Process Mining for Healthcare Decision Analytics with Micro-Costing Estimations

Sander J.J. Leemans[a], Andrew Partington[b], Jonathan Karnon[b], Moe T. Wynn[c]

[a] *RWTH, Aachen, Germany*
[b] *Flinders University, Adelaide, Australia*
[c] *Queensland University of Technology, Brisbane, Australia*

Abstract

Managing constrained healthcare resources is an important and inescapable role of healthcare decision makers. Allocative decisions are based on downstream consequences of changes to care processes: judging whether the costs involved are offset by the magnitude of the consequences, and therefore whether the change represents value for money. Process mining techniques can inform such decisions by quantitatively discovering, comparing and detailing care processes using recorded data, however the scope of techniques typically excludes anything 'after-the-process' i.e., their accumulated costs and resulting consequences. Cost considerations are increasingly incorporated into process mining techniques, but the majority of healthcare costs for service and overhead components are commonly apportioned and recorded at the patient (trace) level , hiding event level detail. Within decision-analysis, event-driven and individual-level simulation models are sometimes used to forecast the expected downstream consequences of process changes, but are expensive to manually operationalise. In this paper, we address both of these gaps within and between process mining and decision analytics, by

better linking them together. In particular, we introduce a new type of process model containing trace data that can be used in individual-level or cohort-level decision-analytical model building. Furthermore, we enhance these models with process-based micro-costing estimations. The approach was evaluated with health economics and decision modelling experts, with discussion centred on how the outputs could be used, and how similar information would otherwise be compiled.
Keywords: process mining, decision analytics, healthcare economics

1. Introduction

Within health systems funded through pooled and constrained resources (e.g., insurance and/or tax revenues), policy and executive decision makers are increasingly aware of whether their processes of care delivery achieve patient health, experience, equity and fairness goals to the extent that they warrant their allocation of funding. As such, the question at the heart of all health service planning and project implementation is whether, to what extent and under what conditions/assumptions, should a process exist or be changed? This is the normative decision problem that all analysis techniques including process mining techniques in the health domain seek to inform.

Process mining is a useful method for discovering and quantifying existing care processes, and provides a data-driven means of studying process behaviour [1]. Discovered processes might also reasonably reflect how health services would be expected to be delivered and perform in the future, under a 'status quo' scenario with no changes in patient demand or resource configuration. Visualisations and the ability to analyse information mined from event-log data enable the formation of hypotheses into potential causal effects between process behaviour and the impact of potential interventions to alter processes and affect outcomes of interest.

However, to establish whether and how processes warrant intervention, additional statistical and forward-looking analyses outside of process mining are often required. Recent work has looked to envelop inferential and predictive methods within process mining algorithms [2, 3], but process mining methods can additionally help provide insight into hidden process information and inform other decision-analytical methods for foresight generation.

Healthcare decision makers and stakeholders require foresight into the accrued costs and consequences expected from observed or potential processes.. It is on the basis of 'expected outcomes' and the price that decision makers are willing to pay for them (i.e., their value for money), that healthcare services are planned and implemented.

The expected value of changes to processescan be evidenced through conformance analysis, but only where there exists a comparative reference process with known and 'acceptably priced' effects. Guidelines and clinical practice protocols can provide reference processes that represent value, but often do not exist at the local-level with the consideration of relevant resources that are idiosyncratic to different decision maker contexts.

The other way for process mining analyses to inform value-based decision making in health is through forward-looking (e.g., simulation) modelling. These analyses/modelling fall outside the 'discovery, conformance and enhancement' capabilities of process mining, but are a crucial part of 'closing of the loop' which ensures insights and methods have practical relevance.

The discovery of simulation models is not a new area for process mining [4]. Recent work in process mining also explored improving the accuracy of simulation models in the event of multitasking and availability constraints [5]. However, while the health domain possesses a large amount of data, it is characterised as flawed in its collection; uncertain in its definitions; it usually captures only proxy measures of desired information; and is sparsely populated for important subgroups of data [6]. It is rare that the process data exists to accurately predict endpoints that inform a value-proposition - namely down-stream costs and patient outcomes. Individual level and discrete event simulation decision-analytic models are difficult to operationalise within tight project constraint, and may only provide little additional information to decision makers [7].

