


Real-Time Data-Driven Maintenance Logistics: A Public-Private Collaboration

Willem van Jaarsveld¹  

Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands

Laurens Bliet 

Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands

Mathijs de Weerd 

Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, The Netherlands

Stella Kapodistria 

Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands

Verus Pronk

Philips Research Laboratories, Eindhoven, The Netherlands

Peter Verleijdsdonk 


Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands

Simon Voorberg 

Department of Information Systems Supply Chain Management and Decision Analysis, NEOMA Business School, Rouen, France

Alp Akçay 

Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands

Paulo da Costa 

Etsy Inc., Dublin, Ireland

Rik Eshuis 

Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands

Uzay Kaymak 


Jheronimus Academy of Data Science, Eindhoven University of Technology, The Netherlands

Geert-Jan van Houtum 

Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands

Sicco Verwer 

Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, The Netherlands

Yingqian Zhang 

Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, The Netherlands

Abstract

The project “*Real-time data-driven maintenance logistics*” was initiated with the purpose of bringing innovations in data-driven decision making to maintenance logistics, by bringing problem owners in the form of three innovative companies together with researchers at two leading knowledge institutions. This paper reviews innovations in three related areas: How the innovations were inspired by practice, how they materialized, and how the results impact practice.

2012 ACM Subject Classification General and reference → Surveys and overviews

Keywords and phrases Data, Maintenance, Logistics, Optimization, Research, Project

Digital Object Identifier 10.4230/OASICS.Commit2Data.2024.5

Funding This work was supported by the “Netherlands Organisation for Scientific Research” (NWO). Project: NWO Big data - Real Time ICT for Logistics. Number: 628.009.012.

Acknowledgements We want to thank all the involved private partners who have contributed to this project.

¹ Corresponding author



© Willem van Jaarsveld, Alp Akçay, Laurens Bliet, Paulo da Costa, Mathijs de Weerd, Rik Eshuis, Stella Kapodistria, Uzay Kaymak, Verus Pronk, Geert-Jan van Houtum, Peter Verleijdsdonk, Sicco Verwer, Simon Voorberg, and Yingqian Zhang; licensed under Creative Commons License CC-BY 4.0

Commit2Data.

Editors: Boudewijn R. Haverkort, Aldert de Jongste, Pieter van Kuilenburg, and Ruben D. Vromans; Article No. 5; pp. 5:1–5:13



OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Companies in maintenance logistics aspire to make better use of the increasing availability of real-time data from the many (inter)connected devices within the internet of things (IoT). Many such companies have in recent years taken a first key step in this direction by investing substantially in a data management infrastructure that ensures central and real-time availability of the raw data generated by the assets as well as key information on maintenance resources; e.g. the real-time location and status of field service engineers and spare parts, the status of repair centers, the availability of remote service engineers in the call center, etc. Companies are eager to leverage this investment to reduce cost and increase operational asset availability, by transitioning from traditional static maintenance logistics plans based on rigid task intervals to dynamic maintenance logistics policies fueled by real-time data.

Against this background, in 2018 the project *“Real-time data-driven maintenance logistics”* was started, as a collaboration between two leading knowledge institutions in the area of data-driven maintenance logistics (Eindhoven University of Technology and Delft University of Technology) and three companies at the forefront of innovations in maintenance logistics: Philips, Fokker Services (FS), and Dutch Railways (NS). The objective of this project according to its proposal was set in two steps:

1. The dynamic identification of actions from real-time data.
2. Organization of the dynamic execution of these actions by appropriately allocating the resources to accomplish them.

A different way of looking at it, based on the outcome of multiple consortium-wide discussions and following the design of an award-winning poster [14], would be the first step requiring a prediction that could be followed up by a prescription. The second step would design the process for this prescription. Hence, the outcome of our research would result in algorithms that identify the appropriate actions and a semi-structured data-driven process model that dynamically accomplishes such actions.

Big data played an important role, in the following sense. All three companies have large amounts of data available, and seek to make their operations data-driven. However, to achieve this, planning must become more nimble, requiring different planning algorithms and approaches. Hence, the general development of data-driven operations was the key motivation for the project.

