API-Based Integration Framework for Dual-LLM Prescriptive Maintenance Report Generation in PHM-Enabled Digital Twin Applications

Atuahene Barimah[1], Chimkakwo Owhor[2], Octavian Niculita[3], Don McGlinchey[4]

[1,2,3,4] *Department of Engineering, Glasgow Caledonian University, Glasgow, G4 0BA, UK*

atuahene.barimah@gcu.ac.uk,cowhor300@caledonian.ac.uk,*octavian.niculita@gcu.ac.uk*,*d.mcglinchey@gcu.ac.uk,*

ABSTRACT

Digital Twins (DTs) have emerged as key enablers of Prognostics and Health Management (PHM) for predicting asset failures and optimizing maintenance strategies. However, translating predictive insights into actionable prescriptive maintenance plans remains a significant implementation challenge. This paper proposes an API-based framework that extends PHM-enabled DTs by incorporating a prescriptive maintenance layer aligned with enterprise operational constraints, including workforce scheduling, inventory availability, cost considerations, and compliance requirements. The framework integrates a generator model to produce maintenance recommendations and a checker model to evaluate report quality against operational criteria using a sequential model loading approach. Prescriptive maintenance reports using enterprise data as context are generated for a hydraulic system undergoing MCD scenarios. The framework proposed in this paper provides a low-cost implementation for integrating LLMs for prescriptive maintenance reporting for DT applications. This study contributes to LLM implementation use cases for DT applications for fleet asset management.

1. INTRODUCTION

Data has increasingly become central in enabling analytics in modern engineering systems as the availability of operational system data enables deeper insights into the behaviour of complex industrial assets. For digital transformation implementation use cases, the deployment of sensors, industrial Internet-of-Things (IIoT) platforms, and cloud-based data infrastructures enable continuous monitoring of system performance and asset health (Barimah, Onu, Niculita, Cowell, & McGlinchey, 2025). These developments have significantly influenced the design of Prognostics and Health Management (PHM) systems, where system data is analyzed to monitor degradation, detect anomalies, and predict potential failures before they occur (Vogl, Weiss & Helu, 2019).

Digital Twin (DT) technology further enhances these capabilities by enabling a virtual representation of physical assets that evolves alongside the physical system throughout its lifecycle. The concept of a Digital Twin was first introduced by Grieves and Vickers as a framework that links a physical system with its virtual representation through continuous data exchange and lifecycle information management (Grieves & Vickers, 2017). By synchronizing operational data with simulation and analytical models, DT-enabled analytics provide a link between physical assets and analytical models, allowing engineers to evaluate system behaviour and anticipate degradation events before failures occur. Kritzinger, Karner, Traar, Henjes & Sihn (2018); Fuller, Fan, Day & Barlow (2020).

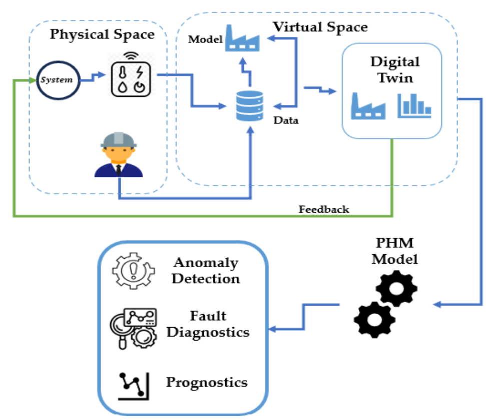

Figure 1. Relationship between DT and PHM applications. (Barimah, Niculita, McGlinchey, Cowell & Milligan, 2024)

Predictive maintenance has therefore emerged as a key PHM strategy for improving operational efficiency by identifying degradation patterns and forecasting equipment failures using historical and real-time sensor data (Carvalho, Soares, Vita, Francisco, Basto & Alcalá, 2019; Barimah et al. 2025). When sufficient operational data is available, predictive models can learn patterns in system behaviour and generalize to unseen operating conditions, thereby improving fault detection and maintenance planning capabilities (Duriez, Brunton & Noack 2017). Such approaches have demonstrated significant potential for reducing unplanned downtime and optimizing maintenance schedules in complex engineering systems (Zonta, da Costa, da Rosa Righi, de Lima, da Trindade & Li 2020).

