Just-in-Time Maintenance with Explainable AI for Industrial IoT: A Comprehensive Survey

Aniket Dhumal*, Mansi Lokhande*, Sharvari Chaure*, Vaishali Kandekar*
*Department of Computer Engineering,
Pune Institute of Computer Technology (PICT),
Dhankawadi, Pune, MH 411046, India

***Abstract*—Combining Just-in-Time (JIT) maintenance with Explainable Artificial Intelligence (XAI) in the Industrial Internet of Things (IIoT) opens up exciting ways to streamline industrial operations. This survey shines a light on how pairing JIT with XAI can boost predictive maintenance in IIoT setups. We zero in on tackling the murky nature of black-box AI models in factories, emphasizing the need for clear, understandable results to meet rules and earn trust. After digging into existing studies, pinpointing where research falls short, we sketch out a solid plan for weaving XAI into JIT maintenance. We dive into how tools like SHAP and LIME work in real-time systems and wrestle with the hiccups of putting them to use. By delivering straightforward, practical insights, this approach aims to sharpen decision-making, trim downtime, and lift efficiency across industrial settings.**

***Index Terms*—Just-in-Time Maintenance, Explainable AI, Industrial IoT, Anomaly Detection, Predictive Maintenance, Machine Learning, Deep Learning**

I. INTRODUCTION

The Industrial Internet of Things (IIoT) is shaking up old-school manufacturing and industrial work by tying together devices and systems for real-time data tracking and analysis (L. D. Xu, W. He, and S. Li, 2014 [1]). This tech shift paves the way for clever maintenance tricks that cut costs and keep things humming along smoothly.

Here's the rub, though: bringing AI into predictive maintenance for IIoT often hits a snag. Those tricky "black-box" models make it tough to figure out why they flag issues, which is a real problem where trust and regulations are non-negotiable (C. Rudin, [2]). Folks on the ground need to grasp the reasoning behind a warning, not just take it on faith.

That's where this survey steps in. We're digging into how Just-in-Time (JIT) maintenance—fixing gear right when it needs it—can pair up with Explainable AI (XAI) to make predictions both spot-on and crystal clear. Unlike past efforts obsessed with nailing accuracy, we're all about shedding light on the "why" behind AI calls, leaning on tools like SHAP and LIME to fit industrial needs and build confidence. This fills a gap, pushing predictive maintenance into a space that's as understandable as it is effective.

II. BACKGROUND

A. Industrial Internet of Things (IIoT)

The Industrial Internet of Things (IIoT) hooks up factory machines with sensors, software, and networks to gather and swap data on the fly [6]. It's a game-changer, letting machines chat with each other in real time, powering up automation, and crunching data to make sharper calls that amp up efficiency and spark fresh business ideas.

B. Just-in-Time Maintenance

Just-in-Time (JIT) maintenance comes from lean manufacturing's playbook, aiming to nix waste and boost efficiency by fixing stuff only when it's truly needed [?]. It's all about keeping a close eye on gear in real time, using data smarts to spot trouble early, and timing repairs just right to keep production rolling without hiccups.

C. Explainable Artificial Intelligence (XAI)

Explainable Artificial Intelligence (XAI) is about cracking open AI's decision-making so people can follow along [8]. In places where safety and rules matter most, that's a big deal. Some XAI tricks are built for specific setups—like decision trees you can trace—while others, like LIME and SHAP, work across the board, unpacking predictions after the fact to keep things transparent and reliable.

III. LITERATURE REVIEW

A. Predictive Maintenance in IIoT

Predictive maintenance taps IIoT data to catch equipment hiccups before they hit [9]. It uses stats to spot weird trends, machine learning—like Random Forests or Neural Networks—to learn from past flops, and deep learning—like CNNs or RNNs—to wrestle with complex patterns [10]. These tricks keep gear ticking and downtime slim.

