Optimizing maintenance by learning individual treatment effects

Toon Vanderschueren [1 2] **Robert Boute** [1 3] **Tim Verdonck** [2 4] **Bart Baesens** [1 5] **Wouter Verbeke** [1]

Abstract

The goal in maintenance is to avoid machine failures and overhauls, while simultaneously minimizing the cost of preventive maintenance. Maintenance policies aim to optimally schedule maintenance by modeling the effect of preventive maintenance on machine failures and overhauls. Existing work assumes the effect of preventive maintenance is (1) deterministic or governed by a known probability distribution, and (2) machine-independent. Conversely, this work proposes to relax both assumptions by learning the effect of maintenance conditional on a machine's characteristics from observational data on similar machines using existing methodologies for causal inference. This way, we can estimate the number of overhauls and failures for different levels of maintenance and, consequently, optimize the preventive maintenance frequency. We validate our proposed approach using real-life data on more than 4,000 maintenance contracts from an industrial partner. Empirical results show that our novel, causal approach accurately predicts the maintenance effect and results in individualized maintenance schedules that are more accurate and cost-effective than supervised or non-individualized approaches.

1. Introduction

The goal in maintenance is to avoid machine failures and overhauls, while simultaneously minimizing the cost of preventive maintenance (PM). Maintenance is an important operational problem, estimated to represent 15% to 40% of production costs (Dunn, 1987; Löfsten, 2000). Maintenance policies aim to optimally schedule maintenance by modeling the effect of preventive maintenance on machine failures and overhauls. In practice, maintenance has an imperfect effect and does not make the machine as good as new. A broad spectrum of maintenance effects have been studied in the literature, ranging from perfect maintenance, making the system as good as new, to worst maintenance, where maintenance causes failure (Pham & Wang, 1996).

Existing approaches in imperfect maintenance rely on strong assumptions to model the effect of PM (Alaswad & Xiang, 2017): (1) the effect is assumed to be deterministic or stochastic assuming a certain probability distribution, and (2) the effect is assumed to be machine-independent, i.e., identical for all machines. This work proposes a general, data-driven maintenance policy that relaxes both assumptions and learns the effect of maintenance conditional on a machine's characteristics. The benefit of our approach is that it allows (1) to flexibly learn maintenance effects from observational data that is biased due to an existing maintenance policy, and (2) to design a machine-specific PM schedule based on these learned effects.

We contribute by proposing a novel prescriptive framework for maintenance that prescribes the maintenance frequency based on the estimated effect of PM on the machine's number of overhauls and failures. Instead of assuming a model of the PM effect, we frame maintenance as a problem of causal inference. To this aim, we leverage state-of-the-art machine learning methods for causal inference that learn models to estimate a machine's potential outcomes for different PM frequencies from biased, observational data. Moreover, we formulate a prescriptive policy that uses the potential outcomes to decide on the optimal PM frequency so as to minimize the total cost of failures and interventions. Empirically, we contribute by demonstrating the use of the presented prescriptive framework on a dataset consisting of more than 4,000 maintenance contracts of industrial equipment provided by an industrial partner.

2. Problem overview

This work aims to solve the problem faced by a provider of full-service maintenance contracts, where a client's as-

[*]Equal contribution [1]Decision Sciences and Information Management, Faculty of Economics and Business, KU Leuven, Leuven, Belgium [2]Applied Mathematics, Department of Mathematics, University of Antwerp, Antwerp, Belgium [3]Technology and Operations Management Area, Vlerick Business School, Leuven, Belgium [4]Research Centre for Operations Management, Faculty of Economics and Business, KU Leuven, Leuven, Belgium [5]Department of Decision Analytics and Risk, Southampton Business School, Southampton, United Kingdom. Correspondence to: Toon Vanderschueren <toon.vanderschueren@kuleuven.be>.

Published at the ICML 2022 Workshop on Spurious Correlations, Invariance, and Stability. Baltimore, Maryland, USA.

set is maintained at a predetermined price (Deprez et al., 2021). The provider's goal is to minimize the total contract cost. To this aim, the service provider needs to decide on each contract's (usage-based) PM frequency, prior to contract start, based on contract information. PM aims to prevent two types of events: *overhauls* (unplanned, comprehensive maintenance interventions where large parts of the machinery need to be replaced) and *machine failures* (resulting in downtime and corrective maintenance).