To keep the project relevant, in the project execution we took a more flexible approach to identifying research topics. In particular, the project featured meetings with the entire consortium, approximately every 6 months. During those meetings, we reported the results of the previous period, as well as plans for the coming period. More importantly, companies provided feedback on those plans, and contributed ideas on how the various developments linked to challenges within each company. Those ideas enabled researchers to refine and make concrete the research directions set out in the project proposal. This led to various concrete research challenges, that were actively researched by various researchers funded by the project. Amongst others, we researched

1. How to include integer constraints in expensive optimization problems.
2. How to optimize information gathering when decision making.
3. How to learn from data to do maintenance prognostics and routing.

For each of these endeavors, we seek to highlight:

- (i) How the challenge arose from the interaction with the participating companies.
- (ii) How the challenge was formalized into research.

- (iii) How data was used/collected, including adjustments and enrichments.
- (iv) Which algorithms existed and which new algorithms were developed in the project.
- (v) Experiments and results.
- (vi) The learnings and practical validation with the companies.

We discuss these 6 elements for integer constraints in expensive optimization problems in Section 2, for optimized information gathering in Section 3, and for learning maintenance prognostics and routing in Section 4. We reflect and conclude in Section 5.

2 Integer constraints in expensive optimization problems

2.1 Domain

The Train Unit Shunting Problem (TUSP) is a complicated planning problem in railway operations, where trains are moved to a shunting yard to be maintained, cleaned and inspected, see Figure 1. The research group at NS has several algorithmic solutions available for this problem, each with its own set of parameters that need to be tuned correctly. However, tuning these parameters requires not just domain knowledge of railways, but also knowledge of the planning algorithms. Such broad knowledge is difficult to find in a team, let alone in one employee. Automated support for tuning the parameters of planning algorithms is therefore very relevant to achieve top performance in planning problems.

An issue is that tuning the parameters of planning algorithms, such as the algorithms used for TUSP, is very time consuming, and can be regarded as an expensive optimization problem due to the required computational resources. Predicting the performance of certain parameter values in advance would help in the tuning process. This can be done with surrogate-based optimization (SBO) techniques, which are particularly well suited for solving expensive optimization problems [17, Chapter 10]. These SBO techniques use machine learning to predict the performance of parameter values and to find the optimal values.

Traditional SBO techniques assume that parameters can be tuned to any real-valued number between a lower and upper bound. Some parameters in TUSP algorithms, however, are integer-valued, meaning that there is only a finite number of possibilities available for them. We call such a restriction an integer constraint. How to best deal with integer constraints in expensive optimization problems is an open research question.

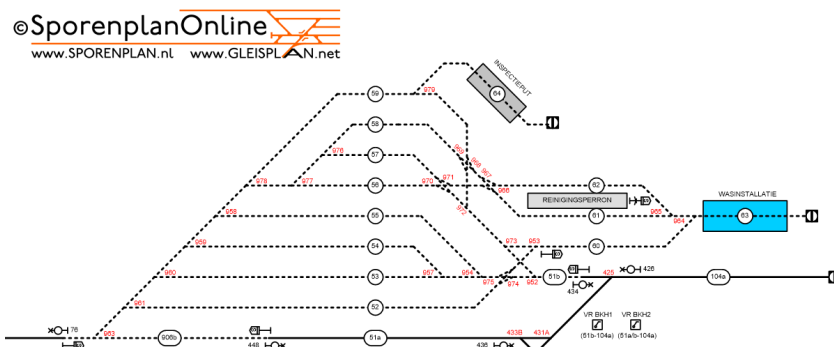


Figure 1 Example of a train shunting yard considered in the TUSP, from www.sporenplan.nl.

2.2 Modeling

Motivated by the problem of parameter tuning of TUSP algorithms, we study expensive optimization problems in general, particularly those with integer constraints. The objective (e.g., algorithm performance) is treated as a black-box and may suffer from noise, for example due to randomness in the TUSP algorithms or disturbances when executing the proposed solution. This way, the problem is reduced to tuning the parameters in such a way that the objective is maximized.

2.3 Data

We consider several simulators and parameter tuning problems that are expensive to run: a robust Traveling Salesman Problem, a wind farm simulator, a pipe shape simulator, an industrial gas filter simulator, a hospital simulator, and hyperparameter tuning for an XGBoost machine learning model. These have been implemented in an open-source software package, EXPOBench [5]. Since gathering data from these simulators is computationally expensive, all data resulting from the project has been made available open-source [7], providing a new big dataset to the community. Additionally, the parameters of several TUSP algorithms have been tuned automatically in MSc projects, namely two TUSP solvers [21, 20] and a TUSP instance generator [16].