Despite these advances, most predictive maintenance implementations focus primarily on fault detection and failure prediction rather than providing actionable maintenance recommendations. Maintenance engineers frequently receive alerts indicating potential system degradation, but must still determine the appropriate maintenance actions, resource allocation, and execution timeline required to address the issue. This limitation creates a gap between predictive PHM analytics and practical maintenance decision-making processes (Bousdekis, Lepenioti, Apostolou & Mentzas 2021). In real industrial environments, maintenance decisions depend not only on predicted asset health but also on operational constraints such as technician availability, spare parts inventory, maintenance policies, and regulatory compliance requirements (Cachada, Barbosa, Leitão, Geraldes, Deusdado, Costa, Teixeira, Teixeira, Moreira, Moreira & Romero 2018).

Recent research has explored the integration of Digital Twin architecture with advanced analytics to address some of these challenges. For example, recent PHM research has investigated the use of hybrid modelling approaches for analyzing complex systems undergoing multi-component degradation scenarios within Digital Twin environments (Barimah et al., 2024; Li, Chan, Zaman, Apostolou & Conroy, 2022). Such studies highlight the importance of combining data-driven models with system knowledge to improve diagnostic and prognostic performance.

More recently, Artificial Intelligence techniques have introduced new possibilities for enhancing decision-support capabilities within PHM systems. Large Language Models (LLMs) have demonstrated strong capabilities in interpreting heterogeneous data sources and generating structured textual outputs (Brown, Mann, Ryder, Subbiah, Kaplan, Dhariwal, Neelakantan, Shyam, Sastry et al. 2020; Zhao, Zhou, Li, Tang, Wang, Hou, Min, Zhang, Zhang et al. 2023). Unlike traditional machine learning models that primarily process numerical datasets, LLMs can interpret both structured and unstructured information, including maintenance documentation, operational logs, and system data. This capability creates opportunities for transforming predictive maintenance insights into prescriptive maintenance recommendations that provide clear guidance for maintenance engineers (Wang & Li 2023).

However, integrating LLM capabilities with Digital Twin infrastructures requires reliable mechanisms for exchanging operational data between system components. Application Programming Interfaces (APIs) provide an effective architecture for enabling such integration by facilitating data ingestion, processing, and communication between PHM analytics platforms and AI models (Rasheed, San & Kvamsdal, 2020). Through API-based architecture, Digital Twin outputs can be combined with operational resource data and interpreted by LLM systems to generate structured prescriptive maintenance reports. This paper addresses the gap between predictive PHM analytics and actionable maintenance planning by proposing an API-based framework that integrates Digital Twin health predictions with Resource Availability Data to enable prescriptive maintenance reporting.

2. OBJECTIVES OF STUDY

This paper aims to develop and evaluate an API-based framework for automated prescriptive maintenance report generation by integrating Digital Twin health predictions with Resource Availability Data through a dual-LLM approach. The reason for this is that it allows stakeholders seeking to protect their sensitive data from being stored in third party LLM service providers to integrate LLM models into their DT applications while having autonomy over their sensitive data. The objectives of this study are listed below:

1. Design an API-based framework that integrates Digital Twin health predictions with Resource Availability Data to enable automated prescriptive maintenance report generation.
2. Implement a dual-LLM architecture using a generative adversarial network (GAN) approach consisting of a generator model that produces initial prescriptive maintenance recommendations and a checker model that evaluates report quality against operational requirements, including maintenance action clarity, completeness of resource allocation, and execution feasibility.
3. Validate the proposed framework using DT predictive analytics for multi-component degradation scenarios for a hydraulic system.
4. Demonstrate automated generation of prescriptive maintenance reports using enterprise level context.

Section 3 in this paper covers the methodology for deploying prescriptive maintenance report generation. Sections 4 and 5 present and discuss the results for implementing the API based dual LLM framework. Finally, the paper concludes with contributions and future research work.

3. METHODOLOGY

3.1. System Architecture Overview

The methodology for developing the prescriptive maintenance model involves integrating a pre-trained LLM prescriptive maintenance strategy as a microservice part of a PHM enabled DT deployment strategy for a hydraulic system (Barimah A. , Niculita, McGlinchey, & Cowell, 2023) as used by the author in a previous publication (see Figures 2 and 3). Requests from another application within the cluster send API calls to the LLM prescriptive maintenance application using the sequence shown in Figure 3. The pre-trained model is evaluated using synthetic hydraulic pump data representing typical PHM degradation scenarios. The proposed model is integrated with DT predictive analytics to improve operational maintenance planning processes.