B. Challenges of Black-Box AI Models

Rolling out black-box AI in factories stirs up some real headaches:

- **Regulatory Compliance Issues**: Strict rules in some industries—like the EU's GDPR—demand clear reasoning behind automated calls (S. Wachter, B. Mittelstadt, and C. Russell, 2017 [3]). Black-box setups muddy that up.
- **Operational Hesitations**: If workers can't peek under the hood of a prediction, they're slow to trust it, sidelining AI's potential (D. Gunning, 2019 [4]).

- **Debugging Difficulties**: Without a clue how a model thinks, fixing glitches or biases turns into a costly guessing game (W. Samek, T. Wiegand, and K.-R. M¨uller, 2017 [5]).

These snags make a loud case for AI that's open and easy to follow.

C. Explainable AI Techniques

XAI dishes out ways to make sense of AI's inner workings. SHAP leans on game theory to rank how much each piece of data sways a prediction [11]. LIME zooms in on single calls, whipping up simple explanations for any model [12]. Layer-wise Relevance Propagation (LRP) picks apart predictions to spotlight key inputs [13], and counterfactuals play "what if" to show how tweaks flip outcomes [3]. Together, they cut through AI's fog.

D. Integration of XAI in Predictive Maintenance

Mixing XAI into predictive maintenance brings some hefty wins. It lays bare why a failure's looming, so workers know what's up. Those clear breakdowns guide smart fixes and let teams tweak models with real-world feedback, making the whole system sharper and more trustworthy.

E. Implemented Systems and Case Studies

Big names are jumping on this. Siemens weaves XAI into Mindsphere for clear maintenance insights [14]. GE Predix uses it to make gear analytics transparent [15]. Hitachi's Lumada and IBM's Watson IoT lean on explainable models to fine-tune asset care [16], [17]. Research is in on it too—like using SHAP to predict turbine troubles [18]. These real-world wins show XAI's muscle in maintenance.

IV. Proposed Framework

A. Overview

Our framework ties JIT maintenance with XAI in IIoT setups (Figure 1). It's got six pieces:

1) **Data Collection and Ingestion**
2) **Data Preprocessing and Feature Engineering**
3) **Anomaly Detection and Predictive Modeling**
4) **Explainability Module**
5) **Maintenance Decision Support**
6) **User Interface and Visualization**

B. Data Collection and Ingestion

A solid IIoT system starts with sensors pulling real-time stats—like temperature or vibration—and networks like MQTT or OPC UA shuttling it around fast. Time-series databases keep that flood of data tidy and ready for action.

C. Data Preprocessing and Feature Engineering

Messy data needs a cleanup—scrubbing noise, patching holes, and scaling it so models don't trip over it. Then we tease out handy bits—like averages or frequency quirks—to juice up prediction power.

D. Anomaly Detection and Predictive Modeling

Spotting trouble means picking the right tools—think Autoencoders or Isolation Forests—training them on what's normal, and letting them flag oddballs as they pop up in real time. It's about nipping issues in the bud.

E. Explainability Module

Here's where XAI gets busy. SHAP cranks out real-time breakdowns of what's driving predictions, with visuals like summary plots to make it click (Lundberg and Lee, 2017 [11]). LIME jumps in with quick, local explanations when SHAP's too slow for the pace (Ribeiro, Singh, and Guestrin, 2016 [12]). To keep it snappy, we lean on precomputing, shortcuts, and GPU muscle (Arrieta et al., 2020 [19]).

F. Maintenance Decision Support

A recommendation engine dishes out fix ideas based on what's flagged and why, timing repairs to dodge chaos, and lining up crews and parts so it all goes smooth. It's practical know-how in action.

G. User Interface and Visualization

Dashboards give workers a live look at gear health, anomaly warnings, and XAI insights. Alerts ping them when stuff's urgent, and they can poke around the data to get the full picture. It keeps everyone sharp and ready.