Let each contract be a tuple $(\mathbf{X}, T, O, F)$. $\mathbf{X} \in \mathcal{X} \subset \mathbb{R}^d$ denotes the characteristics of the machine and contract. The treatment $T \in \mathcal{T} \subset \mathbb{R}^+$ is the PM frequency: the number of preventive maintenance interventions that will be applied per running period (e.g., maintaining every 200 running hours). The contract's number of overhauls and failures per running period are denoted as $O \in \mathcal{O} \subset \mathbb{R}^+$ and $F \in \mathcal{F} \subset \mathbb{R}^+$. We adopt the potential outcomes framework (Rubin, 2004; 2005) and denote the overhauls O and failures F per running period for maintenance frequency t as $O(t)$ and $F(t)$.

The objective is to find the optimal PM frequency that minimizes the total cost per running period. We assume a usage-based maintenance cost similar to (Faccio et al., 2014). A machine i's cost per running period given PM frequency consists of the combined costs of PM, overhauls and failures:

$$c_i = \underbrace{c_t\, t_i}_{\text{PM}} + \underbrace{c_o\, o_i}_{\text{Overhauls}} + \underbrace{c_f\, f_i}_{\text{Failures}} . \qquad (1)$$

We assume that the costs of PM, overhauls and failures are deterministic and known: $c_t, c_o, c_f \in \mathbb{R}^+$.

To assist the provider's decision-making, an observational data set is available with information on n past contracts: $\mathcal{D} = \{(\mathbf{x}_i, t_i, o_i, f_i)\}_{i=1}^n$. Past maintenance decisions were made according to an unknown existing policy, resulting in selection bias. Because of this, the challenge of learning a predictive model for obtaining unbiased estimates of the potential outcomes from $\mathcal{D}$ is to adjust for selection bias.

3. Methodology

Our methodology consists of a predict-then-optimize framework, see Figure A2 for an overview. To estimate each contract's cost for different PM frequencies $c_i(t_i)$, its potential outcomes need to be estimated, i.e., its number of overhauls $o_i(t)$ and failures $f_i(t)$ for a PM frequency t_i, given characteristics $\mathbf{x}_i$. The first step is to learn a machine learning model for *predicting* potential outcomes from historical, observational data on similar full-service contracts $\mathcal{D}$. In the second step, these estimated outcomes are used to *optimize* the PM frequency and resulting total cost.

3.1. Assumptions

The challenge in estimating potential outcomes from observational data is dealing with selection bias. This requires two standard assumptions to learn unbiased estimates of the potential outcomes (Imbens, 2000; Bica et al., 2020):

Assumption 3.1. Consistency. $Y = Y(t)$ for all $t \in T$.

Assumption 3.2. Overlap. For all $\mathbf{x} \in \mathcal{X}$ with $p(\mathbf{x} > 0)$ and $t \in \mathcal{T}$: $0 < p(t|\mathbf{x}) < 1$.

Assumption 3.3. Unconfoundedness. Potential outcomes $O(t)$ and $F(t)$ are independent of the PM frequency T conditional on machine characteristics X: $\{O(t), F(t) \mid \forall\, t \in \mathcal{T}\} \perp\!\!\!\perp T|\mathbf{X}$.

3.2. Predicting preventive maintenance effects

First, we need to predict each machine's potential outcomes $o_i(t)$ and $f_i(t)$ for PM frequency t_i given its characteristics $\mathbf{x}_i$. Therefore, we aim to find models $g_o : \mathcal{X} \times \mathcal{T} \to \mathcal{O}$ and $g_f : \mathcal{X} \times \mathcal{T} \to \mathcal{F}$ defined by $\theta_o, \theta_g \in \Theta$. The goal is to obtain unbiased estimates of the potential outcomes:

$$g_o(t, \mathbf{x}) = \mathbb{E}\left[O(t)|\mathbf{X} = \mathbf{x}\right], \qquad (2)$$

$$g_f(t, \mathbf{x}) = \mathbb{E}\left[F(t)|\mathbf{X} = \mathbf{x}\right]. \qquad (3)$$

We learn g_o and g_f using SCIGAN, a state-of-the-art machine learning approach for predicting potential outcomes given a continuously-valued treatments (Bica et al., 2020). SCIGAN learns g in two steps. First, a generative adversarial network (GAN) is trained to model the distribution of the potential outcomes: the generator is trained to generate counterfactual contracts that cannot be distinguished from factual, observed contracts by the discriminator. In a second phase, the GAN is used to augment the observed training data with generated counterfactual samples. This way, the augmented data set contains all potential outcomes, including both the factual outcomes and the generated, counterfactual outcomes. Because of this, selection bias is no longer a problem and, using this augmented data set, a predictive model g_θ can be trained to predict the potential outcomes in a supervised manner. For this, we use a neural network.