2.4 Algorithms and experiments

We compared with the following baseline SBO techniques: Bayesian optimization with Gaussian processes [19], SMAC [15], HyperOpt [2], CoCaBO [18], and DONE [4]. The DONE algorithm has been adapted such that it can deal with integer constraints, leading to two new techniques that were developed during the project: IDONE [8] and MVRSM [6].

Experiments consisted of running different SBO techniques on different simulators and parameter tuning problems. An open-source software package (EXPOBench) was created for this purpose [5], and the resulting dataset was also publicly released.

2.5 Practice

Domain experts can now tune the parameters of their algorithms even when different types of parameters (continuous, integer) are involved, due to the newly developed techniques, and different solutions are available for domain experts to compare. Furthermore, we have obtained valuable insights into the relation between the type of SBO algorithm and the type of parameters in the problem, providing more guidelines on when to use which technique. The first results on algorithms from NS showed the potential for improvement in parameter tuning, but the algorithms need to be made more consistent before they can really benefit from the research. More research into the generalization aspects of the SBO algorithms is also required, to make sure they can operate in a variety of situations. Besides using the newly developed techniques at project partner NS in MSc projects, these techniques have also been applied in other places such as Redeia in Spain, and in other software packages such as fast CMA-ES [26]. The research has been presented at several companies, conferences and workshops.

3 Optimizing information gathering

3.1 Domain

Work in maintenance is often outsourced to specialized suppliers, and the process of finding and safeguarding the availability of such suppliers, which in turn enables those suppliers to engage in capacity management, is information- and knowledge-intensive [24, see also]. In particular, operators of high-tech equipment such as trains and aircraft send out quotations to a diverse range of potential suppliers. Processing such quotations is time consuming. For suppliers, not every quotation is sufficiently interesting to invest the time required to process it.

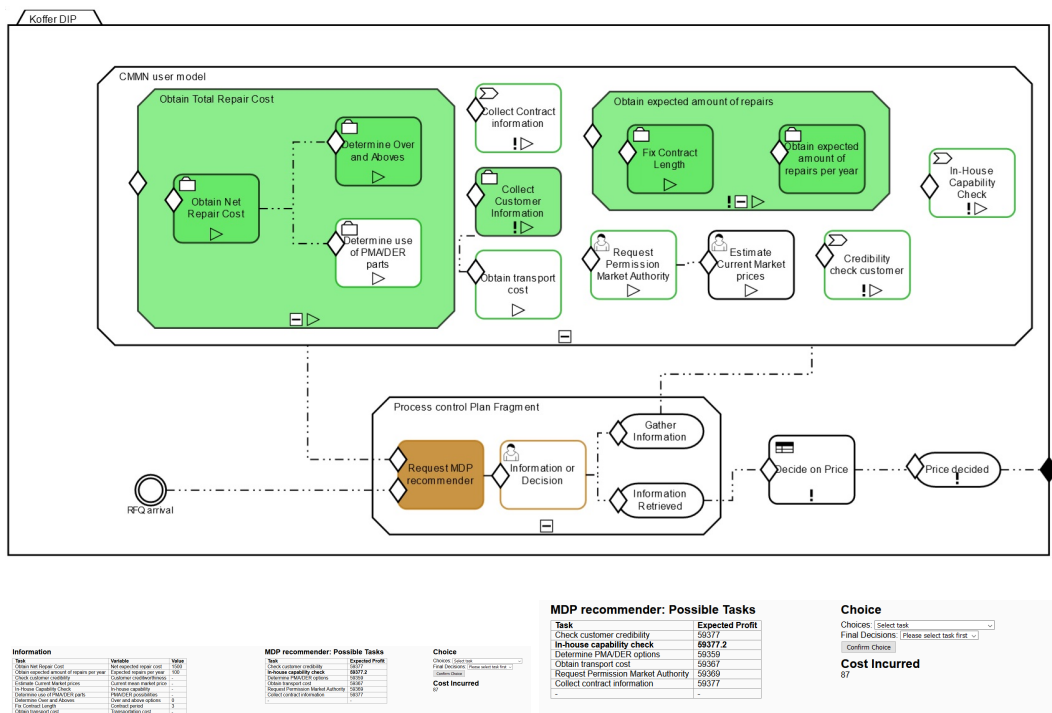
Most time is invested in gathering the information based on which the decision, to quote and what to quote, is made. There is no fixed procedure for this, since every quote is perceived to be unique. This information gathering costs time and money. Therefore, companies such as FS require an approach that helps them guide their employees during the quotation process, by either recommending when to stop since the quote indicates a likely loss for FS, or by recommending which information to keep collecting. Three key and interlocking challenges in the efficient processing of quotations are:

1. How to avoid the collection of information that is not useful.
2. How to ensure that sufficient information is gathered before the deadline for submitting a quote passes.
3. The order in which the information is collected should be according to what is the most valuable information to optimize the quote.