V. Implementation Example

A. Scenario

Picture a plant keeping tabs on key machines with IIoT sensors, aiming to spot hiccups and fix them right on time.

B. Data Collection

- **Sensors**: Collect data on temperature, vibration, pressure, RPM.
- **Data Rate**: High-frequency sampling (e.g., 1 Hz).

C. Modeling

- **Anomaly Detection Model**: Use an Autoencoder neural network trained on normal operation data.
- **Explainability**: Apply SHAP to interpret anomalies.

D. Challenges and Model Optimization

We ran into bumps along the way:

- **Data Quality Issues**: Spotty or noisy sensor readings called for tough cleanup to keep the model on track.
- **Model Selection and Optimization**: Finding the sweet spot between complexity and clarity took work—an Autoencoder with just enough layers nailed it without overcooking the data.
- **Computational Constraints**: Real-time needs pushed us to slim down the model and speed it up with tricks like cutting dimensions.

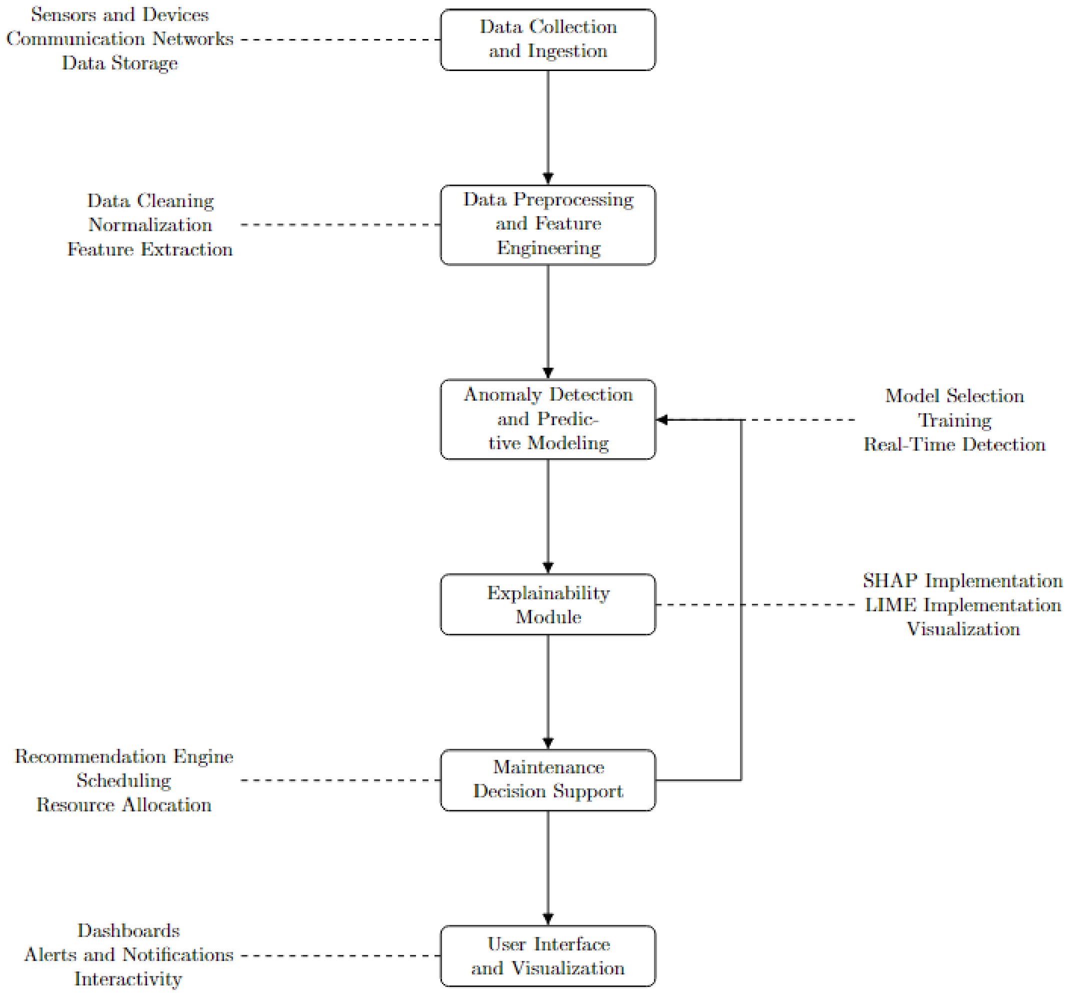

Fig. 1: Proposed Framework for JIT Maintenance with XAI in IIoT Systems

E. Explainability with SHAP

SHAP made the difference, step by step:

1) **Anomaly Detection**: The Autoencoder flagged high-error cases.
2) **SHAP Analysis**: SHAP pinned down what each feature added to the mix.
3) **Interpretation**: Big SHAP scores for vibration and temperature pointed the finger there.
4) **Actionable Decisions**: Crews zeroed in on bearings and cooling setups.
5) **Feedback Loop**: Those insights sharpened the model for next time.

F. Visualization

