200 hours as a single category, effectively truncating the upper bound at 200 hours.

3.2. Data Dimensionality Reduction

The final step in our data preparation was to reduce the scale and dimensionality of our data. Functional data analysis (FDA) was selected for the purpose of dimensionality reduction due to the large quantity of data that needed to be processed. FDA effectively condensed the data while retaining as much information as possible, allowing us to manage noise across power cycles and treat the observations as a continuous function rather than discrete points. This approach enabled the interpolation of data recorded at different time intervals and the alignment of related signals on the same time scale. Additionally, the scikit-fda library leverages the scikit-learn library's framework, enabling seamless integration with existing packages through a familiar API.

Principal component analysis (PCA) and linear discriminant analysis (LDA) were considered as alternative dimensionality reduction methods; however, these resulted in excessive information loss, leading to the correlation behavior with failure being lost in the reduced data, producing inaccuracies. For example, while LDA can typically handle noisy data, it did not perform well in this case, as too much information would have been lost. Unlike FDA, PCA and LDA reduce the number of features, not the number of records. In contrast, FDA preserves the number of features while allowing control over the number of manageable records for our analysis via interpolation, making it the more suitable method for this investigation.

3.3. FDA Basis Representation and Interpolation

FDA allows for a more interpretable representation of the data due to its use of basis functions. To produce the basis function representation the power cycle data for each aircraft was first stored in a Functional Data Grid (FDataGrid) object. FDataGrid is a class from the *scikit-fda* library designed for functional data analysis (FDA) that represents data as a series of curves. The datagrid was then converted to be represented using a B-spline basis. B-spline was selected for this work because it is flexible for non-periodic data by allowing for the generation of a smooth function from complex data such as the signals being received from the jet.

After selecting B-splines as the basis functions, the optimal number of functions to use in the investigation needed to be determined. Using too few basis functions would capture insufficient information, whereas using too many would yield diminishing returns, increase the risk of over-fitting, and raise computational cost. The optimal number of basis functions was found by splitting the dataset on aircraft into training and validation sets on a 70/30 split as well as using stratify to ensure equal proportions of class labels. The functional representation was fitted with varying numbers of basis functions. The number that produced the lowest reconstruction error in cross-validation was identified as the optimal choice. Based on this criterion, 50 basis functions were used.

Once the number of basis functions was determined, it was necessary to ensure that the four signals were comparable despite their differing sampling rates. Varying reporting rates introduce complications for statistical analyses, which FDA is well equipped to handle as it can fill in any missing data as the values of the function. Gaps or heterogeneous noise levels can make direct comparisons unreliable. To address this, each functional curve was interpolated into 200 evenly spaced points on the normalized interval ([0,1]), producing a unified grid for all four signals. Power cycles that could not be converted to a functional representation (e.g., due to missing or invalid data) were excluded from the functional datagrid. When necessary, each signal array was amended to include endpoints at 0 and 1, ensuring that every signal shared the same domain.

The functional dataset for each jet is stored in a single FDataGrid object with the highlighted hierarchical structure below. For a given jet there is a collection of power-cycles

$$Jet = [\boldsymbol{PC}_1, \boldsymbol{PC}_2, \ldots, \boldsymbol{PC}_{\boldsymbol{m}}] \tag{1}$$

where m is the number of power cycles from start of operation to the maintenance event. Each power cycle ($\boldsymbol{PC}_i$) is an array of n observation vectors taken from the functional representation.

$$\boldsymbol{PC}_i = [\boldsymbol{O}_{i1}, \boldsymbol{O}_{i2}, \ldots, \boldsymbol{O}_{in}] \tag{2}$$

variables (e.g., the functional data) and the output variable. This makes KNN regression well-suited for situations where the underlying relationship between the inputs and outputs is unknown or complex. The number of neighbors chosen was determined by training the model on varying increments of neighbors, ranging from 25 to 200 with increments of 25, to get an output that was a prediction of the test set alongside descriptive statistics of the model output for the predicted values at different thresholds of flight hours till failure.

