Towards Collaborative Predictive Maintenance Leveraging Private Cross-Company Data

Marisa Mohr[1,2], Christian Becker[2], Ralf Möller[1], Matthias Richter[2]

Abstract: The accuracy of a predictive maintenance model is largely determined by the available training data. This puts such machine learning systems out of reach for small and medium-sized production engineering companies, as they are often unable to provide training data in sufficient quality and quantity. Building a collaborative model by pooling training data across many companies would solve this issue, but this data cannot simply be consolidated in a central location while at the same time preserving data integrity and security. This paper enables a collaborative model for predictive maintenance on cross-company data without exposing participants' business information by connecting two recent methodologies: blockchain and federated learning.

Keywords: Industrial Internet of Things; Machine Learning; Blockchain; Federated Learning

1 Introduction

Data-driven products and machine learning (ML) methods offer large benefits for production engineering companies. For some, application of such methods may even be necessary to remain competitive. Intelligent planning of maintenance windows, for example, decreases the risk of unwanted production downtimes and helps to keep machines in their optimal conditions. However, development of such products requires a large initial investment in model definition and training data acquisition. The latter is especially important, as the prediction quality of an ML model is largely determined by the data used for its training. If these data do not contain the patterns preceding a failure, the model will be unable to predict impending machine failures from new, previously unseen data. In predictive maintenance, this issue is made worse by the relatively rare occurrence of machine failures. Of course, one could deliberately allow machines to degrade to capture more failure patterns, but this is at the very least fiscally irresponsible – especially for small and medium-sized enterprises (SMEs).

A more sensible approach would be to share data across multiple SMEs or between machine manufacturer and machine users in order to train a more powerful model than one contributor could do on their own. There are, however, two key challenges with this approach: data integrity and data security. In this work-in-progress paper, we discuss these challenges and combine two recent methods, blockchain and federated learning, to enable collaborative predictive maintenance.

[1] University of Lübeck, Institute of Information Systems, Ratzeburgerallee 160, 23562 Lübeck, Germany, {mohr,moeller}@ifis.uni-luebeck.de
[2] inovex GmbH, Ludwig-Erhard-Allee 6, 76131 Karlsruhe, Germany, {mmohr,cbecker,mrichter}@inovex.de

2 Use Case: Collaborative Predictive Maintenance

Predictive maintenance systems estimate the time until a machine will likely fail, highlight possible problems with complex machines and identify the parts that need to be repaired. There are several ways to model each of these individual tasks, but the underlying data are typically time series, e.g., of sensor readings. Classification or prediction models, classical or based on deep learning, learn patterns and regularities from a large amount of historical data [Mo20, Ma18]. As more data generally leads to better forecasts [HNP09], companies could work together to pool their data and create collaborative models for predictive maintenance. If machine manufacturers, their machine users and also third parties such as service partners are among the participants, value chains can be expanded to create value networks.

Pooling the data is not without issues though. Consider a production plant equipped with a number of sensors. The data from these sensors are passed through several systems such as the controller, gateway or internal servers and finally stored in a database. Within one company, there is little need to protect the data in this chain against inside manipulation: there is no incentive for manipulation. However, there is a reason for doing so when exchanging data with external partners: a malicious actor could corrupt their data in order to sabotage or subtly change the model predictions in their favour [Ba18]. A second challenge lies in transparent documentation of maintenance procedures. In the event of a machine failure, faulty or incomplete maintenance can be used to draw conclusions about the failure, thus proving both warranty claims by the machine manufacturer and the fault of the machine operator. Finally, the third major challenge is the collaborative training of models for predictive maintenance without disclosing business information from the individual participants [BM20a].

3 Methods

This section introduces the two methodologies that contribute to collaborative predictive maintenance with respect to data integrity and data privacy.