3.2 Modeling

The challenge is approached via a combination of conceptual modeling and optimization. From the conceptual modeling point of view, we adopt an industrial modeling language for dynamic information-centric business processes denoted by CMMN (Case Management Modeling and Notation). Choosing an industrial standard for our approach helps transfer the research into practice. We introduce an approach that can model decision-intensive processes, or more specifically quotation processes, in CMMN. Moreover, we introduce additional inputs such as an information structure which represents the business case that is being developed as part of submitting a quote, i.e. the projected profit as a function of the information retrieved. The result of this part of the model is defined as an Optimizable Decision-Intensive Process (ODIP) consisting of a CMMN representation of the process according to a certain set of constraints that allow for the following approach.

To optimize information retrieval in this ODIP, we develop an approach to convert the CMMN model into a Markov Decision Process (MDP). An MDP is specifically useful in sequential decision processes with uncertain environments, in which the effect of decisions is also uncertain. Given the uncertain outcome of a quotation and the uncertain input of information, an MDP is highly appropriate to model such quotation process optimization problems and more in general ODIPs. Together, the ODIP model and the MDP derived from it, result in a decision support policy that supports the knowledge worker in making the best decision while going through a quotation process. This decision support policy has also resulted in an online demonstration tool in which we show the effectiveness of the complete approach.



■ **Figure 2** A visualization of the demonstration tool that provides decision support to knowledge workers in a quotation process. (OptimizingInformationGathering).

3.3 Algorithms

To optimize the information gathering, we adopt an algorithm for solving the MDP generated from the CMMN process. Using the output of the algorithm, users are recommended to stop the process (since a loss is expected or more information will not result in higher profits) or to continue and collect a specific piece of information. The support is highly flexible: Whenever the user overrides a recommendation, for example because of tacit knowledge based on experience that is not clear from the data, the algorithm adapts and yields new recommendations appropriate for the path adopted by the user. For small size ODIPs, we are able to find the optimal solution since there is a finite horizon to these problems, using backward recursion. For large size problems, we have introduced a deep reinforcement algorithm to still find feasible and good-performing policies.

3.4 Data and experiments

The main inspiration that led us to this model comes from the quotation optimization problem at FS. Based on their definition of the process, their information inputs (which were simplified to keep the true information anonymous and the solution tractable) and their goals, we have constructed a business case that served as a first experiment (see Figure 2). In this figure, one sees a decision process where multiple pieces of information have already been collected (filled green) and where the green outlined cases are information that can be collected if deemed necessary. The bottom right part of the tool gives an insight in the expected profit when collecting a certain type of information. This experiment allowed us to show the feasibility

of the full approach for an existing decision-intensive process. Subsequently, we show the effectiveness and efficiency of our policies in a full-factorial experiment vis-a-vis relevant benchmarks.

3.5 Practice

Using the developed approach, considerable time can be saved when developing quotations. The challenging part is the development of concrete business cases; ideas for this have been developed both at FS and at companies not participating in the project, such as a large logistics service provider in the southeast of the Netherlands. Another issue is scalability. In solving the optimization problem as a MDP, a challenge is that the number of states grows exponentially in the size of the problem. In successive research, we are actively exploring various strategies to tackle this. In one direction, we decompose an ODIP into subproblems based on decision hierarchy. Large decision problems often contain multiple subdecisions that allow us to use this approach. In the other direction, we introduce solution methods that can deal with significantly larger state spaces such as deep reinforcement learning. Such methods increase the applicability of the model to a wider range of problems.