Given the rarity of reference processes that are value-based, and the impracticability of individual-level simulations, further process mining capabilities are needed to enhance and make use of flawed data to inform simple forward-looking decision-analytic models that link processes to costs and outcomes.

In this paper, we introduce process mining approaches to help inform the structure and parameter estimation of forward-looking, decision-analytic models that can be used for the management of healthcare resources. Specifically, we detail:

1. The linkage and transformation of a routine administrative, clinical and costing data set, from a set of 'episode logs' used in activity-based funding, into event logs that describe patient-pathways.

2. The automation, visualisation and description of data within an event

log, so that information can be easily abstracted into simple forward-looking models, such as decision trees and Markov models.

3. The apportionment of case-level attributes on costs, for the purposes of micro-costing at an event level.

To summarise, the key contributions of this paper are:

1. An example of an event log compiled from linked hospital data that is routinely collected across Australia, with standardised nomenclatures, for the purpose of activity-based funding.

2. A new approach that combines process modelling with economic decision tree modelling, to express processes with explicit trace attributes (e.g., demographic data); as well as discovery techniques, model visualisations and conformance checking techniques for such models.

3. Estimation of micro-cost information for activities of interest within the event-log, from aggregate episode-level costs recorded as trace attributes.

4. An approach to help inform real-world decisions for managing healthcare delivery based on 'value for money', by bringing the mining of processes, and the consideration of costs and outcomes, closer together.

The approach is demonstrated using a case study of real-life data from an Australian Local Health Network, and validated through quantitative tests of estimation accuracy and qualitative interviews with healthcare decision-modelling analysts who would use the mined information.

The remainder of this paper is organised as follows: Section 2 discusses related work, Section 3 introduces the context of the case, Section 4 introduces the technical process mining contributions, Section 5 evaluates and discusses the approach and Section 6 concludes the paper.

2. Related Work

In this section, we discuss three areas of related work: economic evaluation in healthcare delivery, process mining in healthcare, and cost-based process mining.

2.1. Economic evaluation in healthcare delivery

Normative economic evaluation, whether model-based or alongside clinical trials, is a core component of health economics and has existed as an applied science to inform healthcare decision making since the mid-1960s [8]. A major challenge in applying economic evaluation methods within healthcare, is that the systems for its delivery continuously and spontaneously (i.e., dynamically) adapt and self-organise based on the complex interaction of multiple stakeholders with overlapping responsibilities and competing interests [9, 10]. In this context, there is some guidance on combining process evaluation and outcomes evaluation [11], but relying most heavily on qualitative and subjective methods [12].

The success of any service intervention hinges on the sense amongst managers of the system that they are able to 'sell' the plans to diverse interest groups to facilitate 'buy-in' according to individual and collective goals [13, 14]. Because of this, economic evaluation and the modelling involved is itself a necessarily dynamic and iterative process, where a model is only as useful as its use in stakeholder engagement [15]. This means conforming to the confines of the short timing, constrained resourcing, a lack of expertise and a generally negative attitudes towards resource rationing in the teams responsible for service developments [16].

As outlined by Brennan et al. [17], there are a number of different types of model structures which can be employed, depending on a preference for cohort level vs. individual level; with interactions or without interactions; continuous state (deterministic) vs. discrete state (stochastic or individual); and untimed vs. timed (either discrete or continuous). The resulting different modelling approaches include decision trees, Markov models, system dynamics, individual event histories and discrete event simulations (DES).

Novel whole-of-hospital DES platforms [18, 19] illustrate the potential for virtual piloting of potential new services with methods/model structure that could be translated at a low-cost to other hospital settings. However, as was the reported experience of Lord et al. [20] with their modelling of full care pathways using DES for economic evaluation, development resources and analysis time can be higher and take longer than anticipated and fail to fit within decision maker and project constraints. There is also a relevant 'value of information' calculation to make, where a more complicated and expensive approach generates similar outputs to relatively simpler processes, any additional modelling expenses are unlikely to be efficient [7]. Trying to forecast the efficiency of alternative modelling structures is a contemporary topic of research and debate [21]. When selecting a structure, the determinative principle should be that of a 'requisite approach' [22], whereby only as much modelling is undertaken to sufficiently minimise decision uncertainty and nothing more. With this in mind, it makes sense to start with as simple a model as possible, and add to it as necessary.

