


PROACTIVE MONITORING OF EQUIPMENT CONDITION IN MULTI-AGENT SYSTEMS USING LSTM NEURAL NETWORKS

MHD Sami Shekh Al Ard; Elena A. Kovalchuk & Vyacheslav V. Potekhin
Peter the Great St. Petersburg Polytechnic University



This Publication has to be referred as: Shekh Al Ard, M[HD Sami]; Kovalchuk, E[lena Alexandrovna] & Potekhin, V[yacheslav Vitalyevich] (2026). Proactive monitoring of equipment condition in multi-agent systems using LSTM neural networks, Proceedings of the 37th DAAAM International Symposium, pp.xxxx-xxxx, B. Katalinic (Ed.), Published by DAAAM International, ISBN 978-3-902734-xx-x, ISSN 1726-9679, Vienna, Austria
 DOI: 10.2507/37th.daaam.proceedings.xxx

Abstract

This paper presents a method for proactive monitoring of equipment condition in multi-agent systems based on Long Short-Term Memory (LSTM) recurrent neural networks. The proposed approach enables real-time prediction of the State of Health (SoH) and Remaining Useful Life (RUL) of critical equipment components, facilitating autonomous decision-making about reconfiguration of agents in the system before failure occurs. The method is validated using the NASA Battery Dataset as a representative benchmark for degradation processes in automated systems. A multi-output LSTM model trained on charge-discharge cycle data achieves SoH prediction RMSE of 1.67% with $R^2 = 0.94$ and RUL mean absolute error of 36 cycles — outperforming GRU, 1D-CNN, and vanilla RNN baselines. The proactive approach is demonstrated through a case study of an orbital small spacecraft cluster, where proactive reconfiguration triggered by LSTM predictions reduces characteristic velocity (ΔV) expenditure by 18% (23.0 vs. 28.0 m/s) and restores formation positioning RMSE to 1.65 km. The modular architecture with a standardised data interface ensures applicability across diverse automated system domains.

Keywords: LSTM; multi-agent systems; predictive maintenance; remaining useful life; state of health

1. Introduction

Modern automated manufacturing and process control systems incorporate a growing number of critical electromechanical and electronic components whose condition directly determines system reliability and operational efficiency. Unplanned equipment failures cause production losses, safety incidents, and costly emergency maintenance interventions. Predictive maintenance — the practice of monitoring equipment health in real time and predicting failures before they occur — has therefore become a key research and industrial priority [1, 2].

Traditional condition monitoring relies on threshold-based or model-driven diagnostics that detect failures only after degradation becomes measurable. More advanced approaches use machine learning to predict the Remaining Useful Life (RUL) and State of Health (SoH) of equipment from time-series sensor data. Among these, Long Short-Term Memory

(LSTM) recurrent neural networks have emerged as the leading method for temporal degradation modelling due to their ability to capture long-range dependencies in sequential data [3, 4].

However, most existing work treats condition monitoring and system control as separate problems. This paper proposes an integrated approach in which LSTM-based equipment health predictions are directly fed into an autonomous control decision layer. When the predicted RUL falls below a critical threshold, the control system proactively initiates a reconfiguration or compensatory action — before the equipment actually fails. This proactive strategy is contrasted with the conventional reactive approach, where control action is taken only after failure has occurred.

The proposed method is validated on the NASA Battery Dataset [5] and demonstrated in a case study of autonomous formation control of a small spacecraft cluster, where battery failure prediction drives proactive cluster reconfiguration. The paper answers the five DAAAM research questions: (1) What is the problem? — reactive cluster control causes high fuel costs after failure; (2) What have others done? — battery prognostics and formation control studied separately; (3) What have we done? — developed and validated an integrated LSTM-prognostics and potential-field-control system; (4) What is new? — 18% ΔV saving through proactive LSTM-triggered reconfiguration; (5) What is next? — multi-failure scenarios and hardware-in-the-loop validation.

