Prescriptive Process Monitoring for Cost-Aware Cycle Time Reduction

Zahra Dasht Bozorgi*, Irene Teinemaa[†], Marlon Dumas[‡], Marcello La Rosa*, Artem Polyvyanyy*

*The University of Melbourne, Australia

zdashtbozorg@student.unimelb.edu.au, {marcello.larosa,artem.polyvyanyy}@unimelb.edu.au

[†]Booking.com, The Netherlands irene.teinemaa@booking.com

[‡]University of Tartu, Estonia

marlon.dumas@ut.ee

Abstract—Reducing cycle time is a recurrent concern in the field of business process management. Depending on the process, various interventions may be triggered to reduce the cycle time of a case, for example, using a faster shipping service in an order-to-delivery process or calling a customer to obtain missing information rather than waiting passively. However, each of these interventions comes with a cost. This paper tackles the problem of determining if and when to trigger a time-reducing intervention in a way that maximizes a net gain function. The paper proposes a prescriptive monitoring method that uses orthogonal random forests to estimate the causal effect of triggering a time-reducing intervention for each ongoing case of a process. Based on this estimate, the method triggers interventions according to a userdefined policy. The method is evaluated on two real-life datasets.

Index Terms-prescriptive process monitoring, cycle time

I. INTRODUCTION

Reducing cycle time (i.e. the time spanned between the start and completion of a case) is a recurrent improvement objective in the field of business process management [1]. Depending on the process, there may be several *interventions* (a.k.a. *treatments*) that workers may perform to reduce cycle time. For example, in an application-to-approval process, giving a call to a customer to obtain missing information may speed up a case. However, such interventions come at a cost, and they can only be done if a suitable worker is available.

In this setting, this paper tackles the following problem: *Given an intervention that in general reduces the cycle time of a case, for which cases and at which point in the execution of a case, should we trigger this intervention to maximise the total net gain?* Here, the total net gain is the sum of the differences between the benefit of each intervention and its cost. To tackle this problem, we propose a prescriptive monitoring method that triggers interventions based on continuous observation of the process. The method relies on a causal effect estimation technique, namely orthogonal random forests, to estimate the effect of an intervention on the cycle time of a case. When the estimated causal effect exceeds a user-defined threshold (the *policy*), an intervention is triggered. The method incorporates an approach to select the policy based on estimated net gain. We evaluate the method using two real-life datasets.

The next section motivates the approach via an example. Sect. III reviews related work. Sect. IV introduces preliminary concepts and notations. The proposed method is described in Sect. V, while the evaluation is discussed in Sect. VI. Sect. VII draws conclusions and discusses future work directions.

II. MOTIVATING EXAMPLE

We consider a loan origination process that starts when a client submits a loan application. The submitted documents are screened. If some documents are missing, a request for further information is sent to the customer either by email or phone. Handling missing document events via email takes less effort. Historical data, though, shows that customers respond faster to phone calls than to emails, thus reducing the cycle time. But given that phone calls are more costly, workers would rather only call a customer if this would speed up the process.

One approach to determine which customers to call (and when) is by predicting the remaining time of each case and triggering an intervention on those cases that are expected to be most delayed w.r.t. a cycle time target [2]–[4]. However, this approach may be ineffective. Suppose that a case is likely to have a long cycle time because the employee handling it is busy. If so, making a phone call to obtain the missing information (the intervention) would consume more resources with little effect on cycle time. A more suitable approach is to estimate the effect of a phone call on the cycle time of each case and to direct the interventions to cases with the highest estimated effect. This paper pursues this approach.

III. RELATED WORK

A. Prescriptive Process Monitoring

Various prescriptive process monitoring methods have been proposed. Teinemaa *et al.* [2] propose to trigger an intervention when the probability of a case leading to a negative outcome is above a threshold optimised w.r.t. a net gain function. The method by Weinzierl *et al.* [5] uses predictive models to determine which activity, among the most likely next activities in the case, is correlated with higher values of a given KPI. Metzger *et al.* [4] use deep learning models to generate predictions about an ongoing case and feed these predictions to an online reinforcement learning technique, which triggers an intervention based on the predictions and their reliability. The above methods tackle the problem of identifying which cases need an intervention, while our approach aims to identify cases that can benefit most from a given intervention.

B. Causal Inference in Process Mining

Causal inference has been widely applied in the field of process mining. Hompes et al. [6] propose an approach to discover cause-effect relations between aggregate characteristics of a process (e.g. frequency of activity) and process performance indicators (e.g. mean cycle time). Koorn et al. [7] present a method to identify causal relations between a response and an effect. They use statistical tests to discover action-response-effects, where the action defines a subgroup of cases and the response is a treatment that enhances the probability of the effect. Polyvyanyy et al. [8] present causality mining, a systematic approach to discovering causal dependencies between events encoded in large datasets. Narendra et al. [9] use structural causal models to confirm potential cause-effect relations identified by analysts. Qafari et al. [10] use structural equation models to test for the existence of a causal relation between an attribute and an outcome. They later expand the use of SEMs in [11] for counterfactual explanation of caselevel predictions. The above studies do not quantify the effect of a treatment at the level of an individual case. Hence, they cannot be applied to individualized prescriptive process monitoring, which is the focus of this paper.