3.3. Optimizing the maintenance cost

The optimal PM frequency is a trade-off between costs resulting from planned PM on the one hand and costs resulting from overhauls and failures on the other hand. However, using the potential outcomes $o_i(t_i)$ and $f_i(t_i)$, it can be seen that the overhauls and failures can be written as functions of the PM frequency t_i. Therefore, the predicted potential outcomes can be used to directly estimate the costs incurred at different PM frequencies. This is achieved by rewriting all terms in Equation (1) (PM, overhauls and failures) as a function of the PM frequency t_i:

$$c_i(t_i) = c_t\, t_i + c_o\, o_i(t_i) + c_f\, f_i(t_i). \tag{4}$$

Each machine's optimal PM frequency t_i^* is found as the level that minimizes the expected cost: $t_i^* = \arg\min c_i(t)$. To account for heterogeneity in the contracts, this optimal PM frequency is optimized for each specific machine.

4. Results

We validate our methodology empirically using real-world data on full-service maintenance contracts. This section describes our semi-synthetic setup and the choice of evaluation metrics. Our data set contains more than 4,000 full-service maintenance contracts. For each contract i, we have information on the machine, the contract and the maintenance that was performed. Detailed information on this data is provided in the Appendix A2 (see Table A1).

4.1. Semi-synthetic setup

A good estimator should accurately predict both the observed outcome, the number of failures that did occur at maintenance frequency t_i, as well as the unobserved outcomes, the number of failures if the machine had received less or more maintenance. In practice however, not all potential outcomes are observed, which makes evaluation of causal models hard. Because of this, we rely on semi-synthetic data to evaluate our model. This approach is commonly used in both causal inference (see e.g., Berrevoets et al., 2020) and maintenance (e.g., Deprez et al., 2021).

Potential outcomes $o_i(t)$ and $f_i(t)$ are generated based on the observed characteristics $\mathbf{x}_i$. For the overhauls, we have:

$$o_i(t) = 7\,\sigma\left(\underbrace{\mathbf{v}_o^\intercal \mathbf{x}_i}_{\text{Base rate}} - \underbrace{\frac{1}{10}\,\sigma\left(\mathbf{w}_o^\intercal \mathbf{x}_i\right)t}_{\text{PM effect}} + \underbrace{\epsilon_o}_{\text{Noise}}\right) \tag{5}$$

where $\mathbf{v}_o, \mathbf{w}_o \sim \mathcal{U}\left((0,1)^{d\times 1}\right)$ and $\epsilon_o \sim \mathcal{N}(0,1)$. The 7 rescales the average number of overhauls to roughly same number in the original data. For failures, we similarly have:

$$f_i(t) = 9\,\sigma\left(\mathbf{v}_f^\intercal \mathbf{x}_i - \frac{1}{10}\,\sigma\left(\mathbf{w}_f^\intercal \mathbf{x}_i\right)t + \epsilon_f\right) \tag{6}$$

with $\mathbf{v}_f, \mathbf{w}_f \sim \mathcal{U}\left((0,1)^{d\times 1}\right)$ and $\epsilon_f \sim \mathcal{N}(0,1)$.

Using the semi-synthetic setup, the test set contains the potential outcomes for all possible values of $t_i \in \mathcal{T}$ using these equations. Conversely, the training and validation sets include only one observed outcome for one observed t_i. The training, validation and test sets respectively consist out of 50%, 25% and 25% of the data. Hyperparameter optimization is based on the mean squared error on the observed outcomes in the validation set. An illustration of a generated data set is shown in Figure A3.