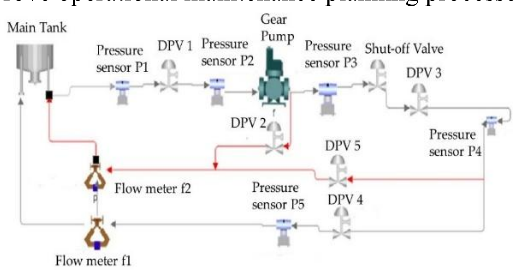

Figure 2: Hydraulic System Testbed Schematic (Barimah et al. 2023).

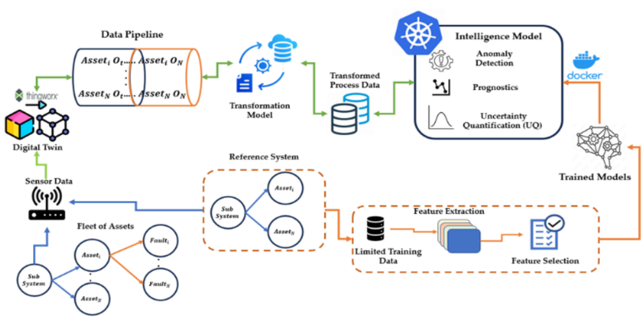

Figure 3: Microservices approach to scalable DT deployment (Barimah et. al 2025)

The LLM framework combines asset health information obtained from the DT analytics applications with operational resource availability constraints to generate prescriptive maintenance reports automatically. The constraints in this case are largely Resource Availability Data which provides operational constraints related to maintenance personnel and other enterprise level resources. The sequence of events in the API workflow for the prescriptive report generation is shown in Figure 4.

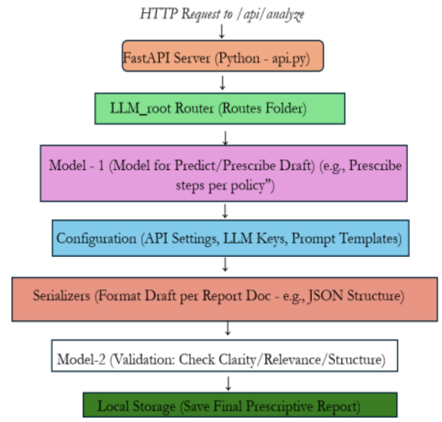

Figure 4: API client request workflow for the prescriptive microservice

To generate prescriptive maintenance reports using a Large Language Model, the conceptual architecture involves a dual-model reasoning approach consisting of a generator component and a checker component. The generator model—which is trained using the process in Figure 5—is responsible for producing the initial prescriptive maintenance report based on the output of the DT analytics and Resource Availability Data. The report generated describes the required maintenance actions for the identified faulty pumps and specifies the operational resources required to perform the maintenance activities with enterprise context. The checker model performs a secondary evaluation of the generated report to assess whether the recommendations satisfy predefined operational criteria. This includes the clarity of maintenance actions, completeness of resource allocation, and feasibility of the proposed maintenance schedule.

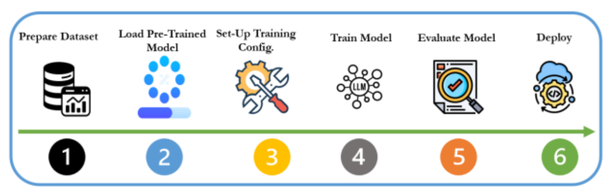

Figure 5: Workflow for Training pretrained models for prescriptive maintenance report generation.

The generator and checker operate as separate LLM components to enable independent reasoning and validation processes (see Figure 6). However, due to computational constraints associated with large-scale LLM model training, the current implementation employs a lightweight transformer-based model (Phi-3-mini-4k-instruct) to perform both roles. Through prompt engineering techniques, the model is first used to generate prescriptive maintenance reports and subsequently to evaluate the generated output against the defined operational criteria. The proposed conceptual design is to be model-agnostic, meaning that the underlying language model can be replaced with more advanced models such as Qwen2.5-Math-72B-Instruct or other large-scale LLM architectures during deployment using higher compute resources without modifying the overall system architecture. The output produced by the LLM module is a structured prescriptive maintenance report that includes recommended maintenance actions, allocation of available technicians, spare parts requirements, and execution timelines.