In previous work [12], we use action rule mining to extract candidate treatments correlated with a positive outcome. That work deals only with binary case outcomes, whereas in this paper we deal with a numerical target variable (cycle time).

IV. PRELIMINARIES AND DEFINITIONS

A. Event logs and traces

We refer to an instance of a process execution as a *case*. A case consists of a set of *events*, where an event represents the execution of an activity. Each event has three attributes: a *case identifier* specifying which case the event belongs to, an *activity* that triggered the event, and a *timestamp* specifying when the event occurred. An event may have further attributes, such as a resource that carried out the activity or an event type.

Definition 1 (Event, Trace, Event Log). An *event* is a tuple $(a, c, t, (d_1, v_1), \ldots, (d_m, v_m))$, $m \in \mathbb{N}_0$, where a is an activity name (label), c is a case identifier, t is a timestamp, and $(d_1, v_1) \ldots, (d_m, v_m)$ are attribute-value pairs. A *trace* is a finite sequence $\sigma = \langle e_1, \ldots, e_n \rangle$, $n \in \mathbb{N}_0$, of events with the same case identifier. An *event log* L is a multiset of traces.

As we aim to estimate the causal effect of an intervention on the cycle time of a case, we are interested in events that occurred before the intervention. We use the notion of a kprefix to capture such preceding events.

Definition 2 (k-Prefix). A *k-prefix* of a trace $\langle e_1, \ldots, e_n \rangle$, $n \in \mathbb{N}_0$ is a sequence $\langle e_1, \ldots, e_k \rangle$, $0 \le k \le n$.

B. Causal Inference and Causal Machine Learning

The field of causal inference is concerned with determining the independent effect of a phenomenon on an outcome of interest. Two causal inference frameworks are discussed in the literature. In the *causal graphical models* framework, a causal graph is constructed by a domain expert. Given a causal graph, it is possible to ascertain whether a causal estimand is identifiable and if so, it can be estimated using automated methods. While this framework is focused on the identification of causal effects, the *Neyman-Rubin potential outcomes* framework is focused on the estimation of causal effects of interventions. Hence, in this paper, we use the *Neyman-Rubin potential outcomes* framework.

An intervention (a.k.a. *treatment*) is represented by a binary variable $T \in \{0, 1\}$, where T = 1 indicates that the treatment is applied and T = 0 that it is not. Each case in a log has two potential outcomes: the outcome under treatment (Y^1) and the outcome under no treatment (Y^0) . Given this, the *conditional average treatment effect* (CATE) is defined as follows.

Definition 3 (Conditional Average Treatment Effect). Let X be a set of attributes that characterise a case.

$$CATE: \theta(x) = \mathbb{E}[Y^1 - Y^0 | X = x]$$

CATE is a causal estimand, meaning that in order to estimate it, we need to have access to both Y^0 and Y^1 . However, in real-life, it is impossible to follow both potential realities. This is known as the "fundamental problem of causal inference" [13]. Thus, to estimate the causal estimand, we express it via statistical estimands, such as $\mathbb{E}[Y | T=1, X]$ and $\mathbb{E}[Y | T=0, X]$, which can be estimated from data. According to the potential outcomes framework, CATE can be expressed via statistical estimands only if the *exchangeability, positivity, consistency,* and *no interference* conditions hold.

The exchangeability (a.k.a. ignorability) condition means that given the pre-treatment attributes X, treatment assignment is independent of the potential outcomes. In other words, after conditioning on X, the treatment assignment should be as good as random, which ensures that the treated and not treated groups are exchangeable, that is:

$$Y^1, Y^0 \perp\!\!\!\perp T \mid X.$$

The positivity condition means that for every set of values for X, treatment assignment is not deterministic. So, every subgroup of interest has some chance of getting either treatment:

$$P(T = t | X = x) > 0$$
, for all t and for all x.

The consistency condition states that the potential outcome under treatment T = t is equal to the observed outcome if the actual treatment received is T = t:

$$Y = Y^t$$
 if $T = t$ for all t .

According to no interference, the potential outcomes of one subject are not affected by the treatments received by others.