Regardless of the structure chosen, the measurement and valuation of resources and their costs is an integral component of any economic modelling. There are top-down and bottom-up costing methods, depending on whether collected at a cost-centre or the individual patient level; and micro and gross costing methods, depending on level of detail in different types of resources utilised [23]. There is a great deal of variation in methodological and reporting quality in costing studies for economic evaluation [24]; however, the main emphasis is on transparently reporting resource quantities and the source of unit prices, so that any direction of bias can be interpreted by a reader [25].

When looking to inform decisions at a local level that are dependent on idiosyncratic resource availabilities and organisational context, bottom-up approaches are preferred [26]. Analytical and cost accounting is practised within Australian hospitals according to the business rules and guidelines within the Australian Hospital Patient Costing Standards [27] and increasingly employ Patient-Level Costing Information Systems (PLICS). This costing data informs the prices paid or reimbursed for Diagnosis Related Groups (DRGs) or other units of activity. While there is a breakdown for both direct and indirect costs for 'buckets' of different types of resources (e.g., workforce, consumables, overheads), such data is commonly only available at an episode or trace level rather than an event level. While the costs for goods and services are easily accounted from the bottom-up, the majority of healthcare costs associated with service and overhead components are often apportioned through subjective and opaque methods. Further, the structure and definitions of data within PLICS are not necessarily aligned with the ‘states’, ‘events’, ‘actions’ or ‘decisions’ on which a decision-analytic model might ideally be based.

2.2. Process Mining in Healthcare

Healthcare is a domain which has received substantial attention from the process mining research community. Healthcare processes can be described as “a series of activities aimed to diagnose, treat and prevent any diseases in order to improve a patient’s health” [28]. Healthcare processes are considered highly complex with significant variations over time [29] due to the patient-centric nature of treatment pathways. Process mining has been used to discover processes, analyse performance, and check conformance of medical treatment processes and healthcare organisation processes [28]. Another application of process mining is to compare the behaviour of processes between healthcare organisations. For instance, Partington et al. [30] describe approaches to performing comparative analysis using process mining for cohorts of patients suffering chest pains in four Australian hospitals. In Australia, several studies analysed the behaviour of healthcare processes in pre-hospital, emergency and in-hospital using routinely collected hospital datasets from South Australia and Queensland (e.g., [30, 31, 32]).

Attempts to apply process mining techniques in healthcare has highlighted several challenges and opportunities due to the unique characteristics of these healthcare processes [33, 34, 28, 1]. Mans et al. [33] address three issues affecting process mining in healthcare, namely data correlation from multiple healthcare systems, typical questions of interest for healthcare stakeholders, and identification of data quality issues. Two of these questions are of particular relevance to this work: how to identify (1) the most-followed and exceptional paths and (2) key differences in care paths followed by different patient groups. In [34], the authors review 37 studies in which process mining was applied to clinical pathways. The studies are classified according to whether they attempt to (i) discover actual execution pathways of different clinical pathways (process discovery), (ii) analyse variants of execution pathways, or (iii) evaluate and improve clinical pathways. The authors conclude that at the time of writing, challenges remain such as improving process mining algorithms so that they are efficient enough to deal with the unstructured processes (clinical pathways) and are able to discover models from which a good explanation of the variants can be obtained. Rojas et al., [28] also conducted a literature review of 74 research articles on process mining in healthcare and found that there is a lack of good visualisations of process models for complex and less-structured healthcare processes. Recently, Martin et al. [1] recommended ten key actions for process mining researchers based on the discussions of 18 experts; both researchers and healthcare practitioners. This work incorporates several of these recommendations when developing new analysis techniques: present the unique value proposistion (RC-2), starting from real-world healthcare problems (RC-3), taking into account healthcare specifics during technique development (RC-5), expressing the trustworthiness of output (RC-6), providing a holistic process view (RC-7), and developing multi-perspective approaches (RC-8).

Process mining thus enables data-driven process improvements in healthcare. However, there still remains limited uptake of process mining in healthcare organisations [1]. In particular, the potential of process mining for health economic decision making remains unexplored.

2.3. Cost Considerations in Process Mining

Although the primary focus of data-driven process improvement tends to be time-based, there are several works that explore how cost estimations and predictions can be achieved using process mining techniques.