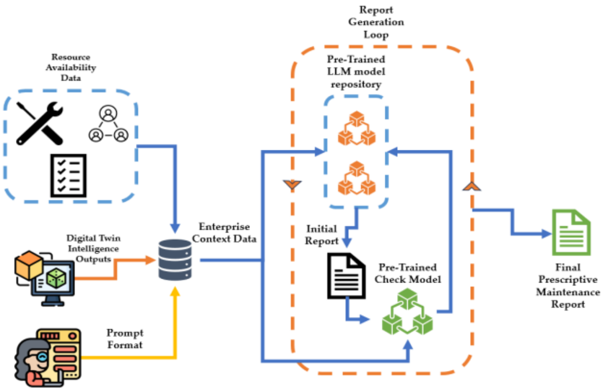

Figure 6: Prescriptive Maintenance Model Overview

3.2. Digital Twin and Resource Availability Data

In this study, Digital Twin asset health predictive outputs for multi component degradation scenarios (MCD) are done using the microservices framework (see Figure 3) used by the author in a previous publication. The performance level of each asset is predicted using a hybrid DT predictive analytics model to predict asset health in MCD scenarios. Fault identification is performed by applying a predefined threshold to the state value, where assets with state values below the threshold are classified as degraded assets requiring maintenance intervention. The Digital Twin, therefore, acts as the diagnostic layer of the framework by identifying faulty assets within the fleet of assets. Table 1 presents the MCD test scenarios used with Figure 7 showing the output of the predictive analytics for T13.

Table 1: Test Degradation Scenarios (Barimah et. al, 2024)

Dataset	Degradation (Component 1)	Degradation (Component 2)	Operational Speed Range (RPM)
T13	Pump (Medium-severity) at 45% DPV2 opening	Constant degradation of Nozzle: 30-70% DPV1 opening	700
T15	Filter (High severity) at 32% DPV1 opening	N/A	700 to 950
T20	Pump (Medium-severity) at 55% DPV2 opening	Nozzle (High severity) at 30% DPV4 opening	700 to 950

```
{
  "request_id": 5,
  "_id": "69c0185a230abd31946f8932",
  "status": "ok",
  "SPC_pred_C103A": 1,
  "components_prediction": {
    "Time": 35.08,
    "Theta": 127.0,
    "Pump_Operating_Point": 1.2466541528701782,
    "Filter_Health": 0,
    "Filter_Performance": 0.91,
    "Pump_Health": 1,
    "Pump_Performance": 0.59,
    "Pump Volumetric Efficiency (%)": 15.25,
    "Valve_Health": 0,
    "Valve_Performance": 0.98,
    "Nozzle_Health": 1,
    "Nozzle_Performance": 0.45,
    "Pipe_Health": 0,
    "Pipe_Performance": 1.81
  },
  "execution_time_minutes": 0.16
}
```


Figure 7: Hybrid DT predictive analytics output for T13

In addition to Digital Twin outputs, the framework integrates Resource Availability Data (RAD) representing operational constraints associated with maintenance planning (see Figure 8 and Table 2). RAD includes information about both human and material resources required to perform maintenance activities. The human resource component describes the available maintenance workforce, including the number of senior and junior technicians available to perform maintenance tasks. This information enables the framework to allocate appropriate personnel to the recommended maintenance activities. The material resource component captures the availability of spare parts required for pump maintenance and component replacement. Spare parts availability is critical in maintenance decision-making because repair or replacement actions cannot be executed if the required components are unavailable. By combining Digital Twin health predictions with Resource Availability Data, the framework incorporates both asset condition monitoring and operational feasibility constraints into the maintenance decision process.

Figure 8: Components of Resource Availability Data

Table 2: Description of Resource Availability Data used for Prescriptive Maintenance Reporting

Resource	Description
Pre-trained LLM Model	Generator Model 1. microsoft/Phi-3-mini-128k-instruct 2. Number of Tokens: 450 3. Model size: 7.12GB
	Discriminator Model 1. microsoft/Phi-3-mini-128k-instruct

	2. Number of Tokens: 50 3. Model size: 7.12GB
	Loop count: Three (3) iterations
Resource Constraints	Compute: 8GB RAM
Regulations	1. Supply of Machinery (Safety) Regulations 2008. 2. Provision and Use of Work Equipment Regulations 1998 (PUWER) 3. UKCA (UK Conformity Assessed) marking as of January 2025. 4. Health and Safety Executive (HSE).
Maintenance Records and Policy	Enterprise level maintenance policy

3.3. API Integration and Structured Prompt Construction

To enable seamless communication between the Digital Twin environment, resource databases, and the reasoning component, the design was framed to employ an API-based integration architecture (See Figure 9). Application Programming Interfaces (APIs) provide a flexible mechanism for integrating heterogeneous data sources and coordinating the flow of information between system components. Within the proposed design framework, the API layer performs three primary functions. First, it ingests Digital Twin outputs representing asset health conditions and operational resource data describing available maintenance personnel and spare parts. Second, the API structures the received data into a standardized format suitable for downstream processing. Finally, the API orchestrates the workflow by forwarding the processed inputs to the reasoning component responsible for generating maintenance recommendations.