4 Learning maintenance prognostics and routing

4.1 Domain

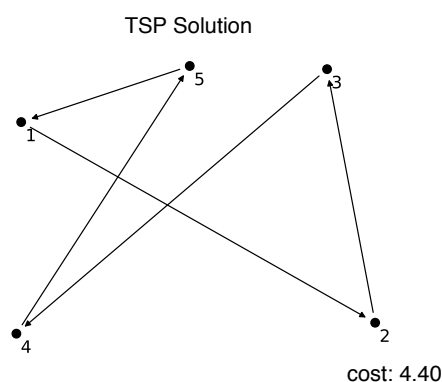
Many prediction and optimization problems arise to obtain policies that maintain a network of industrial assets with minimum maintenance and travel costs. Traditionally, the prediction and optimization problems are decoupled from each other, and prediction uncertainties are incorporated into the decision making models; see [1] for a recent example in the context of maintenance. Furthermore, the methods for solving the prediction and optimization problems rely on the modeling and solution capabilities of decision makers, which can be inaccurate or inefficient for the problems at hand. The motivation of this work is to move from this traditional view to the adoption of data-driven methods based on machine learning leading to accurate and efficient solutions. The scope of our work includes the problems in two distinct areas, prognostics and routing, that come together in the maintenance of a network of industrial assets.

4.1.1 Prognostics

The field of prognostics focuses on predicting the remaining useful life (RUL) of equipment, which is critical for optimizing maintenance schedules and minimizing downtime. One of the primary challenges in prognostics is the difficulty in determining the best time to repair a given asset. In many cases, sensor data from equipment is available; however, parsing this information can be time-consuming and complex due to the high volume and variety of data streams. Furthermore, for newly introduced equipment, run-to-failure data is often unavailable, making it even more challenging to accurately predict the RUL. On the other hand, historical (run-to-failure) data may exist for older assets, providing valuable insights for estimating their RUL. Hence, the development of advanced prognostic algorithms capable of leveraging both sensor data and historical data is crucial for optimizing repair schedules and improving the overall reliability of equipment.

4.1.2 Routing

In the context of equipment maintenance and repair, routing can be formulated as a Traveling Salesman Problem (TSP) where the goal is to determine the most efficient route for visiting a network of assets that require inspection or maintenance. Figure 3 illustrates the solution of a TSP instance. Given the combinatorial nature of the TSP, finding an optimal solution is computationally intractable for large-scale problems. Therefore, we propose adopting a machine learning (ML) perspective to learn improvement heuristics, which can be used to search near-optimal tours. By employing ML techniques such as reinforcement learning, we aim to develop algorithms capable of efficiently exploring the possible solutions and converging to high-quality solutions for the routing problem. This approach not only enables more effective utilization of available resources but also the incorporation of asset prognostics into the routing decision making process.



■ **Figure 3** A Traveling Salesman Problem solution. Indices represent the order in which nodes, representing locations in the network, are visited.

4.2 Modeling

In the field of prognostics, our focus is on developing methods that require minimal human intervention and do not necessitate prior knowledge of the underlying failure mechanisms. The proposed methods are designed to leverage historical run-to-failure data, which is assumed to be available for older equipment, and sensor data for the assets of interest. By modeling this setup as a transfer learning problem, we aim to use the assets with both sensor and run-to-failure data to improve prognostics assets with only sensor data [13]. To effectively capture dependencies in the sensor data, we employ machine learning techniques [12].

In the context of routing, we initially examine classic TSP improvement heuristics as a starting point and baseline for our proposed methods. We hypothesize that a look-ahead policy based on simple operators, can outperform a short-sighted one. Our objective is to learn such policies autonomously, relying solely on executing actions and observing their long-term results. Moreover, we do not have access to an exact solver, as generating optimal tours can become computationally intractable for large-scale problems [9, 10]

4.3 Data and experiments

In the context of prognostics, we utilize simulated data from aircraft turbofan engines operating under various conditions and fault modes. The experimental setup consists of four datasets (i.e., FD001-FD004), each containing different combinations of fault modes

and operating conditions. Figure 4 illustrates two normalized sensor values just before a failure occurs. In each experiment, one dataset is designated as the source domain, and we attempt to learn the RUL of the remaining three datasets (target domains), which have different fault modes and operating conditions. This approach allows us to test the model’s effectiveness in transferring knowledge across operating conditions and fault modes. For each target domain dataset, we assume no access to the observed RUL of the assets and attempt to make predictions based solely on the sensor data from the target domain and observed RUL information from the source domain.




Figure 4 Distribution of normalized sensor values before a failure. Sensor distributions differ for the same assets under different operating conditions and fault modes.