2. Related Work

Predictive maintenance and equipment condition monitoring have been studied extensively using both model-based and data-driven approaches. Early model-based methods relied on physics-of-failure models and extended Kalman filters, which require detailed system knowledge and are difficult to generalise across equipment types. Data-driven methods, particularly machine learning, have largely superseded these in practice due to their flexibility and accuracy [1].

Among data-driven approaches, LSTM networks have consistently achieved state-of-the-art results on RUL prediction benchmarks. Hochreiter & Schmidhuber [3] originally introduced the LSTM architecture to address the vanishing gradient problem in recurrent networks. Li et al. [4] demonstrated multi-scale LSTM feature extraction for bearing RUL estimation, achieving significant improvements over traditional methods. The NASA Battery Dataset [5], introduced by Saha & Goebel, has become the standard benchmark for battery SoH and RUL prediction algorithms, enabling direct comparison across methods.

Integration of health monitoring with autonomous control has been studied in manufacturing contexts. Hrnjica & Softic [6] demonstrated explainable AI for predictive maintenance in industrial settings, providing a precedent for integrating neural network prognostics with autonomous control algorithms. In the domain of multi-agent systems and formation control, potential field methods [7] provide a well-established framework for autonomous coordination of multiple agents. Multi-criteria optimisation of manoeuvre planning under constraints was addressed in [8]. The present work is novel in directly coupling LSTM health predictions to a potential field formation controller, enabling proactive autonomous reconfiguration demonstrated quantitatively with real NASA degradation data [11].

3. LSTM-Based Condition Monitoring Module

3.1 Data Preparation and Health Indicators

The NASA Battery Dataset [5] provides experimental time-series data for lithium-ion battery charge-discharge cycles under controlled conditions (24°C, 2 A constant-current discharge, 2.7 V cutoff, CC-CV charge to 4.2 V). This dataset is representative of the degradation dynamics encountered in automated systems with onboard power subsystems. Four batteries were selected: B0005, B0006, B0007, and B0018. Each battery was mapped to one spacecraft in the simulated cluster, forming a four-spacecraft scenario that preserves real differences in degradation rates without artificial data duplication. The complete dataset characteristics are summarised in Table 1.

Battery	Initial capacity, Ah	Final capacity, Ah	Degradation, %	Cycles	Mapped to
B0005	1.86	1.33	−33.7	148	Satellite 1
B0006	2.04	1.19	−40.7	148	Satellite 2
B0007	1.89	1.43	−28.4	148	Satellite 3
B0018	1.86	1.34	−32.9	112	Satellite 4

Table 1. NASA Battery Dataset characteristics and spacecraft mapping

Two primary health indicators are defined. State of Health (SoH) quantifies the relative capacity loss:

$$SoH(t) = C(t) / C_0 \times 100\% \quad (1)$$

where $C(t)$ is the measured discharge capacity at cycle t and C_0 is the initial nominal capacity. End-of-life (EoL) is defined at $SoH = 70\%$. Remaining Useful Life (RUL) is defined as the number of cycles remaining before EoL, converted to operational hours:

$$RUL_{hours} = RUL_{cycles} \times T_{cycle} \quad (2)$$

where T_{cycle} is the average cycle duration in hours. A third indicator, failure probability P_{fail} , is defined as the predicted probability of EoL occurring within the next $N_h = 100$ hours, enabling probabilistic threshold-based decision-making. Battery B0006 shows the highest degradation (-40.7%), physically justifying its selection as the critical element in the reconfiguration scenario.

3.2 Multi-Output LSTM Architecture and Training

The proposed LSTM model simultaneously predicts three health indicators: $SoH_{percent}$, RUL_{hours} , and P_{fail} . The architecture consists of two stacked LSTM layers with 64 units each, followed by a fully connected dense output layer with three neurons and linear activation. A sliding window of the last 10 consecutive measured capacity values serves as the input sequence, providing sufficient temporal context while maintaining computational efficiency.