A dashboard lays it out:

- **Real-time Sensor Data**: Graphs tracking readings over time.
- **Anomaly Alerts**: Flagging trouble spots.
- **SHAP Explanations**: Bar charts showing what's behind it.
- **Maintenance Recommendations**: Tips on what to do next.

VI. Discussion

A. Benefits

This setup trims downtime with timely fixes, saves bucks by dodging extra repairs, and wins trust with clear breakdowns. It's a triple threat for smoother, cheaper operations.

B. Challenges

Good data's a slog to get, tricky models can muddy the waters, and plugging this into old systems takes cash and know-how. It's a steep climb worth tackling.

C. Ethical and Legal Considerations

Rules like GDPR push for clear answers (S. Wachter et al., 2017 [3]), while standards like ISO/IEC 27001 and IEC 61508 lock in safety and accountability [22]–[24]. XAI keeps it all above board in high-stakes spots.

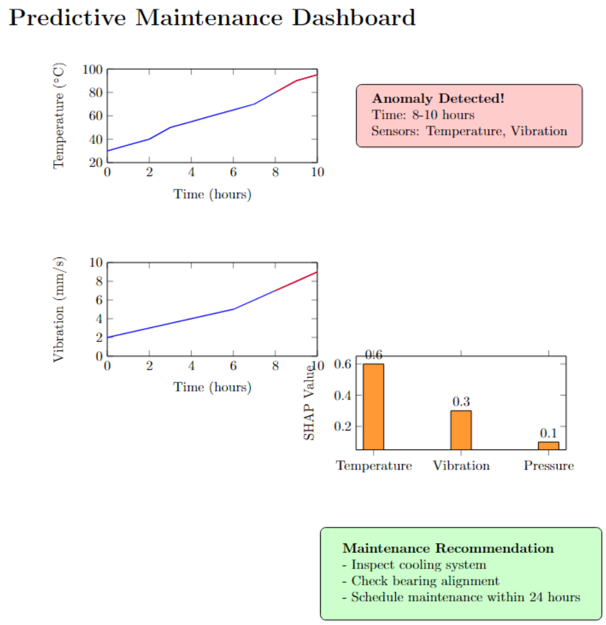

Fig. 2: Dashboard Displaying Anomaly Detection and SHAP Explanations

D. Future Directions

Edge computing could zip up anomaly spotting onsite. Smarter XAI might tailor insights to specific trades. And setting standards could make this the go-to way. There's plenty of room to grow here.

VII. Conclusion

Teaming Just-in-Time maintenance with Explainable AI in IIoT setups could shake up predictive maintenance for the better. Clear, sharp predictions mean smarter calls, less idle time, and slicker workflows. Sure, data and setup hurdles loom, but the gains in clarity and efficiency make it a no-brainer for more digging and doing.