On the one hand, we evaluate the prescribed maintenance frequencies using the maintenance frequency t_i that was observed in practice for the observed outcomes in the training and validation set. On the other hand, we want to evaluate our policy for different levels of selection bias. For this, we control the level of selection bias in the semi-synthetic data using an approach similar to (Bica et al., 2020). Selection bias is simulated by assigning PM frequencies from a beta distribution as follows:

$$t_i \sim 20\,\text{Beta}\left(1 + \frac{\lambda\delta_i}{10}, 1 + \lambda\delta_i\right) \tag{7}$$

where $\delta_i = \sigma(\mathbf{w}_b \mathbf{x}_i)$ with $\mathbf{w}_b \sim \mathcal{U}\left((0,1)^{d\times 1}\right)$. δ_i ensure that treatment assignment is based on observed features $\mathbf{x}_i$. This way, λ controls the level of selection bias. $\lambda = 0$ results in $\text{Beta}(1,1)$ or the uniform distribution, which implies random maintenance assignment. Higher values of λ imply more selection bias with $\lambda = 30$ resulting in a maintenance distribution similar to the observed distribution. An illustration of the observed distribution and generated distributions for different values of λ is shown in Figure A4.

4.2. Evaluation

Evaluation is done using three different metrics. First, we evaluate the ability of the machine learning model to accurately predict a contract's potential outcomes. This is measured using the mean integrated square error (MISE) (Silva, 2016; Schwab et al., 2020):

$$\text{MISE} = \frac{1}{n}\sum_{i=1}^{n}\int_0^m \left(y_i(t) - \hat{y}_i(t)\right)^2\, \mathrm{d}t. \tag{8}$$

Second, we want to evaluate the accuracy of the prescribed maintenance frequencies. To this end, we consider a variant of the policy error (PE) (Schwab et al., 2020) that compares the prescribed maintenance frequency with the ideal level:

$$\text{PE} = \frac{1}{n}\sum_{i=1}^{n}\left(t_i^* - \hat{t}_i^*\right)^2. \tag{9}$$

Third, we evaluate the prescribed maintenance frequency in terms of costs using the policy cost ratio (PCR) that compares the costs of the estimated optimal maintenance frequency with the ideal level:

$$\text{PCR} = \frac{1}{n}\sum_{i=1}^{n}\frac{c_i(\hat{t}_i^*)}{c_i(t_i^*)}. \tag{10}$$

For all metrics, a lower value is better with 0 being the optimal value for MISE and PE and 1 for PCR.

	MISE	
	Overhauls	Failures
SCIGAN	$\mathbf{7.71 \pm 0.60}$	$\mathbf{14.16 \pm 1.68}$
MLP	10.25 ± 1.33	18.27 ± 3.65

	PE	PCR
SCIGAN–ITE	$\mathbf{2.40 \pm 0.46}$	$\mathbf{1.07 \pm 0.01}$
MLP–ITE	4.36 ± 1.25	1.11 ± 0.02
SCIGAN–ATE	8.77 ± 1.07	1.24 ± 0.04

Table 1: **Empirical evaluation.** We compare performance for the different policies over five runs. We evaluate each model's ability to predict the potential outcomes $o_i(t)$ and $f_i(t)$ (MISE), as well as each policy's ability to accurately prescribe the maintenance frequency (PE) and minimize costs (PCR). For all metrics, a lower value is better.

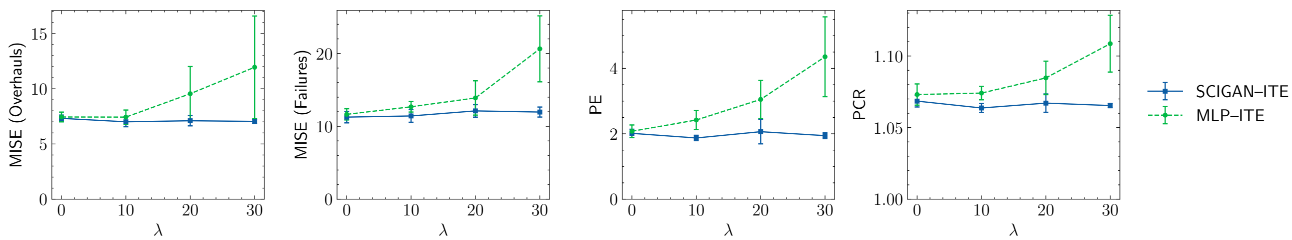

Figure 1: **Results for varying levels of selection bias.** Selection bias is controlled by λ (Equation (7)). Although SCIGAN performs similar to the MLP for lower values of λ, it performs better when the data is biased. Conversely, SCIGAN–ITE performs similar both when assignments are randomized ($\lambda = 0$) or biased at levels similar to the observed data ($\lambda = 30$).