KNN regression was used to predict which bucket each

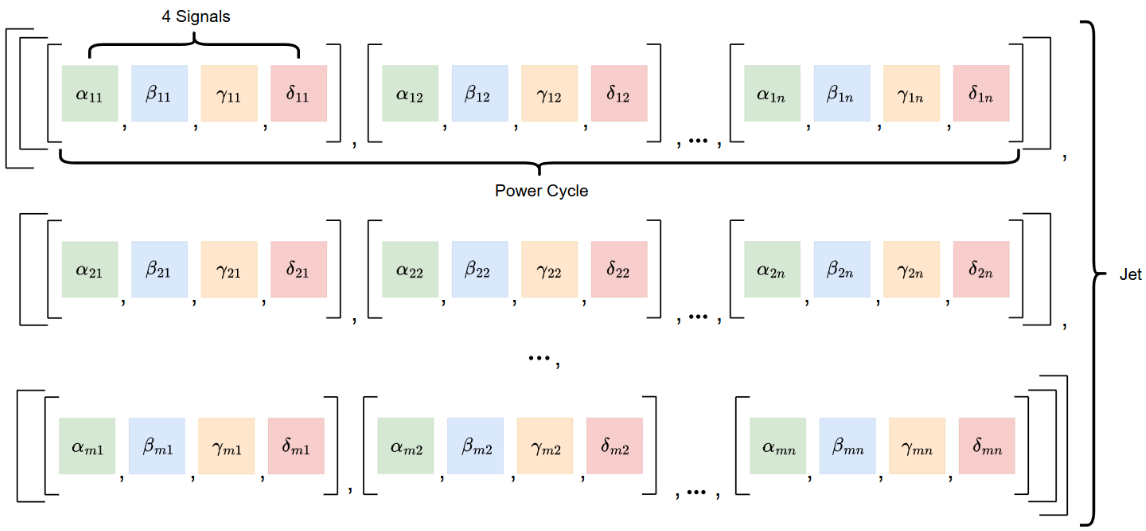

Figure 1: The input data for the KNN model consisted of # of jets, where each jet consisted of m number of power cycles and each power cycle consisted of n number of data points for each of the four signals.

Equation 2 above represents the i^{th} power cycle where i ranges from 1 to m and n is the number of selected grid points; in this work n = 200. Every observation vector ($\boldsymbol{O}_{ij}$) contains the four signals evaluated at the same grid point, where j ranges from 1 to n, α is fan airspeed, β is the fan speed, γ is the fan voltage, and δ is the fan current.

$$\boldsymbol{O}_{ij} = [\alpha_{ij}, \beta_{ij}, \gamma_{ij}, \delta_{ij}] \tag{3}$$

Thus, the FDataGrid object can be viewed as an array-of-arrays, see Figure 1 for visualization. This standardization guarantees a common range for all signals and facilitates the subsequent fKNN model described in the following section.

4. MODELING: FUNCTIONAL KNN REGRESSOR

The functional datagrid was divided into training and test sets as mentioned in Section 3.3, where the former was used to train a functional k-Nearest Neighbors (fKNN) Regressor with 25 neighbors. The choice of KNN regression was motivated by its non-parametric nature, which means that it does not impose any assumptions about the underlying functional form of the relationship between the input

power cycle belonged to. Different minimum thresholds were evaluated to determine when a prediction should be considered a 1 or a 0, where a 1 is the given label for an impending failure and 0 is the given label for a healthy part. Effectively these thresholds produced the bound at which if the model produced an output below this threshold it was treated as a 0 and when the model produced an output at or above the threshold it was treated as a 1. This threshold was used in combination with the count of predictions to determine how many predictions it took to correctly identify when a power cycle is within 50 actual flight hours till failure. The thresholds were selected from the range of values that were produced for the range of predicted model output.

After training the model, the predictions were evaluated for how many failure predictions were required before a true positive was detected. Looping through a list of values up to 30, the model would output precision metrics and predicted values for each KNN model to identify the optimal parameters for the number of nearest neighbors, the threshold, and the number of predictions required. In the following section, selected results for hyperparameters and the interpretation of those results are provided.