Training was performed on batteries B0005, B0006, and B0007 (combined: 444 data points); evaluation was conducted on the held-out test battery B0018 (112 points). The Adam optimiser was used with initial learning rate 0.001, batch size 16, and early stopping with patience of 20 epochs monitoring validation loss. Dropout regularisation (rate 0.2) was applied after each LSTM layer to prevent overfitting. The multi-output architecture ensures that SoH , RUL , and P_{fail} share the same temporal feature representation, improving generalisation compared to three independent single-output models.

Training converged in 87 epochs with a final training loss of 0.0031 and validation loss of 0.0038, indicating no significant overfitting. Evaluation results on B0018: SoH RMSE = 1.67%, $R^2 = 0.94$; RUL MAE = 36 cycles (approximately 36 days). Results are transmitted via a standardised CSV interface: timestamp | unit_id | RUL_{hours} | $SoH_{percent}$ | failure_probability, providing a modular decoupled interface between the prognostic and control modules.

3.3 Comparative Evaluation of Neural Network Architectures

To justify the choice of LSTM, the proposed model was benchmarked against three alternative neural network architectures on the same train/test split with identical hyperparameters: Gated Recurrent Unit (GRU) — a simplified gated recurrent architecture with fewer parameters; one-dimensional Convolutional Neural Network (1D-CNN) — a feed-forward architecture extracting local temporal features; and vanilla Recurrent Neural Network (RNN) — a simple recurrent architecture without gating. Results are summarised in Table 2.

Method	SoH RMSE, %	RUL MAE, cycles	Train time, s	Parameters
LSTM (proposed)	1.67	36	42	~33k
GRU	1.89	41	38	~25k
1D-CNN	2.34	53	21	~18k
Vanilla RNN	3.12	68	35	~8k

Table 2. Comparison of neural network architectures for battery health prediction

LSTM achieves the best SoH RMSE (1.67%) and RUL MAE (36 cycles), outperforming GRU by 12%, 1D-CNN by 40%, and vanilla RNN by 87% on SoH accuracy. The improvement over GRU is modest but consistent, reflecting the benefit of the full LSTM gating mechanism for long-range degradation trajectory modelling where degradation patterns span tens of cycles. The vanilla RNN's poor performance confirms the importance of gating for this task. The training time overhead of LSTM versus GRU (42 vs 38 seconds) is negligible for offline training. The results confirm LSTM as the preferred architecture.

3.4 Sensitivity of Decision Thresholds

The proactive control trigger is activated when $RUL_hours < N_h$ or $failure_probability > \theta$. The choice of N_h and θ directly affects the trade-off between early warning lead time and false alarm rate. To analyse this sensitivity, the control trigger was evaluated for $N_h \in \{50, 100, 150\}$ hours and $\theta \in \{0.5, 0.7, 0.9\}$. Increasing N_h from 50 to 150 hours increases advance warning from 2 to 6 days but raises false alarm probability from 3% to 11%. The threshold $\theta = 0.7$ with $N_h = 100$ hours was selected as the operating point offering zero missed detections on the test set, 4% false alarm rate, and sufficient lead time for autonomous reconfiguration.

4. Proactive Control Algorithm for Multi-Agent Systems

4.1 Agent Dynamics and Potential Field Control

The autonomous control layer receives health predictions from the LSTM module via the standardised CSV interface and uses them to drive proactive formation reconfiguration. Position errors of agents with respect to their target positions are determined continuously [9]:

$$e_i = r_i - r_i^{ref} \quad (3)$$

where r_i is the current position and r_i^{ref} is the target position of unit i . The control force is computed by the potential field method [9]:

$$F = F_{att} + F_{rep} + F^d \quad (4)$$

where F_{att} is the attractive component reducing position error; F_{rep} is the repulsive component maintaining safe separation; and F^d is the damping component controlling approach velocity. The attractive component is implemented as a linear spring force proportional to position error, while the repulsive component activates only when inter-unit distance falls below a safe threshold d_{safe} . This combination ensures smooth convergence to target positions without oscillation. Individual signal magnitudes in maintenance mode are computed by normalisation [10]:

$$\Delta V_i = (e_i / e_{max}) \cdot \Delta V_{as}^e \quad (5)$$

A multi-criteria objective function simultaneously minimises position error, relative velocity, and fuel consumption:

$$J = \omega_p \sum_i \|r_i - r_i^{ref}\|^2 + \omega_v \sum_i \|v_i\|^2 + \omega_f \sum_i \Delta V_i \quad (6)$$

where ω_p , ω_v , ω_f are weighting factors set to 0.5, 0.3, 0.2 respectively, reflecting the priority of positioning accuracy over velocity and fuel. This formulation ensures the algorithm seeks a compromise between formation holding accuracy and fuel conservation, rather than pursuing the ideal position as quickly as possible. When reconfiguration is initiated, target assignment is optimised by minimising total displacement:

$$J = \sum_i d_i \quad (7)$$

The final signal for each unit during reconfiguration is determined by an integral performance indicator balancing position error and velocity:

$$J_i = \omega^d (e_i / e_{max}) + \omega_v (v_i / v_{max}) \quad (8)$$

$$\Delta V_i = J_i \cdot \Delta V_{as}^e, \quad \Delta V_{min} \leq \Delta V_i \leq \Delta V_{max} \quad (9)$$

The algorithm implements two-phase synchronous control: all signals are computed first, then executed simultaneously, eliminating order-dependent artefacts and ensuring consistent system-level behaviour.

4.2 Proactive Reconfiguration Trigger

The control system polls the LSTM module every 3600 seconds. A unit is flagged critical when:

$$RUL_hours < 100 \text{ OR } failure_probability > 0.7 \quad (10)$$

Upon detection: (1) the critical unit is excluded from the target formation; (2) a new target configuration is computed as a regular (N-1)-gon preserving the formation baseline L; (3) each remaining unit is assigned an updated target and drives toward it autonomously. All decisions are made onboard without ground station involvement, eliminating the communication latency that would delay a reactive strategy by hours or days.

5. Case Study: Orbital Cluster Formation Control

To demonstrate the integrated system, a 148-day simulation of a four-unit orbital cluster of small spacecraft (SSC) was conducted. Each spacecraft is equipped with a lithium-ion battery from the NASA dataset (B0005, B0006, B0007, B0018 respectively), and the LSTM module monitors all four subsystems in real time with a query interval of 3600 seconds. The cluster maintains a square formation with baseline $L = 500$ km for distributed aperture radar Earth observation [11, 12].

Spacecraft No. 2 (battery B0006, maximum degradation: -40.7%) was flagged critical at day 79. The LSTM module predicted EoL 100 hours in advance with $P_{fail} > 0.7$, triggering autonomous proactive reconfiguration. Fig. 1 shows four key moments of the simulation: the initial square formation ($t = 0$), the proactive reconfiguration trigger ($t = 79$ days), actual failure of spacecraft No. 2 ($t = 84$ days), and the stabilised equilateral triangle formation at end of simulation ($t = 148$ days). The remaining three spacecraft (No. 1, 3, 4) converge autonomously to the new triangle configuration, preserving the virtual formation centre and baseline L.



Fig. 1. Cluster configuration states: square formation $\rightarrow$ proactive reconfiguration $\rightarrow$ triangle formation

Fig. 2 compares cumulative ΔV expenditure between proactive and reactive strategies over the full 148-day simulation. The proactive approach consumed 23.0 m/s versus 28.0 m/s for the reactive baseline — a saving of 5.0 m/s ($\sim 18\%$). This advantage arises because proactive reconfiguration is initiated while formation deviations are still small (spacecraft No. 2 is still partially functional), requiring lower corrective impulses than emergency post-failure correction when spacecraft No. 2 has already drifted from its nominal position.