Our proposed maintenancy policy uses SCIGAN to learn the individual treatment effects (ITE) and will be referred to as SCIGAN–ITE. This policy is benchmarked against two other policies (see Table A2). MLP–ITE is a policy based on a neural network (MLP) that learns o_i and f_i given $\mathbf{x}_i$ and t_i in a completely supervised manner without adjusting for selection bias. This allows us to assess whether there is a benefit of adjusting for selection bias. SCIGAN–ATE is a generalized policy setting a single optimal t^* for all contracts based on the average (instead of the individual) maintenance effect. This allows to assess the potential benefit of an individualized policy customized for each machine.

4.3. Empirical results

In this section, we present the results of the semi-synthetic experiments based on more than 4,000 maintenance contracts, as put forward in Sections 4.1 and 4.2. The goal is to answer two research questions. (1) Does an individualized approach outperform a general approach? (2) Does a causal, prescriptive approach outperform a supervised, predictive approach? These are examined for both the observed PM frequencies as well as under varying levels of selection bias.

Observed PM frequencies We present the results for the different methodologies given the maintenance frequency t_i that was observed in practice in Table 1 and Figure A5. For both failures and overhauls, SCIGAN more accurately predicts the potential outcomes in terms of MISE compared to MLP, the supervised approach. Moreover, the individualized, prescriptive approach (SCIGAN–ITE) most accurately prescribes the optimal PM frequency compared to the supervised (MLP–ITE) and non-individualized approach in terms of policy error. Finally, SCIGAN–ITE also results in lower costs compared to MLP–ITE and SCIGAN–ATE. The improved performance of SCIGAN–ITE compared to MLP–ITE illustrates the importance of adjusting for selection bias when learning from observational data. Moreover, the relatively worse performance of the average approach, SCIGAN–ATE, indicates the benefit of an individualized, machine-dependent policy for imperfect maintenance that takes into account machine characteristics.

Different levels of selection bias We compare SCIGAN–ITE's and MLP–ITE's performance for different levels of selection bias using Equation (7) in Figure 1. SCIGAN achieves similar predictive performance, in terms of MISE, and quality of decision-making, in terms of PE and PCR, for the entire range of operating conditions ranging from randomized PM assignments ($\lambda = 0$) to realistic levels of selection bias ($\lambda = 30$). The supervised approach, MLP, results in accurate predictions and decisions when preventive maintenance is randomized, but deteriorates when λ increases. This illustrates that, even when conditioning on confounders, not adjusting for selection bias results in worse performance when data is limited (Alaa & Schaar, 2018).

5. Conclusion

This work proposes a novel, generally applicable prescriptive maintenance policy that models maintenance by learning a maintenance effect conditional on the machine's characteristics from observational data. This is achieved by relying on state-of-the-art machine learning methodologies for causal inference. The benefit of our approach is that, unlike existing approaches, our methodology does need strong assumptions regarding the maintenance effect, but is instead able to learn it from observational data using flexible machine learning models. We validate our approach with semi-synthetic experiments using real-life data on more than 4,000 maintenance contracts. Our proposed approach outperforms supervised and non-individualized approaches in terms of both accuracy and cost of the prescribed maintenance schedules. These findings show that our proposed approach offers a powerful and flexible policy for individualized maintenance and highlight the importance of dealing with selection bias when learning from observational data.

Acknowledgements

We would like to thank the anonymous reviewers for the useful comments & feedback. This work was supported by the BNP Paribas Fortis Chair in Fraud Analytics, FWO research project G015020N, and FWO PhD Fellowship 11I7322N. The computational resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation – Flanders (FWO) and the Flemish Government – department EWI.

References

Alaa, A. and Schaar, M. Limits of estimating heterogeneous treatment effects: Guidelines for practical algorithm design. In *International Conference on Machine Learning*, pp. 129–138. PMLR, 2018.

Alaswad, S. and Xiang, Y. A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability engineering & system safety*, 157:54–63, 2017.