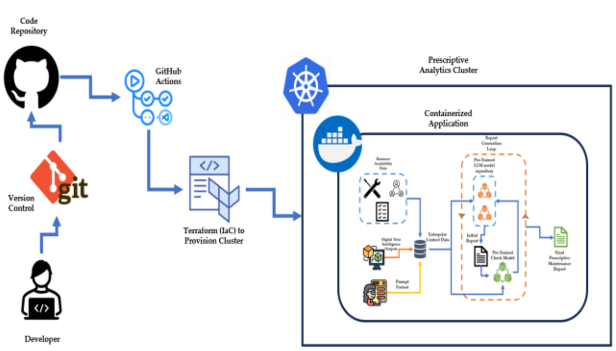

Figure 9: Deployment Framework for Prescriptive analytics

Following data integration, the framework integrates a structured prompt that serves as the input to the LLM reasoning module. The prompt is generated by combining the Digital Twin outputs with the Resource Availability Data to produce a concise representation of the maintenance scenario. The prompt contains information describing the number of faulty pumps, the affected components, and the available maintenance resources. The structured prompt also specifies the required format for the generated output to ensure that the LLM produces a consistent prescriptive maintenance report. This structured input formulation allows the LLM to interpret system conditions and generate recommendations that consider both asset health and operational resource constraints.

4. RESULTS

The proposed framework was evaluated using synthetic hydraulic pump degradation data representing typical Prognostics and Health Management (PHM) scenarios involving single and multi-component degradation (MCD). The objective was to assess the capability of the Digital Twin (DT)-enabled prescriptive maintenance framework to transform asset health predictions and operational resource constraints into actionable maintenance recommendations. The evaluation focused on the complete deployment architecture which integrates DT analytics, maintenance decision support, and enterprise-level operational management within a distributed microservice environment. Figure 10 presents the end-to-end DT deployment architecture comprising three primary network domains: the Client Asset Network, Enterprise Network, and Analytics Network. The Client Asset Network contains the physical assets, programmable logic controllers (PLCs), and associated sensing infrastructure responsible for generating operational process data. The data is transmitted to the Enterprise Network, which hosts the DT platform and provides the operational interface through which users, maintenance teams, and DT developers interact with the system. The DT platform maintains process visualization, asset health monitoring, and system-level decision support capabilities. The Analytics Network hosts cloud-native DT intelligence services deployed as microservices within a Kubernetes environment. These services include the analytics microservice responsible for diagnostic and prognostic reasoning, maintenance microservice responsible for maintenance planning and report generation, and supporting user-management services. Together, these components provide the diagnostic, prognostic, and prescriptive intelligence required to support PHM-enabled decision making.

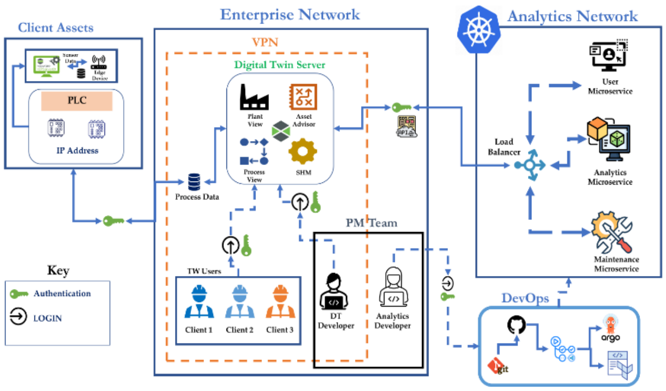

Figure 10: End-to-end Digital Twin deployment

The identified degradation states from the analytics microservice are combined with Resource Availability Data—including maintenance personnel availability and spare-parts inventory information—to support maintenance planning. The maintenance microservice generated prescriptive maintenance reports for each degradation scenario. An example report is presented in Figure 11, while Table 3 summarizes the reports generated across multiple MCD test scenarios. The generated reports identify the affected assets, such as pumps and nozzles, and provide recommended maintenance actions, required technical resources, proposed intervention schedules, and estimated operational costs. In addition, each report includes a summary of the number of degraded assets addressed and the projected maintenance expenditure associated with the recommended actions. The results show the capability of the proposed DT architecture of successfully integrating diagnostic and prognostic outputs with enterprise resource context prescriptive maintenance recommendations. Also, the distributed deployment architecture used illustrates how diagnostic, prognostic, and prescriptive services can be deployed as scalable microservices, supporting the practical implementation of DT-enabled maintenance decision support within industrial environments.