For the routing experiments, we consider simulated data containing TSP instances of various sizes (20, 50, 100). We construct a simulator capable of handling a batch of TSP instances and making decisions corresponding to improvement operators based on 2-edge swaps. Data is collected following policies in the simulator, whereby a sequence of operators creates a history of TSP solutions and selected operators for a batch of TSP instances. We train and evaluate the proposed methods for different TSP sizes, learning policies for a 50-node TSP and evaluating on a 100-node TSP to assess how well the model generalizes to different sizes. We repeat a similar setup for other routing problems such as the Vehicle Routing Problem. Additionally, we test the proposed method on real-world TSPLib instances, which are not seen during the training process, in order to evaluate the model’s generalization to instances with different location distributions.

4.4 Algorithms

In the field of prognostics, we develop a deep transfer learning algorithm which means we apply deep learning techniques in transfer learning algorithms. To evaluate the performance of the proposed algorithm, we compare it with various methods, such as transfer component analysis (TCA), and correlation analysis (CoRaL). Furthermore, we investigate the effectiveness of the proposed algorithm in comparison to methods trained only on the source domain (Source-only), and standardized data approaches.

In the routing context, we develop a deep reinforcement learning algorithm designed to select 2-edge swaps for the TSP, the multiple Traveling Salesman Problem, and the Vehicle Routing Problem. A 2-edge swap means that we swap two parts of the initial route and check if this improves the performance. We compare the performance of this algorithm against several benchmarks, including exact solvers, classic heuristics, and both supervised and

reinforcement learning methods. This extensive comparison allows us to assess the relative strengths and weaknesses of our proposed approach, as well as its potential for application in real-world routing and scheduling problems.

4.5 Practice


In the context of prognostics, domain experts can benefit from our proposed deep transfer learning algorithm when faced with similar business cases, such as the presence of RUL labeled data for older assets and the need to obtain predictions for new assets that lack observed failures. As part of the project, [25] has applied our method in a use case at FS, and investigated how it can improve health state predictions of a specific aircraft component under varying working conditions. The strength of our method lies in its ability to be applied even in the absence of information about the underlying degradation of assets. Through the evaluation of test data, we have demonstrated that our proposed method outperforms other off-the-shelf transfer learning algorithms and yields more accurate predictions than training solely on the source domain.

In the realm of routing problems, our proposed heuristics offer domain experts an effective solution for TSP-like problems. The learned policies outperform other hand-crafted heuristics while operating under a similar computational budget. A notable advantage of our algorithm is its capacity to be applied in batches, enabling the efficient simultaneous solving of multiple instances. Experimental results further reveal that our algorithm can be employed to solve more general TSP instances, and can be adapted to fit specific instance distributions depending on the application, such as drilling problems, cargo delivery, maintenance scheduling and chip design.

Motivated by the success of deep reinforcement learning for solving the TSP efficiently, the approach has been extended for a real-life service-logistics use case of Philips, where the routing of a field service engineer is optimized in the presence of imperfect alerts on the health condition of physical assets [11]. Based on the discussions with Philips, we developed the Dynamic Traveling Maintainer Problem with Alerts (DTMPA) on asset networks. In the DTMPA, we rank the quality of information retrieved from the alert using various information levels. We propose a wide range of heuristics for the DTMPA to cover each information level. The numerical results show that the deep reinforcement learning heuristic outperforms the others, requiring least information. The DTMPA is extended to the case of multiple field service engineers who must coordinate actions to minimize asset unavailability whilst maximizing coverage to anticipate future events in the network [22]; see Figure 5 for a schematic depiction of that case. We employ an iterative deep reinforcement learning algorithm to directly improve sophisticated dispatching strategies by learning e.g. a repositioning strategy from simulated data. A key advantage of this approach is scalability since data collection can be distributed over multiple compute nodes, which enables us to solve industrial-scale problems in a reasonable time. Numerical experiments based on the Dutch hospital network show that the algorithm quickly produces near-optimal policies and the trained policies demonstrate robustness against minor model modifications.

5 Conclusion and reflection

In this paper, we discussed the various research endeavors that were part of the project “*Real-time data-driven maintenance logistics*”. Reflecting on the process, we next list a few observations and learnings we obtained from this project.



■ **Figure 5** A multi-maintainer service-logistics model: Service engineers respond to IoT-generated alerts of various degrees of severity that must be appropriately ranked considering risk, urgency and opportunity.

First of all, working together with companies worked well for this type of project. The companies each bought into the project by promising a monetary contribution. As a consequence, they each had the objective to get something out of the project, and project participants from the companies were aligned on this with senior management. As a consequence, company representatives were interested in the results and actively participated in project meetings. A good example of this is the award-winning poster [14], which was built based on intensive consortium-wide discussions.