The positivity and consistency conditions can be verified from data and by ensuring that there are no multiple definitions of the treatment under study. However, verification of the exchangeability condition is not straightforward. In observational data, such as event logs, there often exist variables that influence both the treatment assignment and the outcome. The existence of these *confounding variables* (or *controls*) creates a non-causal association between T and Y, which can invalidate the study. The best way to circumvent this problem is to conduct a randomised experiment (A/B test). As we are working with event logs, where randomisation of the treatment is not ensured, we *assume* that exchangeability holds. Under this assumption, the observed variables X contain sufficient information needed to adjust for confounding. The adjustment can then be carried out during the estimation step.

In many real-life use cases, including business processes, the no interference condition can often be violated as well. For instance, skipping an activity in one case might leave the process worker available to perform activities in another case, and vice versa. As a result, the potential outcomes of the second instance are affected by the treatment applied in the first. However, it is common to assume that a violation of the no interference condition has only a small influence on the causal effect estimates and is considered negligible in practice.

Next to identification, the other core problem in causal inference is estimating CATE from observational data. A large body of works from recent years focus on using machine learning methods for CATE estimation. A popular method is the work by Athey *et al.* [14]. which is a flexible non-parametric estimation method based on *generalised random forests*. However, it does not allow for high dimensional set of confounders. This problem is addressed in a method called *double machine learning* proposed by Chernozhukov *et al.* [15]. Oprescu *et al.* [16] generalise the ideas from these two approaches in an approach called *orthogonal random forests*.

C. Orthogonal Random Forests

We use orthogonal random forests (ORFs) to estimate the effects of treatments on cases. When applied to the problem of heterogeneous treatment effect estimation, ORFs assume the data D to be in the form $D = \{(T_i, Y_i, W_i, X_i)\}_{i=1}^n$, for n observations. For each observation i, T_i is the received treatment, Y_i is the observed outcome, W_i represents potential confounding variables, and X_i are the features capturing heterogeneity. The ORF method makes the assumptions captured in the following structural equations:

$$Y = \theta(X) \cdot T + f(X, W) + \epsilon,$$
$$T = g(X, W) + \eta,$$

where f(X, W) models the outcome Y when no treatment is applied, $\theta(X)$ is the treatment effect function (CATE), and g(X, W) captures the relationship between treatment T, confounders W, features X, and ϵ and η are unobserved noises, such that $E[\epsilon | X, W, T] = 0$ and $E[\eta | X, W, \epsilon] = 0$.

ORF follows the residualisation approach used in double machine learning. Double machine learning, as the name suggests, consists of two stages. In the first stage, a model is fit to predict Y from X and W, and another model is fit to predict T from X and W, that is:

$$\hat{Y} = \mathbb{E}[Y \mid X, W]$$
 and $\hat{T} = \mathbb{E}[T \mid X, W].$

In the second stage, the effects of confounding are removed by fitting a model to predict $Y - \hat{Y}$ from $T - \hat{T}$. The treatment effect function $\theta(x)$ is defined as follows:

$$\mathbb{E}[Y - Y - \theta(x) \cdot (T - T) \mid X = x] = 0.$$

V. Approach

We aim to design a recommender system that seeks to minimise the cycle time of a process case in a cost-aware manner. At the core of our system is a causal estimation module that estimates the effect of a given treatment on a target metric. The target is to decrease the cycle time of a process case. We assume that the treatments are binary. For instance, suppose that in the loan origination process discussed in Sect. II, the default way of handling missing documents is by sending emails to clients to request additional documents, while some customers get phone calls instead. In this scenario, the binary treatment options are sending emails (T = 0) or making phone calls (T = 1). We refer to cases that got a phone call as the treatment group and the cases that got an email as the control group. Making a phone call should, in general, speed up the process. However, it is not feasible to call every client because of the associated additional costs. Thus, we want to find the best policy for deciding which clients should get phone calls.

Our approach consists of two phases, as depicted in Figure 1. In the *offline phase*, case prefixes are extracted from the log and used to train a causal effect estimation model. Then, the best model threshold is selected. This threshold defines a treatment policy. In the *online phase*, the causal effect estimator and the selected policy are used to recommend whether to treat or not to treat an ongoing case based on its estimated treatment effect and the selected treatment policy.

Next, we describe the three steps of the offline phase, namely *data pre-processing*, *causal model construction*, and *policy selection*), followed by the online phase.

A. Data Pre-processing

1) Data Cleaning: First, we pre-process the log to remove incomplete cases. We then repair missing attribute values. For a numeric attribute, we set missing values to the median value of that attribute. For a categorical attribute, we set missing values to the most observed value for that attribute (the mode).

2) k-Prefix Extraction: To construct feature vectors, we select a single prefix of length k from each case. In cases where the treatment assignment is present in the data, k is the length of the prefix before the actual treatment time. For instance, given a trace $\sigma = \langle e_1, \ldots, e_n \rangle$, suppose that e_{k+1} indicates the treatment event where $k + 1 \leq n$. In this case, the selected prefix is $\langle e_1, \ldots, e_k \rangle$. Thus, k can be different for different cases. In cases where the treatment did not occur, we estimate the treatment point by randomly drawing from the distribution of prefix lengths when the treatment is present.