In [35], the authors proposed a generic framework to associate cost information with event data by developing a cost model, annotating logs with cost functions and then undertaking cost predictions using a transition system. Low et al [36] then present a genetic algorithm with heuristics to generate alternative process executions and compares these execution costs based on a given cost model.

In [37], the authors propose process model-enhanced cost, and cost prediction based on production volume and time prediction for manufacturing processes. In [38, 39, 40], the authors extended the approach of [35] to process model notations and proposed a context-aware cost-data analysis approach using process mining. Their approach enables modelling of cost information at process model and activity levels. In [39], the selection of a classification algorithm to associate cost information based on context together with a case study using the process of a maternity department is presented.

In addition to linking cost considerations with process mining, several works present approaches to consider cost implementations during process simulation, process execution and process monitoring.

In [41], the authors present a novel framework for cost-informed process execution together with data requirements and technical challenges to support it. In [42], the authors describe a casual effect estimation technique to determine the intervention options that reduce the cycle time of a case while maximising the total net gain. The total net gain is considered as the sum of the differences between the benefit of each intervention and its cost.

3. Case Context

In this section, we discuss the context of the particular healthcare decision analysis context that we consider in the remainder of this paper: we describe particular analysis questions and the available data.

The overarching research question for the study is: *How can we design process mining methods that can learn or inform simple decision-analytic model structures and parameters, from routinely collected health data?*. Specifically, we are interested in answering the following questions.

- How can we link the patient pathways with observed costs and other outcomes attributes for a given population?
- How can we mine and visualise existing or *in situ* patient pathways to describe a 'status quo' comparator, from which any potential new or reconfigured service may be adapted and compared?
- How can we estimate the costs of specific events within existing patient pathways, from data recorded at an episode- or case-level and information held within the process model?
- Are the mined outputs informative for model builders, given how they would usually access similar information?

First, we put together an event log from routinely collected data from a Local Health Network, for a population of those who present to a hospital and receive an initial diagnosis of chest pain (ICD10 code R07). These patients are subsequently treated through the Emergency Department and may be admitted and receive additional acute and sub-acute care. For the accuracy experiments, we also used further diagnoses.

We subsequently experimented with developing new process mining methods to address the research questions. We evaluated the implementation and interpretation of models using quantitative validation, and interviews with economic decision analysts.

The data used for this study was limited to what is routinely collected by hospital and health systems, to manage and report on process behaviour. Routinely collected data means that it is the same data attributes available in every Australian hospital. As outlined in Figure 1, there were five datasets linked together deterministically using anonymised patient identifiers: `Emergency Department Data Collection`, `Admitted Patient Data Collection`, `Ward Transaction Data`, `Patient Costing Data Collection` and `Mortality Data Collection`. The data spans patient journeys across the Emergency Department and admitted care, whether they are emergency or elective, and acute or sub/non-acute in nature. Because episodes of care can reflect administrative changes, without a patient actually leaving the hospital, these multiple 'episodes of care' were linked together into a single journey where they were found to have events with adjoining date and time stamps.