Fig. 3 shows the formation positioning RMSE over time. The peak of 450.43 km at day 80 reflects the intentional reconfiguration manoeuvre — the deliberate relocation of three spacecraft from square vertices to triangle vertices. By day 85, RMSE drops to 15.65 km as the potential field controller drives each unit toward its new target. At simulation end (day 148), RMSE stabilises at 1.65 km — identical to the accuracy achieved in the original square formation, confirming complete geometric restoration. Key quantitative results are summarised in Table 3.



Fig. 2. Cumulative ΔV expenditure: proactive vs. reactive control strategies

Fig. 3. Formation positioning RMSE: reconfiguration transient and convergence

Parameter	Proactive strategy	Reactive strategy	Improvement
Reconfiguration trigger, days	79	84 (after failure)	5 days earlier
Total ΔV expenditure, m/s	23.0	28.0	-18% (-5.0 m/s)
Peak RMSE during reconfig., km	450.4	450.4	—
RMSE at day 85, km	15.65	N/A (failed)	—
Final RMSE at day 148, km	1.65	—	Restored

Table 3. Quantitative comparison of proactive vs. reactive reconfiguration strategies

6. Discussion

The results demonstrate three key advantages of the proposed integrated proactive monitoring and control approach compared to the reactive baseline. First, the LSTM model provides accurate and reliable health predictions (SoH RMSE 1.67%, $R^2 = 0.94$) consistently outperforming alternative neural network architectures. The multi-output prediction of SoH, RUL, and failure probability provides a richer decision basis than single-indicator approaches and enables probabilistic reconfiguration decisions with quantified uncertainty.

Second, the proactive control strategy delivers measurable energy savings of 18% (−5.0 m/s). In spacecraft operations, every metre per second of characteristic velocity saved translates directly to either extended mission lifetime or reduced propellant mass at launch. For a cluster of four spacecraft with typical propellant budgets of 50–80 m/s, a saving of 5 m/s per reconfiguration event represents approximately 6–10% of the total propellant reserve. Over multiple reconfiguration events in a longer mission, the cumulative saving becomes strategically significant.

Third, the modular CSV-interface architecture allows independent development, testing, and replacement of either the prognostic module or the control algorithm without redesigning the integrated system. This modularity is a key enabler for deployment across diverse automated system domains. The same integration pattern is directly applicable to industrial robot cells (degrading joint actuators), autonomous vehicle fleets (degrading battery packs), and process control systems (degrading sensor or actuator elements).

A limitation of the current approach is the reliance on a single degradation dataset for LSTM training and evaluation. The NASA Battery Dataset, while widely used, represents a specific chemistry and operating regime. Real deployment would require retraining or fine-tuning on application-specific data. Additionally, the current work assumes a single-failure scenario; simultaneous degradation of multiple units has not been addressed. Future work will explore transfer learning across equipment types, online model adaptation, multi-failure scenarios, and higher-fidelity orbital dynamics models.

7. Conclusions

A method for proactive monitoring of equipment condition in multi-agent systems using LSTM neural networks has been developed and validated. The multi-output LSTM model achieves SoH RMSE of 1.67% ($R^2 = 0.94$) and RUL MAE of 36 cycles on the NASA Battery Dataset, outperforming GRU by 12%, 1D-CNN by 40%, and vanilla RNN by 87% on SoH accuracy. The standardised CSV interface between the prognostic and control modules ensures modularity and scalability.

Integration of the LSTM prognostic module with a multi-agent potential field formation controller enables proactive autonomous reconfiguration triggered 100 hours before predicted failure. The case study of orbital small spacecraft cluster control demonstrates 18% ΔV savings (23.0 vs. 28.0 m/s) over a reactive baseline, with formation positioning RMSE restored to 1.65 km after reconfiguration. The approach generalises across diverse automated system domains. Future work will address multi-failure scenarios, online LSTM adaptation, transfer learning, and hardware-in-the-loop validation with higher-fidelity orbital models.