Figure 1: Overview of the proposed approach

3) Prefix Encoding: To apply the machine learning method for estimating treatment effects, we need to encode trace prefixes (ongoing cases) as fixed-sized feature vectors. Encoding static case attributes is straightforward since their values do not change during the execution of a case. However, traces often contain dynamic event attributes whose values change as the case unfolds. Various methods have been proposed to encode dynamic attributes [17], including aggregation encoding, last-state encoding, index-based encoding, and tensor encoding [18]. We use aggregation encoding for the activity type and resource attributes and last-state encoding for event attributes. Hence, the feature vector constructed from a kprefix contains one numeric feature for each activity type or resource A, the number of times A appears in the k-prefix. For other event attributes (e.g. "offered loan amount" or "type of loan offer"), the feature vector encodes only the last value of this attribute (e.g. the last offered loan amount). If an encoded attributes is a categorical attribute, we apply one-hot encoding to represent it as a numeric feature vector - a common practice when using tree-based ensemble models, such as ORF.

4) Temporal and Workload Features: Given that we seek to estimate treatment effects on cycle time, we include temporal information in the feature vectors. Specifically, we include the month, weekday, and hour of the timestamp of the last event in the prefix, the time between the first and the last event in the prefix, and the time between the last two events in the prefix (i.e. the inactivity period prior to the most recent event).

We also include the number of active cases as a feature to act as a proxy for the current workload in the process. We do so because workload is a potential confounder, influencing both the cycle time and the treatments (during high workloads, cases may be delayed, and the effect of treatments might be weaker than usual). Finally, we encode the difference between the first event in the prefix (the start time of the case) and the first event in the log (the start time of the log's timeframe) since the process might behave differently at different points in time; hence, this feature can be a confounder.

B. Causal Model Construction

This step of the offline phase takes as input the encoded k-prefixes and learns a function that, given a k-prefix $\langle \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(k)} \rangle$, returns a point estimate of the treatment effect θ and the corresponding confidence interval.

We use ORFs to implement the causal model. The use of ORFs has several advantages. ORFs support non-parametric estimation of the target variable while allowing for a high-dimensional set of confounding variables W. This is useful in the context of process mining since traces typically have

a large number of event attributes, which can lead to feature explosion. The problem is exacerbated when there is a large number of resources and categorical attributes. Finally, due to ORFs being asymptotically normal (i.e. the distribution of the estimated treatment effects approaches a normal distribution as the sample size grows), they allow for the construction of valid confidence intervals based on bootstrapping.

An ORF requires four inputs: the target variable Y, the treatment indicator T, features capturing heterogeneity X, and the confounding variables W. Since this recommender system aims to decrease cycle time, the outcome Y is the cycle time of the process. We assume that a binary treatment is previously identified, which is hypothesised to cause a decrease in cycle time. In the loan origination example, this is a phone call.

As the proposed approach can be applied at an operational level, we would like to know the treatment effect when it is time for the process worker to decide whether to treat a case or not. Therefore, effect heterogeneity is captured via static case-level attributes and all event attributes (including the activity name and the resource), and the inter-case features available at the decision time. This means that X is a feature vector that captures all of this information about the k-prefix after the suitable encodings have been applied, as described above. In this study, we assume that all of the available features derived from the event log are also potential confounders, meaning that all the features present in X are included in W. However, in general, X and W do not need to be the same. A domain expert can remove some features from W if she believes they do not influence the treatment decision or outcome.

Armed with these definitions of Y, T, X, and W, the next step is to train an ORF. The input data is temporally split into the train and test sets. A separate test set is required for the evaluation of the trained model, and for selecting the best policy. As mentioned in the preliminaries, we need to learn the treatment effect function θ by solving the same set of equations as in double machine learning. This is a two-phase approach that requires \hat{Y} and \hat{T} to be modeled locally at X = x, obtaining the estimates \hat{Y} and \hat{T} . In the second stage, the treatment effect is estimated by minimising the residual of Y(denoted by $Y - \hat{Y}$) on the residual of T (denoted by $T - \hat{T}$) locally at X = x using the squared loss:

$$\theta(x) = \operatorname{argmin}_{\theta} \mathbb{E}[(Y - Y - \theta \cdot (T - T))^2 \,|\, X = x]$$

ORFs allow the use of flexible models for the estimation of \hat{T} and \hat{Y} . These models can be chosen based on their performance on specific datasets. Also, at prediction time, we have the option to use different models to residualise the treatment and the outcome. Fitting \hat{Y} and \hat{T} locally around the target feature X ensures that more weight is put on samples that are similar in the feature space. A random forest with a causal criterion is constructed to capture the similarity metric in the X space, so that samples that are similar in the feature space X have similar treatment effects.