3.1 Blockchain: Protection against Forgery, Consensus, and More Concepts

Sharing data in a consortium requires that the data is verifiable and can't be forged. Ideally, data would be signed where it is recorded (i.e., in the sensor) before it is transferred to a downstream system. However, commercially available systems that implement such approaches may be incompatible between different manufacturers. Korb et al. [Ko19] propose to install a microcontroller directly between a sensor and a downstream system instead. By dispensing with proprietary operating systems, there is minimal scope for manipulation and unwanted changes in the processing logic of the data. By implementing a

blockchain on the microcontroller, it can be ensured that the data is forgery-proof using suitable signing methods and cryptographic methods.

A more important aspect is the consensus mechanism offered by a blockchain. In short, consensus mechanisms are protocols that ensure that all participants are synchronized with each other and agree on which transactions are legitimate and included in the blockchain. This procedure enables a transparent maintenance log based on error messages and maintenance entries. The tuple of an entry in the maintenance log, for example, an error message or a performed maintenance activity, is cryptographically hashed and the hash is written to the blockchain. In the maintenance log, the hashes of the respective entry and the hashes of the previous entry, if any, are linked together. This allows complete traceability of the entries and the participants are able to check the authenticity of the entries themselves at any time. This procedure does not prevent unauthorized manipulation of the data, but would make it obvious to the other participants. Ultimately, this mechanism helps to build trust between parties, which benefits everyone in the end.

3.2 Federated Learning: Collaborated Models with Private Data

Still, participants may be reluctant or even unable to share their data with collaborators, as doing so might expose trade secrets or violate data protection regulation. Federated learning (FL) solves this issue. In FL, each participant trains an ML model on their private data and using their own hardware [Mc16, Li20]. These models are then aggregated by a central curator to form a unified global model that has learned from every participants' private data without ever directly accessing it.

The FL algorithm works in several rounds. Initially, the central curator selects a list of the participating machines – the workers – using some candidate selection algorithm, e.g. [Hu20], to ensure that only cooperative participants with suitable data may take part in the training. The current global model is sent to all participating workers. They then use their private data to derive an updated local model and send it back to the curator. Finally, the curator aggregates the local models to an updated global model, e.g., using a weighted mean, where the weights are proportional to the number of training samples used by the workers. This process is repeated until the model is accurate enough, the number of rounds exceeds a threshold, or some other termination criterion is fulfilled.

With FL, the training data is not required to be centralized, but can instead remains with the owners. Nevertheless, private data may still be extracted from the global model as demonstrated by Carlini et al. [Ca18]. This is a necessary by-product of all ML, because an ML model is essentially a compact representation of the training data. This issue can be mitigated using differential privacy techniques [Ca18, Li20], e.g., by adding noise to the local models before sending them to the curator. Of course, this will degrade the overall model performance; how to decide the trade off between model quality and data privacy depends on the participants, the use case, the data, etc. and is out of scope of this paper.

4 Integration

The application of blockchain, federated learning and differential privacy operates on two logical components: While some components are accessible via the Internet, others are only present in private networks. Figure 1 shows an exemplary architecture as to be implemented in the research project "KOSMoS". Note that these are only logical components, not physical or digital system parts.

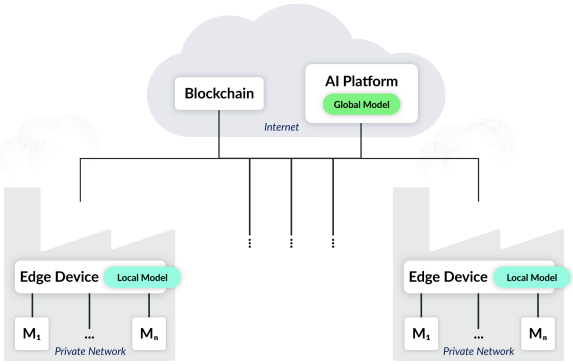

Fig. 1: Exemplary architecture with machines M_i in local networks connected to introduced methods.

The blockchain component of the system serves as a common, trusted distributed database of machine manufacturers and machine users. In the simplest case, the blockchain is centrally operated with only one node. However, there are usually several nodes, where each node contains a complete version of the blockchain. In our case, nodes can be located in the IT of a service provider, a machine manufacturer or even the machine user [BM20b].