Table 3: Prescriptive Maintenance reports generated for Test scenarios in Table 2 (see Appendix A.1)

Test Scenario	Prescriptive Maintenance Report
T13	gcu_asset_report_T13.txt
T15	gcu_asset_report_T15.txt
T20	gcu_asset_report_T20.txt

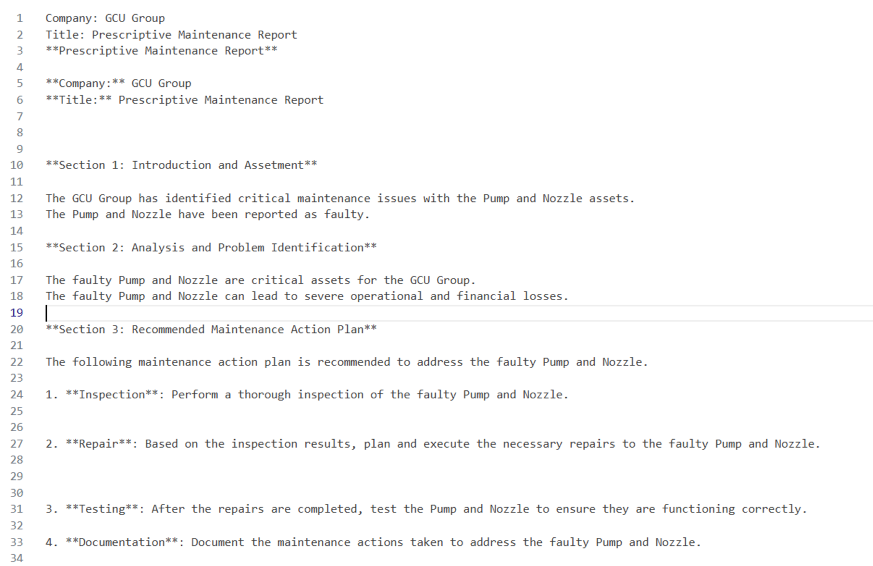

```
Company: GCU Group
Title: Prescriptive Maintenance Report
**Prescriptive Maintenance Report**

**Company:** GCU Group
**Title:** Prescriptive Maintenance Report



**Section 1: Introduction and Assetment**

The GCU Group has identified critical maintenance issues with the Pump and Nozzle assets.
The Pump and Nozzle have been reported as faulty.

**Section 2: Analysis and Problem Identification**

The faulty Pump and Nozzle are critical assets for the GCU Group.
The faulty Pump and Nozzle can lead to severe operational and financial losses.

**Section 3: Recommended Maintenance Action Plan**

The following maintenance action plan is recommended to address the faulty Pump and Nozzle.

1. **Inspection**: Perform a thorough inspection of the faulty Pump and Nozzle.


2. **Repair**: Based on the inspection results, plan and execute the necessary repairs to the faulty Pump and Nozzle.



3. **Testing**: After the repairs are completed, test the Pump and Nozzle to ensure they are functioning correctly.

4. **Documentation**: Document the maintenance actions taken to address the faulty Pump and Nozzle.
```

Figure 11: Example of Generated Prescriptive Maintenance Report

5. DISCUSSION

The generated prescriptive maintenance report demonstrates the feasibility of integrating Digital Twin health predictions with pre-trained LLMs reasoning to support prescriptive maintenance decision-making. Unlike traditional predictive maintenance systems that primarily focus on fault detection or remaining useful life estimation, the proposed framework translates asset health predictions into actionable maintenance plans that incorporate operational constraints specific to a defined enterprise context. One important observation from the generated reports is the ability of the system to integrate multiple types of operational information within a single reasoning process. By combining Digital Twin outputs with Resource Availability Data, the framework generates maintenance recommendations that consider both asset health conditions and available operational resources such as maintenance personnel and spare parts inventory. This integration is particularly important in industrial environments where maintenance planning must balance technical requirements with operational feasibility. The use of prompt engineering to guide LLM reasoning also demonstrates the potential of language models for structured engineering decision support. The structured prompts used in the framework enable the model to interpret maintenance scenarios and generate consistent report formats containing relevant operational details. This capability highlights the potential of LLMs to function as reasoning layers within Digital Twin-enabled PHM systems.