C. Policy Selection

Oftentimes, applying treatments to cases comes at a cost. Making phone calls to customers is an example of this. Having a treatment effect estimate for every running case can help decision makers separate the cases that would benefit from the treatment from those that would not be affected or would be negatively affected by it. Particularly, if the cost of the treatment is not far below from its benefit, it is important to carefully select a policy to determine which cases to treat. Accordingly, in this step, we propose a policy selection module that maximises the net benefit given the cost of the treatment and the treatment effect estimates. Furthermore, our proposed policy selection approach provides a way to evaluate the ORF model that is used to estimate the treatment effects.

The first step in the policy selection module is to build a *Qini curve*. Qini curves and the closely related *uplift curves* provide a way to evaluate CATE estimators when the ground truth treatment effect is not available, which is the case in all real-world datasets [19]. Since our target variable in this study is continuous and the goal is to reduce it, we use a modified version of the Qini curve presented below.

Given an estimator $\hat{\theta}$ and k-prefixes σ_i (with the treatment decision happening at the k-th event), let π be the ascending ordering of the traces according to their estimated treatment effect, i.e., $\hat{\theta}^{\pi}(\sigma_i) \leq \hat{\theta}^{\pi}(\sigma_j), \forall i < j$. In addition, $\pi(n)$ is used to denote the first n percent of traces from the ordering.

Let $R_{\pi(n)}$ be the sum of the duration of cases in $\pi(n)$, i.e., $R_{\pi(n)} = \sum_{i \in \pi(n)} Y_i$. Furthermore, let $R_{\pi(n)}^{T=1}$ and $R_{\pi(n)}^{T=0}$ be, respectively, the sums of the durations of cases in the treated and control groups in $\pi(n)$. Also, to denote the number of traces in the treated and control groups in $\pi(n)$, we use $N_{\pi(n)}^{T=1}$ and $N_{\pi(n)}^{T=0}$, respectively. We now define the Qini curve as:

$$Qini(n) = R_{\pi(n)}^{T=0} \times \frac{N_{\pi(n)}^{T=1}}{N_{\pi(n)}^{T=0}} - R_{\pi(n)}^{T=1}$$

Qini(n) is the expected total reduction in cycle time, given that the top *n* percent of cases selected by the ORF model are treated (e.g. we make phone calls to the top *n* percent of cases). The Qini curve can be plotted by computing Qini(n) for different values of *n*. In addition to evaluating our ORF model, we use the Qini curve for policy selection by incorporating the cost and the benefit multipliers to create *net value curves*. Suppose *v* is the value of reducing the cycle time by one unit of time and *c* is the cost of applying the treatment to one case. Then, the net gain of applying the treatment to $\pi(n)$ is defined as follows:

$$gain(n) = v \times Qini(n) - c \times N_{\pi(n)}^{T=1}$$

Similar to the Qini curve, the net value curve can be drawn by computing gain(n) for different values of n. The user has the option to view the curve and to select the optimal policy based on organisational constraints. For example, suppose the organisation has a target of achieving a net benefit of x. In that case, the net value curve will provide the minimum proportion of the population that needs to be treated to achieve that goal. If there are no constraints, the policy that yields the highest net value can be selected automatically.

D. Online phase

In the online phase, the applicability of the treatment for an ongoing case with an observed k-prefix is assessed. If the treatment is not applicable at that point, the assessment is repeated after the next event. If the treatment is applicable, the k-prefix of the case is encoded as a feature vector using the same approach as in the offline phase. Then, the treatment effect is estimated using the pre-trained ORF model. If the estimated net gain is sufficiently high according to the preselected treatment policy, the prescriptive monitoring system recommends applying the treatment.

Let us come back to the example loan origination process. First, we take a log of this process, clean it, encode the caselevel and event-level features, and create a treatment attribute, where the value of this attribute is zero for cases that got an email and 1 for cases that got a phone call. We then divide this pre-processed dataset into training and testing sets. We use the training set to train an ORF model. This trained model can then take an incomplete case as input and returns, as a number, the estimated effect of calling the customer. We then estimate the effect of the phone call on all the cases in the test set. We use these estimates to create a net-value curve. The policy maker uses the curve to decide the percentage of future cases that will get the phone call based on organisational constraints.

In the online phase, a new application is created by the customer. After the execution of each activity, the applicability of the phone call at that stage is assessed. In the initial stages of document processing, there is no need for the phone call, so we move on to the next activity. If we encounter missing documents in this case, we encode the information about this case using the same method as in the offline phase and estimate the treatment effect using our trained ORF. Suppose that the selected policy is to treat half of the cases. If the estimated effect for this new case is in the top 50% of the cases we have had so far, we call this customer; otherwise, we send an email.