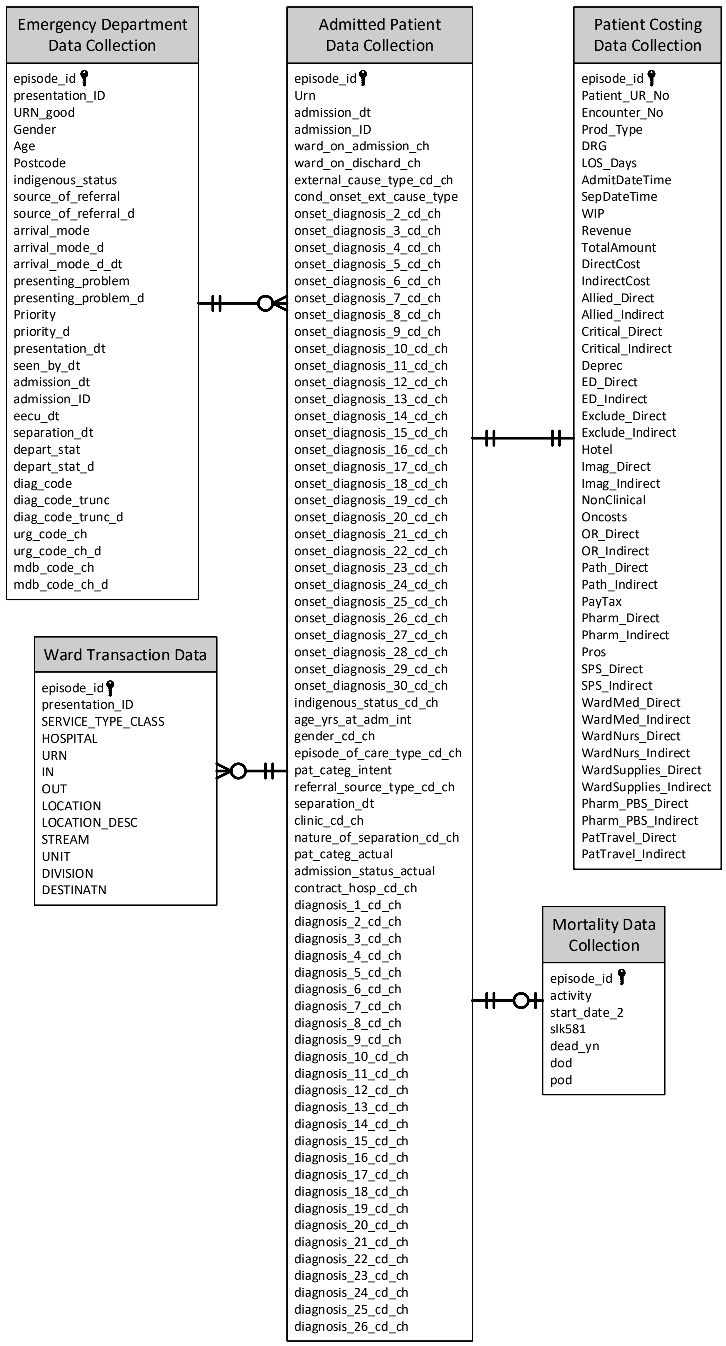

Figure 1: Entity Relationship Diagram of the linked data used within this study, and commonly available within Australian Local Health Networks

The data is transformed into an event log adhering to the XES standard using the FilterTree tool[1] and rather standard conversion steps, including a selection of to-be kept activities with pre-fix `service type`. The cost is the sum of the event attributes `allieddirect`, `wardmeddirect` and `wardnursdirect`, which are first lifted to the trace level. Empty and zero costs were removed, such that the new cost-based tools ignore these traces; a future study might focus on the sub-set of traces without attached costs. The resulting event log has 16 149 traces, 34 025 events and describes 24 different types of activities.

4. Process Mining for Decision Analytics

In this section, we introduce a new type of model that combines trace attribute data with a process specification such that decision analytics that are aware of processes can be performed. We first introduce pre-existing concepts. Second, we introduce a new hybrid formalism for process models and trace attribute data. Third, we describe how cumulative trace attributes (e.g. costs) can be attributed to the execution of activities.

4.1. Preliminaries

An *event* denotes the execution of a process step (activity). A *trace* is a sequence of events that bring a case through a process. A trace can be annotated with attributes indicating properties of the trace. We denote the set of all activities with Σ and the set of all traces with $\mathbb{T}$. For instance, $\langle \text{triage}, \text{ED}, \text{discharge} \rangle^{\text{arrival:10-02-2022 11:45, diagnosis:R07}}$ is a trace consisting of

[1] See `http://leemans.ch/filtertree`.

3 events, annotated with 2 attributes[2]. A *language* is a possibly infinite set of traces. A*process model* expresses a language and an *event log* expresses a stochastic language with a finite number of traces.

Tree-based decision models that represent care processes and patient prognoses are one of the simplest, but also most widely used implementations of a decision-analytic model [43]. They consist of decision and chance nodes; branches that reflect the divergent of pathways for cohorts of patients; and at the termination of each branch, the cumulative resource inputs and associated outcomes. Such models are inherently untimed, or rather, the represented pathways occur over an instantaneous discrete period. Time is only characterised within the model, where it may form a node from which separate 'time related' branches are defined (e.g., slow vs. fast). Time can otherwise be included as one of the accumulated effects at the end of a branch e.g., total length of stay.