Berrevoets, J., Jordon, J., Bica, I., van der Schaar, M., et al. Organite: Optimal transplant donor organ offering using an individual treatment effect. *Advances in Neural Information Processing Systems*, 33, 2020.

Bica, I., Jordon, J., and van der Schaar, M. Estimating the effects of continuous-valued interventions using generative adversarial networks. *Advances in Neural Information Processing Systems*, 33:16434–16445, 2020.

Deprez, L., Antonio, K., and Boute, R. Pricing service maintenance contracts using predictive analytics. *European Journal of Operational Research*, 290(2):530–545, 2021.

Dunn, R. Advanced maintenance technologies. *Plant Engineering*, 41(12):80–93, 1987.

Faccio, M., Persona, A., Sgarbossa, F., and Zanin, G. Industrial maintenance policy development: A quantitative framework. *International Journal of Production Economics*, 147:85–93, 2014.

Imbens, G. W. The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710, 2000.

Löfsten, H. Measuring maintenance performance–in search for a maintenance productivity index. *International Journal of Production Economics*, 63(1):47–58, 2000.

Pham, H. and Wang, H. Imperfect maintenance. *European Journal of Operational Research*, 94(3):425–438, 1996.

Rubin, D. B. Direct and indirect causal effects via potential outcomes. *Scandinavian Journal of Statistics*, 31(2):161–170, 2004.

Rubin, D. B. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.

Schwab, P., Linhardt, L., Bauer, S., Buhmann, J. M., and Karlen, W. Learning counterfactual representations for estimating individual dose-response curves. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 5612–5619, 2020.

Silva, R. Observational-interventional priors for dose-response learning. *Advances in Neural Information Processing Systems*, 29, 2016.

A1. Problem formulation

The assumed causal structure of the problem is shown in Figure A1. Machine and contract characteristics X affect the outcomes O and F both directly as well as indirectly through T.

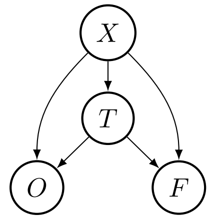

Figure A1: **Causal diagram depicting the relations between the different variables.** X: Machine and contract characteristics, T: Preventive maintenance, O: Overhauls, and F: Failures.

A2. Data

An overview of the data is shown in Table A1. All events are presented per running period, which is a set number of running hours. For reasons of confidentiality, the exact number of running hours per period is not presented. Costs are averaged over all events and re-scaled for reasons of confidentiality. Moreover, the data is preprocessed as follows. Categorical variables are encoded with dummies and $\mathbf{x}_i$ is standardized. The PM frequency, overhauls and failures that occurred throughout the contract are converted to the number of events per running period. Even though a contract's exact number of running hours is not known in advance, an estimate is typically available.

Variable	**Domain**
Machine information	
Type	$\{1, \ldots, 7\}$
Age at contract start	$[0, 39]$
Running hours at contract start	$[2500, 110000]$
Running hours during contract	$[0, 186000]$
Average running hours per year	$[300, 8500]$
Contract information	
Type	$\{1, 2\}$
Duration (days)	$[180, 5850]$
Preventive maintenance per running period	
PM frequency	$[0, 20]$
Outcomes per running period	
Number of overhauls	$[0, 128]$
Number of failures	$[0, 185]$
Average costs (in €)	
Preventive maintenance	73
Overhaul	207
Failure	104

Table A1: **Data overview.** Overview of the available contract information: machine and contract characteristics, preventive maintenance interventions, overhauls, and failures.

A3. Methodology

We show a high-level overview of our predict-then-optimize methodology in Figure A2.

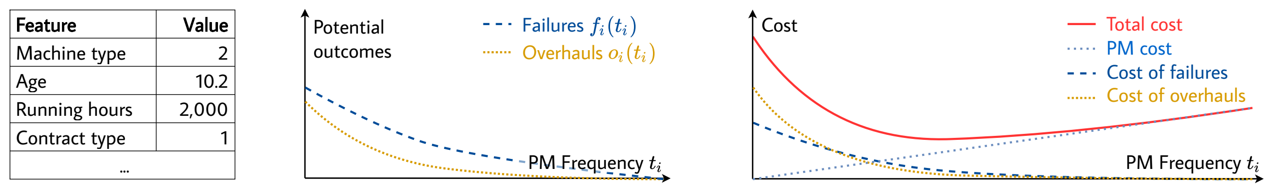

1. Machine information $\mathbf{x}_i$ 2. Predict potential outcomes $o_i(t_i)$ and $f_i(t_i)$ 3. Prescribe the PM frequency t_i to minimize the total cost

Figure A2: **Methodology overview.** We present a high-level overview of our methodology. Contract information $\mathbf{x}_i$ is used to predict the potential outcomes in terms of overhauls $o_i(t)$ and failures $f_i(t)$. Based on these estimates, the total cost for different levels of PM can then be estimated. Finally, the PM frequency is chosen to minimize the total expected cost.