VI. EVALUATION

To assess our approach, we conducted experiments on two real-life logs, measured the estimated causal effect of our recommendations, and compared the results with improvement suggestions identified by non-causal methods commonly used for prescriptive process monitoring [2]–[4].

A. Experimental Setup

The logs we used are the BPI17 and BPI19 logs, available from the 4TU Centre for Research Data.¹ We used these two datasets because they contain interventions that could possibly have a causal effect on the cycle time. Moreover,

¹https://data.4tu.nl

Log	Number of trees	Minimum leaf size	Maximum depth	Sub-sample ratio	Lambda reg
BPI17	200	20	30	0.4	0.01
BPI19	200	10	10	0.7	0.01

Table I: Parameter setting for training the ORF models

each log shows unique features. BPI17 is characterised by a combination of case-level and event-level attributes, while most of the attributes in the BPI19 log are case-level.

BPI17: This log contains traces of a loan application process of a Dutch financial institute. The data contains attributes about the applications and the loan offers made by the bank.

BPI19: This log contains traces from a purchase-to-pay process of a Dutch multinational company. Each case describes a purchase order item from its creation to payment. This log also records activities such as changes, cancellations, and message exchange related to purchase orders.

Evaluation measures: In the literature, CATE estimators are often evaluated by computing Qini curves. The underpinning intuition is that if the CATE is estimated accurately, the cases with a positive outcome in the treated group would have a higher estimated CATE than those in the same group with negative outcomes. Also, cases in the control group with a negative outcome should have a higher CATE than the positive outcome cases in the same group. Thus, a desirable causal model has a Qini curve above the random curve.

Training setting: We split the data into 60%–20%–20% for training, validation, and testing, respectively, by preserving the temporal order between cases. We used the training set to train the ORF model, the validation set to tune the hyperparameters, and the test set to provide an unbiased evaluation of the model. The details of the training settings are shown in Table I.

B. Results

We discuss the results obtained by running the ORF algorithm on the two logs above, and compare our results with noncausal methods based on two predictive models: Lasso and Random Forest (RF). We chose these two methods because they had high performance in predicting cycle time in these two logs. As shown in [20], their performance on these logs is only slightly lower than deep adversarial models, which are the state-of-the-art in predictive monitoring, while having significantly lower training time. To resemble the setting assumed by predictive monitoring techniques, these models were trained on untreated cases only. Specifically, the baseline approach is one that predicts the cycle time given that no intervention is made, and triggers an intervention if the predicted cycle time is above a threshold. This threshold is set in such a way that X% of the cases in the testing set are targeted, where X is an independent variable in the experiment.

1) BPI17: This log shows a number of cases with unusually long duration. This severely impacts customer experience and, in many cases, leads to a negative outcome. However, since this is an interactive process, the cause of delay in many cases is that the bank is waiting for input from the customer to move on with the process. So, one possible intervention is to contact the customer to ask for additional information. While this intervention is effective for many of the cases observed in this log, the cost of calling all customers might be too high, or the company might not have the required resources to call everyone. So, we applied our approach to identify which cases benefit more from receiving this additional phone call.

2) BPI19: In this log, we observe some activities that indicate a change in the purchase order items. Since changes might lead to re-work, we hypothesised that avoiding changes leads to lower cycle times. Specifically, we considered skipping of activity Change Price as our treatment for this experiment. However, fixing the price of purchase orders all at the same time for each case without considering the specific context of each case leads to a rigid order placement procedure and might require extra work at the beginning of the process. Thus, rather than fixing the price at the beginning of every case, it is beneficial to have a targeted approach. Furthermore, our approach provides a recommendation for each case, indicating whether a price change should be permitted or not. In this way, we can avoid the rework in cases where this change is harmful and increases cycle time while preserving flexibility in other cases that are not highly affected by price change.

3) Discussion: The plots in Figure 2 show the Qini curves of the models constructed in our experiments. The black dashed line shows the expected incremental reduction in cycle time if a random policy was used to treat a certain percentage of the cases and the blue, red, and green lines show the reduction if the policy was based on the ORF, RF, and Lasso models respectively. It can be seen that in both datasets, the expected reduction is higher if the ORF-based policy is followed rather than policies based on non-causal predictive models. Notably, the curves for RF and Lasso are below the random policy line indicating that while predictive models are good at identifying which cases will take a long time, they are not necessarily good at identifying which cases should be targeted with the chosen treatments. The Qini curve for the BPI17 log shows that treating 90% of the cases with the highest treatment effect gives the same reduction as treating everybody. Particularly in the curve for BPI19, it can be seen that just treating 90% of cases gives a higher reduction in cycle time than treating everybody. This curve is not monotonically increasing as usual cumulative gain charts because the quantity on the y-axis is the *difference* between treated and non-treated cases, which in some segments can be negative since the treatment hurts the outcome in these cases. In the plot for BPI17, we do not see a decrease because the treatment does not hurt the outcome, but the treatment effect is low enough not to result in any gain in the cases in the last decile.