In order to realize collaborative predictive maintenance models, the central curator is located in a cloud-based AI platform for prediction models. In the federated learning process, its task is to collect the model updates from all participants, aggregate the new model from them and then distribute the resulting model back to the participants. Each machine operator who wants to take part, needs one edge device locally on their hall floor, which is integrated into the private part of the architecture of KOSMoS. The edge device has access to all machine data required for model training and also serves as a communication interface between the hall floor with a set of $n \in \mathbb{N}$ machines M_i, where $i = 1, ..., n$, and the cloud components. Different data structures of the machines and sensors are translated into topic messages so that the models can be trained with a uniform structure.

5 Conclusion and Future Work

Data-driven business models that are used across company boundaries are important for SMEs, and the challenges discussed here are gaining in importance. The methods presented

promise to ensure data integrity and data security at a high level. In detail, however, there are still many open questions that need to be addressed. The KOSMoS approach will be raised to a usable prototype level. In doing so, we will tackle many of the challenges mentioned here and solve them piece by piece. The result will be a reusable system complemented by a framework that will enable SMEs in particular to train collaborative prediction models with their partner companies. The platform will not only be applicable in production-related SMEs but also will be transferable to other industries and use cases.

Acknowledgment The contents of this publication are taken from the research project "KOSMoS - Collaborative Smart Contracting Platform for Digital Value Networks", funded by the Federal Ministry of Education and Research (BMBF) under reference number 02P17D026 and supervised by Projektträger Karlsruhe (PTKA). The responsibility for the content is with the authors.

Bibliography

[Ba18] Bagdasaryan, E.; Veit, A.; Hua, Y.; Estrin, D.; Shmatikov, V.: How To Backdoor Federated Learning. CoRR, abs/1807.00459, 2018.

[BM20a] Becker, C.; Mohr, M.: Federated Machine Learning: über Unternehmensgrenzen hinaus aus Produktionsdaten lernen. atp magazin, (5):18–20, 2020.

[BM20b] Bux, T.; Mohr, M.: Blockchain Lösungen für den produktionstechnischen Mittelstand. WT WERKSTATTSTECHNIK, BD. 111(Nr. 4):201–204, 2020.

[Ca18] Carlini, N.; Liu, C.; Erlingsson, Ú.; Kos, J.; Song, D.: The Secret Sharer: Evaluating and Testing Unintended Memorization in Neural Networks. ArXiv, abs/1802.08232, 2018.

[HNP09] Halevy, A.; Norvig, P.; Pereira, F.: The Unreasonable Effectiveness of Data. IEEE Intelligent Systems, 24:8–12, 2009.

[Hu20] Huang, H.; Lin, K.; Guo, S.; Zhou, P.; Zheng, Z.: Prophet: Proactive Candidate-Selection for Federated Learning by Predicting the Qualities of Training and Reporting Phases. arXiv:2002.00577 [cs, stat], February 2020.

[Ko19] Korb, T.; Michel, D.; Riedel, O.; Lechler, A.: Securing the Data Flow for Blockchain Technology in a Production Environment. IFAC-PapersOnLine, 52(10):125–130, January 2019.

[Li20] Li, T.; Kumar Sahu, A.; Talwalkar, A.; Smith, V.: Federated Learning: Challenges, Methods, and Future Directions. IEEE Signal Processing Magazine, 37:50–60, 2020.

[Ma18] Mao, W.; He, J.; Tang, J.; Li, Y.: Predicting remaining useful life of rolling bearings based on deep feature representation and long short-term memory neural network. Advances in Mechanical Engineering, 10(12), December 2018.

[Mc16] McMahan, H. B.; Moore, E.; Ramage, D.; Hampson, S.; Agüera y Arcas, B.: Federated Learning of Deep Networks using Model Averaging. CoRR, abs/1602.05629, 2016.

[Mo20] Mohr, M.; Wilhelm, F.; Hartwig, M.; Möller, R.; Keller, K.: New Approaches in Ordinal Pattern Representations for Multivariate Time Series. In: Proceedings of FLAIRS-33. 2020.