8. Acknowledgements

The research was carried out at Peter the Great St. Petersburg Polytechnic University. The authors thank the NASA Prognostics Center of Excellence for making the Battery Dataset publicly available.

9. References

- [1] Jardine, A.K.S.; Lin, D. & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance, *Mechanical Systems and Signal Processing*, Vol. 20, No. 7, pp. 1483-1510, DOI: 10.1016/j.ymssp.2005.09.012
- [2] Sikorska, J.Z.; Hodkiewicz, M. & Ma, L. (2011). Prognostic modelling options for remaining useful life estimation by industry, *Mechanical Systems and Signal Processing*, Vol. 25, No. 5, pp. 1803-1836, DOI: 10.1016/j.ymssp.2010.11.018
- [3] Hochreiter, S. & Schmidhuber, J. (1997). Long short-term memory, *Neural Computation*, Vol. 9, No. 8, pp. 1735-1780, DOI: 10.1162/neco.1997.9.8.1735
- [4] Li, X.; Zhang, W. & Ding, Q. (2019). Deep learning-based remaining useful life estimation of bearings using multi-scale feature extraction, *Reliability Engineering & System Safety*, Vol. 182, pp. 208-218, DOI: 10.1016/j.ress.2018.11.011
- [5] Saha, B. & Goebel, K. (2007). Battery Data Set, NASA Ames Prognostics Data Repository, NASA Ames Research Center, Moffett Field, CA. Available from: <https://ti.arc.nasa.gov/tech/dash/groups/pcoe/prognostic-data-repository/> Accessed: 2026-07-01

- [6] Hrnjica, B. & Softic, A. (2019). Explainable AI in manufacturing: a case study of predictive maintenance, Proceedings of the 30th DAAAM International Symposium, pp. 0184-0190, B. Katalinic (Ed.), Published by DAAAM International, ISBN 978-3-902734-22-8, ISSN 1726-9679, Vienna, Austria, DOI: 10.2507/30th.daaam.proceedings.026
- [7] Olfati-Saber, R. (2006). Flocking for multi-agent dynamic systems: algorithms and theory, IEEE Transactions on Automatic Control, Vol. 51, No. 3, pp. 401-420, DOI: 10.1109/TAC.2005.864190
- [8] Sollogub, A.V.; Skobelev, P.O.; Simonova, E.V.; Tsarev, A.V. & Stepanov, M.E. (2012). Models for solving network-centric planning and control tasks for group operations of a small spacecraft cluster, Modelirovanie Sistem i Protessov, No. 1, pp. 33-38
- [9] Filimonov, A.B. & Filimonov, N.B. (2021). Constructive aspects of the potential field method in mobile robotics, Avtometriya, Vol. 57, No. 4, pp. 45-53, DOI: 10.15372/AUT20210406
- [10] Sutton, G.P. & Biblarz, O. (2017). Rocket Propulsion Elements, 9th ed., John Wiley & Sons, Hoboken, NJ, 910 p., ISBN 978-1-118-75388-0
- [11] Sollogub, A.V.; Simonova, E.V. & Stepanov, M.E. (2011). Application of a multi-agent collective control system for a group of small spacecraft in Earth remote sensing tasks, Proceedings of the International Scientific-Practical Conference "Management of Large Systems-2011", Moscow: ICS RAS, Vol. 3, pp. 259-262
- [12] Monti Guarnieri, A.; Giudici, D.; Guccione, P.; Manzoni, M. & Rocca, F. (2020). MIMO SAR formations: orbital diameter and synchronization tolerances, The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, pp. 631-635

Working Paper of 37th DAAAM Symposium