6. CONCLUSION

Despite the promising results of the proposed implementation, several limitations remain. The current evaluation relies on large compute resources to run LLMs locally to fully capture the complexity of real industrial systems. There is a trade-off therefore between enterprise data autonomy and the cost of compute for optimal report generation. The framework proposed in this paper provides a low-cost implementation for integrating LLMs for prescriptive maintenance reporting for DT applications. While the results demonstrate the practical viability of the proposed framework, additional validation is required to assess the quality and trustworthiness of the LLM-generated maintenance reports, preferably using real datasets and the associated analytics capturing equipment degradation as well as specific resource availability data. However, the current evaluation focuses primarily on demonstrating the end-to-end functionality of the DT-to-prescriptive-maintenance pipeline. The contributions of the paper are:

C1. This study describes a practical implementation of LLM use cases for DT applications for fleet asset management.

C2. The paper also presents a microservices GAN based approach for pre-trained LLMs for prescriptive maintenance asset health reporting in the context of DTs.

Future research of this paper will cover more rigorous quantitative evaluation metrics as well as the impact on the prescriptive maintenance reporting process of the increased complexity for the constraints of the resource allocation data. This would contribute to the development of trustworthy and explainable AI-enabled DT systems capable of supporting enterprise-scale maintenance operations.

NOMENCLATURE

API	Application Programming Interface
DT	Digital Twin
DPV	Direct Proportional Valve
GAN	Generative Adversarial Network
IIoT	Industrial Internet of Things
LLM	Large Language Models
MCD	Multi- Component Degradation
PHM	Prognostic and Health Management
RAD	Resource Availability Data
RPM	Revolutions per minute

REFERENCES

Barimah, A.K., Jahanzeb, A., Niculita, O., Cowell, A. and McGlinchey, D., 2025. Benchmarking Control Strategies for Multi-Component Degradation (MCD) Detection in Digital Twin (DT) Applications. Computers, 14(9), p.356.

Barimah, A., Niculita, O., McGlinchey, D., Cowell, A. and Milligan, B., 2024, June. Towards Physics-Informed PHM for Multi-component degradation (MCD) in complex systems. In PHM Society European Conference (Vol. 8, No. 1, pp. 14-14).

Barimah, A.K., Onu, O.P., Niculita, O., Cowell, A. and McGlinchey, D., 2025. Scalable Data Transformation Models for Physics-Informed Neural Networks (PINNs) in Digital Twin-Enabled Prognostics and Health Management (PHM) Applications. Computers, 14(4), p.121.

Barimah, A., Niculita, O., McGlinchey, D. & Cowell, A., 2023. Data-quality assessment for digital twins targeting multi-component degradation in industrial internet of things (IIoT)-enabled smart infrastructure systems. *Applied Sciences,* 13(24), p. 13076

Bousdekis, A., Lepenioti, K., Apostolou, D. and Mentzas, G. (2021) 'A review of data-driven decision-making methods for Industry 4.0 maintenance applications', Electronics, 10(7), p. 828.

Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I. and Amodei, D. (2020) 'Language models are few-shot learners', arXiv [Preprint]. arXiv:2005.14165.

Cachada, A., Barbosa, J., Leitão, P., Geraldes, C.A.S., Deusdado, L., Costa, J., Teixeira, C., Teixeira, J., Moreira, A.H.J., Moreira, P.M. and Romero, L. (2018) 'Maintenance 4.0: Intelligent and Predictive Maintenance System Architecture', in 2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA), Turin, Italy, pp. 139–146.

Carvalho, T.P., Soares, F.A.A.M.N., Vita, R., Francisco, R.d.P., Basto, J.P. and Alcalá, S.G.S. (2019) 'A systematic literature review of machine learning methods applied to predictive maintenance', Computers & Industrial Engineering, 137, p. 106024.

Duriez, T., Brunton, S.L. and Noack, B.R. (2017) Machine Learning Control: Taming Nonlinear Dynamics and Turbulence. Fluid Mechanics and Its Applications, vol. 116. Cham: Springer International Publishing.