5. Results & Discussion

With the clean, functional representation of the data and tuned KNN model, a report of the metrics for the precision of this model is provided. The reported metrics of precision are based on the number of true detections divided by the total number of predicted detections. Reported metrics show how the precision, or positive predictive value, was influenced by the hyperparameters of the number of predictions required before producing a remaining useful life (RUL). This informed what minimum threshold value was needed to classify the power cycle as a prediction. For these results the number of nearest neighbors was kept at 25, which was found to produce optimal results, although, as mentioned before, the model was tuned using values ranging from 25 to 200 in increments of 25 for the nearest neighbors.
Table 1 shows the results for when the number of predictions was 1, 3, or 5, and for each of those the minimum prediction threshold of 0.24, 0.44, and 0.74 are given. These results were given from the test data of approximately 33 aircraft. As a note, the last row is given for completeness, although the precision is not applicable since there were no predictions made.

Table 1: Results from hyperparameter tuning for fKNN model

Number of Positive Predictions	Threshold	*n*th Prediction Required	Precision
15	0.24	1	6%
15	0.24	3	6%
15	0.24	5	6%
11	0.44	1	27%
7	0.44	3	71%
5	0.44	5	40%
2	0.74	1	50%
1	0.74	3	100%
0	0.74	5	N/A

Given a small threshold (0.24) the precision of the predictions is consistently low regardless of the number of predictions required before an RUL is generated. On the other extreme, given a high threshold (0.74), while the precision is better, the number of jets that have a predicted failure are low, thus reducing the impact of this model. From the values provided the minimum threshold value of 0.44 with 3 predictions required for an RUL value is the optimal case for this dataset which predicts a failure within 50 flight hours for 7 jets, and from the historic data it was determined that 71% of those jets had a failure within 50 flight hours after that prediction would have been made. This suggested that the model is able to identify when the fan is within 50 flight hours till failure by the time of the third prediction of 0.44 or greater. While this model may have missed some of the failures, the goal of this model was to provide advanced warning even on a few jets to provide sufficient cost savings by allowing time to order the nacelle fan and schedule maintenance.

6. Conclusion

This paper presents the process of producing a prognostic prototype for an algorithm to monitor the life of the nacelle vent fan on a jet. It has been shown that this method can produce an RUL by utilizing FDA and KNN with some precision, for this dataset and parameters this maxed out at 71% but given more true cases of failure this could be improved. The work presented here is among the first few stages of a larger scale procedure to create an algorithm to track part degradation. From this work, key observations were made to identify how data science and machine learning techniques can be combined to model complex real-world scenarios in a prototype algorithm. It is the intention of the team to investigate other fans to see if the same process would be applicable to similar parts in different locations of the jet.

This project assumed that this part did not transfer among jets during its lifespan. Future work should be done to mitigate this limitation to track by part rather than by jet. Analysis for this project is still ongoing as the prototype algorithm continues to run on historic and current datasets. Improvements to the algorithm will be made to better match the estimated RUL to the actual life remaining of the part based on recorded maintenance events. These adjustments will primarily be focused on tracking the specific part as it moves between jets as well as determining how many failure predictions are required before a true positive was detected. This method can be beneficial for use with highly complex data, with multiple signals, and difficult to detect true failures using FDA and KNN.

Acknowledgement

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E.A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. (2020) SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17(3), 261-272. DOI: 10.1038/s41592-019-0686-2.

Harris, C.R., Millman, K.J., van der Walt, S.J. et al. *Array programming with NumPy*. Nature 585, 357–362 (2020). DOI: 10.1038/s41586-020-2649-2.

Van Rossum, G. (2020). *The Python Library Reference, release 3.8.2*. Python Software Foundation.

Scikit-learn: Machine Learning in Python, Pedregosa et al., JMLR 12, pp. 2825-2830, 2011.

References

Bro, R., & Smilde, A. K. "Principal component analysis." *Analytical Methods*, vol. 6, no. 9, 2014, pp. 2812–2831, https://doi.org/10.1039/c3ay41907j.