A comparison of our proposed methodology, SCIGAN–ITE, with two ablations is shown in Table A2. Our proposed, individual policy, SCIGAN–ITE, prescribes the PM frequency based on the individual treatment effect (ITE) estimated using SCIGAN. This proposed approach is compared with two ablations. The first, MLP–ITE, does not account for selection bias. The second, SCIGAN–ATE, is a general policy based on the average treatment effect (ATE) and does not differentiate based on machine characteristics.

Methodology	Selection bias?	Individualized?
SCIGAN–ITE	✓	✓
MLP–ITE	✗	✓
SCIGAN–ATE	✓	✗

Table A2: Methodologies comparison.

A4. Results

An visualization of our semi-synthetic setup is shown in Figure A3. The simulated selection bias for different values of λ is shown in Figure A4.

(a) Potential outcomes (b) Costs

Figure A3: **Semi-synthetic data.** We show the observed outcomes in the training and validation set with dots and potential outcomes in the test set with a line. The average potential outcomes and cost are shown with a bold line.

A more detailed comparison of the prescribed PM frequencies per policy is shown in Figure A5.

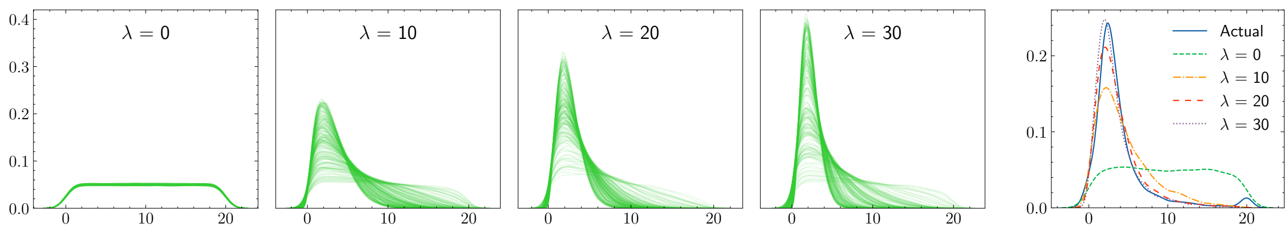

(a) Effect of λ on the individual machines' maintenance distributions

(b) PM frequencies

Figure A4: **Simulating selection bias.** (A4a) We show the distributions that govern the PM frequency for different machines. As these distributions depend on the machine's characteristics, certain machines will more frequently have more maintenance, resulting in selection bias. Moreover, higher values of λ imply more diversity in the distributions and, consequently, more selection bias. (A4b) We show how the PM frequency is distributed among the different machines in reality and as a result of different values of λ. Larger values of λ result in more selection bias with a value of 30 resulting in a PM frequency distribution close to the original.

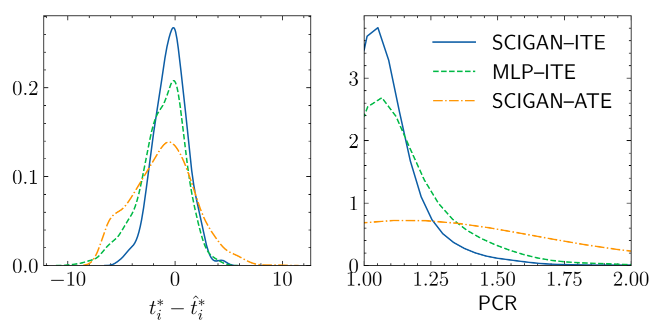

Figure A5: **Evaluating the policies' decisions.** We compare the accuracies and costs for each policy's decisions in terms of the difference between the prescribed and ideal maintenance level (left), as well as the policy cost ratio (right). Results are shown for one representative iteration.