The company input steered the direction of the research on various occasions, and in some cases, this led to research that is challenging theoretically, but that also has a high practical value. A key example is the work on transfer learning [13] for predicting remaining useful life via deep domain adaptation. This work was inspired by a comment from one of the companies that they had ample data for older systems, but not so much for new ones, and that it would be useful to be able to use data for the old systems to predict failures for the new systems. This work unlocked a whole stream of literature on precisely that subject, exactly because it is theoretically challenging. Also, these ideas turned out to be applicable in the company setting [25]. Similarly, the use case on efficiently dealing with requests for quotations was a basis for work that is interesting both in practice and in theory [24, 23].

Some work was useful in ways that were not initially foreseen. For example, the surrogate models discussed in Section 2 were not designed with problems in maintenance logistics in mind, but turned out to be really useful in that context. Relatedly, we note that while each of the junior researchers on the project had their own focus, a substantial overlap both in methodological interests and in problem domain enabled them to learn from each other, which strengthened the entire team and enabled us to deliver interesting interdisciplinary research insights, e.g. the comparison of various traditional optimization techniques and deep reinforcement learning by [10], and the first Traveling Salesman Problem competition [3] which similarly bridges disciplines and was organized via a cooperation within the consortium.

If we look back at the two-step objective of this project we can conclude that we have achieved our goals in both steps. In Sections 2 and 4, we have shown some of the achievements related to the first step regarding identifying actions from real-time data. Also, in Sections 3 and 4, we show how actions can be implemented into a decision making process where we aim at dynamically executing them.

6 Guidance for future collaborative projects

Reflecting on our project’s journey, several key strategies stand out as instrumental for the success of similar collaborative ventures. A critical element was ensuring company buy-in, particularly through promised contributions. This commitment significantly boosted their ongoing engagement and helped align their objectives with the project’s goals. Regular, well-structured consortium meetings also played a vital role. These gatherings maintained company interest and facilitated real-time alignment with their evolving needs, proving essential for sustained collaboration.

Additionally, our project’s flexible planning approach, which emphasized broader goals over specific objectives, allowed us to nimbly adapt to changes and unexpected developments within the companies. This flexibility was key in marrying theoretical exploration with practical applicability, ensuring the project’s relevance and responsiveness.

These insights offer valuable lessons for future research initiatives aiming to bridge the gap between academia and industry. Strategic engagement, consistent communication, and adaptable planning emerge as fundamental components for fostering effective partnerships and achieving meaningful outcomes in such collaborative settings.

References

- 1 Alp Akcay. An alert-assisted inspection policy for a production process with imperfect condition signals. *European Journal of Operational Research*, 298(2):510–525, 2022. doi:10.1016/J.EJOR.2021.05.051.
- 2 James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 115–123. PMLR, 17–19 June 2013. URL: <http://proceedings.mlr.press/v28/bergstra13.html>.
- 3 L Bliiek, P da Costa, R Refaei Afshar, Y Zhang, T Catshoek, D Vos, S Verwer, and F Schmitt-Ulms. The first AI4TSP competition. <https://paulorocosta.gitbook.io/ai4tsp-competition/>.
- 4 L. Bliiek, H. R. G. W. Verstraete, M. Verhaegen, and S. Wahls. Online optimization with costly and noisy measurements using random Fourier expansions. *IEEE Transactions on Neural Networks and Learning Systems*, 29(1):167–182, 2018. doi:10.1109/TNNLS.2016.2615134.
- 5 Laurens Bliiek, Arthur Guijt, Rickard Karlsson, Sicco Verwer, and Mathijs De Weerd. EXPObench: Benchmarking surrogate-based optimisation algorithms on expensive black-box functions. *arXiv preprint arXiv:2106.04618*, 2021. arXiv:2106.04618.
- 6 Laurens Bliiek, Arthur Guijt, Sicco Verwer, and Mathijs De Weerd. Black-box mixed-variable optimisation using a surrogate model that satisfies integer constraints. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, GECCO ’21, pages 1851–1859, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3449726.3463136.
- 7 Laurens Bliiek, Arthur Guijt, Sicco Verwer, Mathijs De Weerd, and Rickard Karlsson. Raw data of the EXPensive Optimization benchmark library (EXPObench). Available at <https://doi.org/10.4121/14247179.v2>, 2021.
- 8 Laurens Bliiek, Sicco Verwer, and Mathijs De Weerd. Black-box combinatorial optimization using models with integer-valued minima. *Annals of Mathematics and Artificial Intelligence*, 89:639–653, 2021. doi:10.1007/s10472-020-09712-4.
- 9 Paulo da Costa, Jason Rhuggenaath, Yingqian Zhang, and Alp Akcay. Learning 2-opt heuristics for the traveling salesman problem via deep reinforcement learning. In *Asian Conference on Machine Learning*, pages 465–480. PMLR, 2020.