Brown, E. R., McCollom, N. N., Moore, E., & Hess, A. "Prognostics and health management a data-driven approach to supporting the F-35 Lightning II." *2007 IEEE Aerospace Conference*, 2007, https://doi.org/10.1109/aero.2007.352833.

Kabashkin, I. (2025). AI and Evolutionary Computation for Intelligent Aviation Health Monitoring. *Electronics*, *14*(7), 1369. https://doi.org/10.3390/electronics14071369.

Kordestani, M, Orchard, M. E., Khorasani, K. and Saif, M, "An Overview of the State of the Art in Aircraft Prognostic and Health Management Strategies," in *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1-15, 2023, Art no. 3505215, doi: 10.1109/TIM.2023.3236342.

Laloë, T. "A k-Nearest Neighbor Approach for Functional Regression." *Statistics & Probability Letters*, vol. 78, no. 10, Aug. 2008, pp. 1189–93, https://doi.org/10.1016/j.spl.2007.11.014. Accessed 20 Mar. 2023.

Qu, L. & Pei, Y. "A comprehensive review on discriminant analysis for addressing challenges of class-level limitations, small sample size, and robustness." *Processes*, vol. 12, no. 7, 2 July 2024, p. 1382, https://doi.org/10.3390/pr12071382

Stanton, I., Munir, K., Ikram, A., & El-Bakry, M. "Predictive maintenance analytics and implementation for aircraft: Challenges and opportunities." *Systems Engineering*, vol. 26, issue 2, 2023, pp. 216-237, https://doi.org/10.1002/sys.21651

Taunk, K., De, S., Verma, S., & Swetapadma, A. "A Brief Review of Nearest Neighbor Algorithm for Learning and Classification." *Proceedings of the International Conference on Intelligent Computing and Control Systems*, IEEE, 2019.

Tharwat, A., Gaber, T., Ibrahim, A., & Hassanien, A. E. "Linear discriminant analysis: A detailed tutorial." *AI Communications*, vol. 30 no. 2, 16 May 2017, pp. 169-190, https://doi.org/10.3233/aic-170729

Wang, J., Chiou, J., & Müller, H., "Functional Data Analysis." *Annual Review of Statistics and Its Application*, vol. 3, no. 1, June 2016, pp. 257–95, https://doi.org/10.1146/annurev-statistics-041715-033624.

Wang, Q., Zheng, S., Farahat, A., Serita, S., & Gupta, C. "Remaining Useful Life Estimation Using Functional Data Analysis." *ArXiv (Cornell University)*, Cornell University, June 2019, https://doi.org/10.1109/icphm.2019.8819420. Accessed 9 Jan. 2024.

Yang, J., Xie, H., Yu, G., & Liu, M. (2021). Achieving a Just–in–time Supply chain: the Role of Supply Chain Intelligence. *International Journal of Production Economics*, *231*(107878), 107878. https://doi.org/10.1016/j.ijpe.2020.107878

Yao, Z. & Ruzzo, W. L. "A Regression-based K nearest neighbor algorithm for gene function prediction from heterogeneous data." *BMC Bioinformatics* 7 (Suppl 1), S11 (2006). https://doi.org/10.1186/1471-2105-7-S1-S11

Zhang, B., Zheng, K., Huang, Q., Feng, S., Zhou, S., & Zhang, Y. "Aircraft Engine Prognostics Based on Informative Sensor Selection and Adaptive Degradation Modeling with Functional Principal Component Analysis." *Sensors*, vol. 20, no. 3, 9 Feb. 2020, p. 920, https://doi.org/10.3390/s20030920. Accessed 5 Feb. 2022.

Zhu, T. & Zhang, J. "A New k-Nearest Neighbors Classifier for Functional Data." *Statistics and Its Interface*, vol. 15, no. 2, Lehigh University, Jan. 2022, pp. 247–60, https://doi.org/10.4310/20-sii650. Accessed 29 July 2025.