The Qini curves are without cost and benefit multipliers. So, we proceeded to plot the net value curves for both models with varying values of v/c, where v is the value of reducing the cycle time by one day, and c is the cost of applying the treatment to one case. Rather than the absolute values of v and c, it is the ratio v/c that affects the shape of the curve. Fig. 3 shows the net value curves with different values of v/c for the



Figure 2: Qini curves for ORF models.

two logs. We observe that as the ratio v/c decreases — i.e. the treatment becomes more expensive relative to the benefit it provides — the net value of treating the cases decreases, and so, it becomes more important to apply the treatment with a more targeted approach. Particularly, in the BPI17 log, the value of applying the treatment decreases more quickly than in the BPI19 log. For example, when v/c = 0.3 the net value of applying the treatment to all cases is close to zero and the highest benefit is gained if half of the cases are treated. But in the BPI19 log the highest gain comes with treating 90% of cases regardless of the value of v/c, and treating all cases still provides a gain even when v/c = 0.3. This is because the cases in the former log are shorter and the treatment has less effect, on average, than in the latter log. Indeed, the mean duration in the BPI17 is 20 days and the average of the estimated effects is 1 day, while in the BPI19 log the mean duration is 150 days and the mean treatment effect is 35 days.

C. Sensitivity Analysis

Like other treatment effect estimation methods, the ORF algorithm works under the assumption of no unobserved confounding. However, when unobserved confounders do exist, our treatment effect estimates might have some bias. So, one threat to the internal validity of our results is the possibility of the presence of unobserved confounders. However, the effects of these unobserved confounders might be mild enough not to change our inference. We perform sensitivity analysis to determine how strong the effects of a confounder need to be in order for our model to change fundamentally. To achieve this, we used the Austen plots sensitivity analysis method by Veitch and Zaveri [21]. We chose this approach because it fully separates sensitivity analysis and modelling of the observed data,



Figure 3: Net value curves for ORF models.

which allows us to use it with any machine learning method. Furthermore, this approach provides an interpretable model of the influence of unobserved confounders, which quantifies the sensitivity of our models to unobserved confounding.



Figure 4: Sensitivity curves for ORF models.

Figure 4 shows the Austen plots for our models for BPI17 and BPI19 logs. The solid black line is the sensitivity curve that indicates values of α , influence on treatment and partial R^2

influence on the outcome. This means that if an attribute has α and partial R^2 values above this curve, it would be sufficient to induce the bias value shown in the plot. The coloured dots show the influence of the observed potential confounders. The plots show that our conclusions from these models are robust to unobserved confounding. Since the coloured dots are well below the sensitivity curve, even the observed confounders do not have sufficient power to change the qualitative conclusions of our experiments. Particularly, the curve for the BPI17 log is very robust to confounding since it allows for the presence of variables that have a high effect on treatment assignment if their influence on the outcome is less than 0.25. Even if an unobserved confounder exists that fully determines the treatment policy, it would need to have an influence above 0.25 on the outcome to induce a bias of one day in the result, which in practice is highly unlikely. The model for BPI19 is more sensitive to the presence of potential confounders if they are highly influential on either the treatment or the outcome. However, since all the observed attributes are shown to be well below the sensitivity curve, the model is still robust if the effects of the hidden confounders are similar to the measured ones. It should be stressed, however, that the presence of strong hidden confounders is always a possibility, and sensitivity analysis is constrained to the available data.

D. Threats to validity

The evaluation comes with a threat to external validity (lack of generalisability) as it focuses on two concrete application scenarios. This can be addressed by conducting further experiments with logs of different characteristics and from different domains. A threat to construct validity is that the experiments are based on observational data. Hence, the estimated treatment effects may not match the true causal effects. To mitigate this threat, we performed sensitivity analysis to assess the robustness of our model to unobserved confounding. A rigorous A/B test should be conducted before deploying the recommendations of our method in an operational setting.

VII. CONCLUSION AND FUTURE WORK

We proposed a prescriptive monitoring method that recommends if and when to apply an intervention (treatment) to an ongoing case to decrease its cycle time. The method relies on orthogonal random forests trained on historical traces to estimate the decrease in cycle time (the treatment effect) given the current state of a case. Based on this estimate, the method calculates the expected gain of the treatment given a cost function and generates recommendations based on a user-defined policy. We evaluated the applicability of our approach based on two real-life logs. The evaluation showed that the proposed approach yields a higher net gain than treatment policies based on non-causal predictive approaches, as previously proposed in the literature. We also showed via sensitivity analysis that the models built on these datasets are robust against potential unobserved confounding effects.