- 10 Paulo da Costa, Jason Rhuggenaath, Yingqian Zhang, Alp Akcay, and Uzay Kaymak. Learning 2-opt heuristics for routing problems via deep reinforcement learning. *SN Computer Science*, 2:1–16, 2021.
- 11 Paulo da Costa, Peter Verleijdsdonk, Simon Voorberg, Alp Akcay, Stella Kapodistria, Willem Van Jaarsveld, and Yingqian Zhang. Policies for the dynamic traveling maintainer problem with alerts. *European Journal of Operational Research*, 305(3):1141–1152, 2023. doi:10.1016/J.EJOR.2022.06.044.
- 12 Paulo Roberto De Oliveira da Costa, Alp Akcay, Yingqian Zhang, and Uzay Kaymak. Attention and long short-term memory network for remaining useful lifetime predictions of turbofan engine degradation. *International Journal of Prognostics and Health Management*, 10(4), 2019.
- 13 Paulo Roberto de Oliveira da Costa, Alp Akçay, Yingqian Zhang, and Uzay Kaymak. Remaining useful lifetime prediction via deep domain adaptation. *Reliability Engineering & System Safety*, 195:106682, 2020. doi:10.1016/J.RESS.2019.106682.
- 14 Paulo De Oliveira Da Costa et al. Real-time data-driven maintenance logistics. 3rd Place Commit2Data poster competition at ICT.OPEN, March 2019.
- 15 Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*, pages 507–523. Springer, 2011. doi:10.1007/978-3-642-25566-3_40.
- 16 R. Martens. Controlling TUSP features during instance generation via Bayesian optimization. Eindhoven University of Technology, Master’s thesis, 2022.
- 17 Joaquim R. R. A. Martins and Andrew Ning. *Engineering Design Optimization*. Cambridge University Press, 2021. doi:10.1017/9781108980647.
- 18 Binxin Ru, Ahsan Alvi, Vu Nguyen, Michael A. Osborne, and Stephen Roberts. Bayesian optimisation over multiple continuous and categorical inputs. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8276–8285. PMLR, 13–18 July 2020. URL: <http://proceedings.mlr.press/v119/ru20a.html>.
- 19 B. Shahriari, Kevin Swersky, Ziyu Wang, R. Adams, and N. D. Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016. doi:10.1109/JPROC.2015.2494218.
- 20 L. J. Van den Nieuwelaar. Automatic algorithm configuration with search heuristics for the train unit shunting problem. Eindhoven University of Technology, Master’s thesis, 2021.
- 21 L. Van der Knaap. Contextual hyperparameter optimization for the train unit shunting problem. Delft University of Technology, Master’s thesis, 2021.
- 22 Peter Verleijdsdonk, Willem van Jaarsveld, and Stella Kapodistria. Scalable policies for the dynamic traveling multi-maintainer problem with alerts. *Eur. J. Oper. Res.*, 319:121–134, 2024. doi:10.1016/J.EJOR.2024.05.049.
- 23 S Voorberg, W Van Jaarsveld, R Eshuis, and GJ Van Houtum. Information acquisition for service contract quotations made by repair shops. *European Journal of Operational Research*, 305(3):1166–1177, 2023. doi:10.1016/J.EJOR.2022.06.048.
- 24 Simon Voorberg, Rik Eshuis, Willem Van Jaarsveld, and GJ Van Houtum. Decisions for information or information for decisions? Optimizing information gathering in decision-intensive processes. *Decision Support Systems*, 151:113632, 2021. doi:10.1016/J.DSS.2021.113632.
- 25 H.G.P. Wismans. Domain adaptation for prognostics in the aerospace industry. Eindhoven University of Technology, Master’s thesis, 2019.
- 26 Dietmar Wolz. FCMAES - a Python-3 derivative-free optimization library. Available at <https://github.com/dietmarwo/fast-cma-es>, 2022.