Acknowledgments

The author wishes to thank colleagues and industry partners for their sharp insights and solid data that fueled this survey.

References

[1] L. D. Xu, W. He, and S. Li, "Internet of things in industries: A survey," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2233–2243, 2014.

[2] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.

[3] S. Wachter, B. Mittelstadt, and C. Russell, "Counterfactual explanations without opening the black box: Automated decisions and the gdpr," *Harvard Journal of Law & Technology*, vol. 31, pp. 841–887, 2017.

[4] D. Gunning, "Darpa's explainable artificial intelligence (xai) program," in *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 2019, pp. 2–2.

[5] W. Samek, T. Wiegand, and K.-R. Müller, "Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models," *arXiv preprint arXiv:1708.08296*, 2017.

[6] K. Ashton, "That 'internet of things' thing," *RFID Journal*, vol. 22, no. 7, pp. 97–114, 2009.

[7] J. K. Liker, *The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer*. McGraw-Hill, 2004.

[8] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (xai)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.

[9] A. K. S. Jardine, D. Lin, and D. Banjevic, "A review on machinery diagnostics and prognostics implementing condition-based maintenance," *Mechanical Systems and Signal Processing*, vol. 20, no. 7, pp. 1483–1510, 2006.

[10] X. Li, Q. Ding, and J. Sun, "Remaining useful life estimation in prognostics using deep convolution neural networks," *Reliability Engineering & System Safety*, vol. 172, pp. 1–11, 2018.

[11] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017, pp. 4765–4774.

[12] M. T. Ribeiro, S. Singh, and C. Guestrin, ""why should i trust you?": Explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.

[13] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PloS one*, vol. 10, no. 7, p. e0130140, 2015.

[14] Siemens, "Mindsphere: The cloud-based, open iot operating system," 2021. [Online]. Available: https://www.plm.automation.siemens.com/global/en/products/mindsphere/

[15] GE Digital, "Predix platform," 2021. [Online]. Available: https://www.ge.com/digital/iiot-platform

[16] Hitachi Vantara, "Lumada industrial dataops and industrial insights," 2021. [Online]. Available: https://www.hitachivantara.com/en-us/products/iot.html

[17] IBM, "Ibm watson iot: Accelerate innovation with the internet of things," 2021. [Online]. Available: https://www.ibm.com/internet-of-things

[18] J. A. Jiménez, I. Arribas, Y. Saez, and Á. Herrero, "Explainable artificial intelligence (xai) for the predictive maintenance of industrial machinery," *Applied Sciences*, vol. 10, no. 18, p. 6422, 2020.

[19] A. B. Arrieta *et al.*, "Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai," *Information Fusion*, vol. 58, pp. 82–115, 2020.

[20] L. Antwarg, R. Miller, B. Shapira, and L. Rokach, "Explaining anomalies detected by autoencoders using shap," *Expert Systems with Applications*, vol. 167, p. 114122, 2021.

[21] V. Arya *et al.*, "One explanation does not fit all: A toolkit and taxonomy of ai explainability techniques," *arXiv preprint arXiv:1909.03012*, 2019.

[22] International Organization for Standardization, "Iso/iec 27001:2013 information technology — security techniques — information security management systems — requirements," 2013. [Online]. Available: https://www.iso.org/standard/54534.html

[23] ——, "Iso 31000:2018 risk management — guidelines," 2018. [Online]. Available: https://www.iso.org/standard/65694.html

[24] International Electrotechnical Commission, "Iec 61508:2010 functional safety of electrical/electronic/programmable electronic safety-related systems," 2010. [Online]. Available: https://webstore.iec.ch/publication/5510

[25] A. Bibal and B. Frénay, "Legal requirements on explainability in machine learning," in *Proceedings of the NIPS 2016 Workshop on Interpretable Machine Learning in Complex Systems*, 2016.