The approach assumes a binary treatment setting, i.e. either a treatment of a given type is applied or not. An avenue for future work is extending this method to accommodate multiple treatments (e.g. choosing between one type of treatment or another) as well as treatments with continuous values. Our method learns the policy for which cases to treat (whom to treat). Another direction for future work is optimising the time of treatment (when to treat). Also, the approach requires the treatment to be pre-defined. Another direction for future work is to design methods to automatically discover candidate treatments from a historical log.

Reproducibility The source code of our tool can be found at https://github.com/zahradbozorgi/prescriptiveMonitoringORF.

Acknowledgments Research funded by the Australian Research Council (grant DP180102839) and the European Research Council (PIX Project). Thanks to Simon Remy, Kiarash Diba, and Luise Pufahl for providing preprocessed datasets.

REFERENCES

- M. Dumas, M. L. Rosa, J. Mendling, and H. A. Reijers, *Fundamentals of Business Process Management, Second Edition.* Springer, 2018.
- [2] I. Teinemaa, N. Tax, M. de Leoni, M. Dumas, and F. M. Maggi, "Alarmbased prescriptive process monitoring," in *BPM*, 2018.
- [3] A. Metzger, A. Neubauer, P. Bohn, and K. Pohl, "Proactive process adaptation using deep learning ensembles," in *CAiSE*, 2019.
- [4] A. Metzger, T. Kley, and A. Palm, "Triggering proactive business process adaptations via online reinforcement learning," in *BPM*, 2020.
- [5] S. Weinzierl, S. Dunzer, S. Zilker, and M. Matzner, "Prescriptive business process monitoring for recommending next best actions," in *BPM*, 2020.
- [6] B. F. A. Hompes, A. Maaradji, M. L. Rosa, M. Dumas, J. C. A. M. Buijs, and W. M. P. van der Aalst, "Discovering causal factors explaining business process performance variation," in *Proc. of CAiSE*, 2017.
- [7] J. J. Koorn, X. Lu, H. Leopold, and H. A. Reijers, "Looking for meaning: Discovering action-response-effect patterns in business processes," in *BPM*, 2020.
- [8] A. Polyvyanyy, A. Pika, M. T. Wynn, and A. H. ter Hofstede, "A systematic approach for discovering causal dependencies between observations and incidents in the health and safety domain," *Safety Science*, 2019.
- [9] T. Narendra, P. Agarwal, M. Gupta, and S. Dechu, "Counterfactual reasoning for process optimization using structural causal models," in *BPM Forum*, 2019.
- [10] M. S. Qafari and W. M. P. van der Aalst, "Root cause analysis in process mining using structural equation models," in *BPM Workshops*, 2020.
- [11] M. S. Qafari and W. van der Aalst, "Case level counterfactual reasoning in process mining," in CAISE Forum, 2021.
- [12] Z. D. Bozorgi, I. Teinemaa, M. Dumas, M. L. Rosa, and A. Polyvyanyy, "Process mining meets causal machine learning: Discovering causal rules from event logs," in *ICPM*, 2020.
- [13] P. W. Holland, "Statistics and causal inference," *Journal of the American statistical Association*, 1986.
- [14] S. Athey, J. Tibshirani, S. Wager *et al.*, "Generalized random forests," *Annals of Statistics*, 2019.
- [15] V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins, "Double/debiased machine learning for treatment and structural parameters," *The Econometrics Journal*, 2018.
- [16] M. Oprescu, V. Syrgkanis, and Z. S. Wu, "Orthogonal random forest for causal inference," in *ICML*, 2019.
- [17] I. Teinemaa, M. Dumas, M. L. Rosa, and F. M. Maggi, "Outcomeoriented predictive process monitoring: Review and benchmark," ACM *Transactions on Knowledge Discovery from Data*, 2019.
- [18] I. Verenich, M. Dumas, M. La Rosa *et al.*, "Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring," *ACM Trans. Intell. Syst. Technol.*, 2019.
- [19] W. Zhang, J. Li, and L. Liu, "A unified survey on treatment effect heterogeneity modeling and uplift modeling," *ACM Comput. Surv.*, 2021.
 [20] F. Taymouri, M. L. Rosa, and S. Erfani, "A deep adversarial model for
- [20] F. Taymouri, M. L. Rosa, and S. Erfani, "A deep adversarial model for suffix and remaining time prediction of event sequences," SDM, 2021.
- [21] V. Veitch and A. Zaveri, "Sense and sensitivity analysis: Simple post-hoc analysis of bias due to unobserved confounding," in *NeurIPS*, 2020.