Eker, O.F., Camci, F. and Jennions, I.K. (2016) 'Physics-based prognostic modelling of filter clogging phenomena', Mechanical Systems and Signal Processing, 75, pp. 395–412.

Fuller, A., Fan, Z., Day, C. and Barlow, C. (2020) 'Digital twin: Enabling technologies, challenges and open research', IEEE Access, 8, pp. 108952–108971.

Grieves, M. and Vickers, J. (2017) 'Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems', in Kahlen, F.J., Flumerfelt, S. and Alves, A. (eds.) Transdisciplinary Perspectives on Complex Systems. Cham: Springer, pp. 85–113.

Kritzinger, W., Karner, M., Traar, G., Henjes, J. and Sihn, W. (2018) 'Digital twin in manufacturing: A categorical literature review and classification', IFAC PapersOnLine, 51(11), pp. 1016–1022.

Li, Y., Chan, D., Zaman, N., Apostolou, E. and Conroy, P., 2022, October. ML Detection and Isolation of Functional Failures using Syndrome Diagnostics. In Annual Conference of the PHM Society (Vol. 14, No. 1).

Rasheed, A., San, O. & Kvamsdal, T., 2020. Digital twin: Values, challenges and enablers from a modeling perspective. IEEE access, Volume 8, pp. 21980-22012.

Vogl, G.W., Weiss, B.A. and Helu, M. (2019) 'A review of diagnostic and prognostic capabilities and best practices for manufacturing', Journal of Intelligent Manufacturing, 30, pp. 79–95.

Wang, H. and Li, Y.F. (2023) 'Large language model empowered by domain-specific knowledge base for industrial equipment operation and maintenance', in Proceedings of the 5th International Conference on System Reliability and Safety Engineering (SRSE), Beijing, China, pp. 474–479.

Zhao, W.X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., Du, Y., Yang, C., Chen, Y., Chen, Z., Jiang, J., Ren, R., Li, Y., Tang, X., Liu, Z., Liu, P. and Wen, J.R. (2023) 'A survey of large language models', arXiv [Preprint]. arXiv:2303.18223.

Zonta, T., da Costa, C.A., da Rosa Righi, R., de Lima, M.J., da Trindade, E.S. and Li, G.P. (2020) 'Predictive maintenance in the industry 4.0: A systematic literature review', Computers & Industrial Engineering, 150, p. 106889.

BIOGRAPHIES

Atuahene Barimah is currently a PhD researcher at Glasgow Caledonian University and a 2026 WCSIM Scholar. His research interests include data-driven maintenance, system reliability, process control, project management, operations research, digital twin design, IIoT, and computational finance.

Chimkakwo Kinikanwo Owhor graduated from the Federal University of Petroleum Resources, Nigeria, with a BSc in Physics (Hons), then worked as a Risk and Safety Analyst in the oil and gas sector before moving into healthcare. He recently completed an MSc in Applied Data Science in Engineering (Distinction) from Glasgow Caledonian University, UK. His research interests include prescriptive maintenance, large language model applications in industrial settings, digital twin technologies, and sustainable development. His current work focuses on resource-efficient LLM-driven frameworks for maintenance decision support in engineering systems.

Octavian Niculita is a Senior Lecturer in Instrumentation with Glasgow Caledonian University. He has a PhD in Industrial Engineering from the Technical University of Iasi, Romania carried out under the EDSVS framework. His current research interests include industrial digitalization, predictive maintenance, PHM system design, integration of PHM and asset design for aerospace, maritime, and oil & gas (surface and subsea) applications. Octav has over ten years of experience in design and development of prognostics and health management applications, having worked on applied aerospace projects funded by The Boeing Company and BAE Systems as a Research Fellow and Technical Lead on his previous appointment with the IVHM Centre at Cranfield University, UK. He is a member of the Prognostics and Health Management Society, InstMC and the IET.

Don McGlinchey graduated from Strathclyde University with a BSc (Hons) Physics before working as a project engineer at Babcock Energy Ltd. He returned to academia and gained an MSc in Bulk Solids Handling Technology and

his Doctorate on a study of the effect of vibration on powder beds. He is currently a Professor in the Department of Engineering at Glasgow Caledonian University where he is the academic leader in teaching, research, and consultancy in the area of multi-phase flow. He has edited two books, and authored over 100 papers, articles, and consultancy reports.

APPENDIX

Appendix A.1 Code for generating prescriptive maintenance reports

GitHub Repository: https://github.com/AtuaheneBarimah/DT_Prescriptive_llm_phme26.git