Branches typically follows natural disease progression or diagnostic/treatment processes, but requires some level of abstraction by an analyst to define the most relevant branches. Ideally, this definition should be based on how a branch would be expected to have an effect on costs and outcomes, but may sometimes also reflect the available data. In the same way that process models can be 'spaghetti-like', so too can decision trees be 'bushy'. This is particularly the case with processes that occur over a long period of time.

[2]note that event and log attributes have been defined as well, however in this paper we will not require these. Event attributes can be lifted to trace attributes if necessary.

4.2. Process Models for Decision Analytics

In this section, we introduce a new hybrid of economic decision trees and process models: *process models for decision analytics* (PMDA), mimicking existing economic decision tree models based on mined process behaviour [44]. We first introduce PMDAs formally, after which we discuss how they can be discovered automatically from event logs, checked for conformance and visualised.

4.2.1. PMDA*s*

Intuitively, a PMDA is a prefix tree in which each node represents a set of traces with trace attributes. Each node can either denote the set containing the empty trace (τ), prepend traces of its children with an activity (α), annotate traces of its children with a trace attribute (ρ), or combine the languages of its children ($\times$).

Definition 1 (PMDA syntax)**.** *Let Σ be an alphabet of activities such that $\tau \notin \Sigma$, let $a \in \Sigma$ be an activity, let t be a trace attribute and let v be a value. Then, $\tau \in \mathcal{N}$ is a* PMDA*, with $\tau \notin \Sigma$. Furthermore, let $N_1 \ldots N_x \in \mathcal{N}$ be at least one* PMDA *node. Then, $\alpha(a, N_1, \ldots N_x) \in \mathcal{N}$ is a* PMDA*; $\rho(t, v, N_1, \ldots N_x) \in \mathcal{N}$ is a* PMDA*; and $\times(N_1, \ldots N_x) \in \mathcal{N}$ is a* PMDA*.*

Next, we define the language of PMDAs recursively:

Definition 2 (PMDA semantics)**.**

$$
\begin{aligned}
\mathcal{L}(\tau) &= \{\langle\ \rangle\} \\
\mathcal{L}(\times(N_1, \ldots N_x)) &= \{\sigma^T | \sigma^T \in N_i\} \\
\mathcal{L}(\alpha(a, N_1, \ldots N_x)) &= \{\langle a, a_1, \ldots a_y\rangle^T \mid \langle a_1 \ldots a_y\rangle^T \in N_i\} \\
\mathcal{L}(\rho(t, v, N_1, \ldots N_x)) &= \{\sigma^{T\{t:v\}} | \sigma^T \in N_i\}
\end{aligned}
$$

Notice that $\times$ combines the languages of its children without adding any behaviour, which is necessary if the root node of a PMDA should describe a language in which some traces differ in their first activity and have completely disjoint trace attributes or values.

As a consequence of the technicalities in this definition, trace variables are overwritten by nodes higher in the tree structure. For instance, the language of the PMDA $\rho(c, 100, \rho(c, 200, \tau))$ is $\{\{\langle\rangle^{c:100}\}$. Furthermore, the trace attributes employ an open-world assumption: a PMDA expresses the trace attributes a trace should have, but not necessarily all trace attributes a trace should have.

For readability, we might omit unambiguous τ nodes in this paper. For instance, $\alpha(a, \alpha(b), \rho(c, 100))$ is a shorthand for the PMDA $\alpha(a, \alpha(b, \tau)$, $\rho(c, 100, \tau)$ and its language is $\{\langle a, b\rangle, \langle a\rangle^{c:100}\}$.

Thus, a PMDA expresses a finite language of traces, where traces can be annotated with attributes. It is not possible to denote the empty language.

Obviously, when disregarding trace attributes and ρ nodes, a deterministic lexicographically sorted PMDA consisting of τ and α nodes is language unique [45], that is, there is only one such PMDA expressing that particular language (up to sorting of child nodes).

Notice that PMDAs differ from standard process modelling formalisms such as Petri nets [46], BPMN [47] and process trees [45], as these formalisms cannot express constraints on trace attributes. A recent related approach are context-aware process trees [48], which allow for the expression of event-based constraints in process trees. PMDAs have a different purpose: to align as close as possible with current practice in economic modelling. Consequently, the control-flow constructs of PMDAs are deliberately limited to sequence and exclusive choice, which is a sub-set of context-aware process trees that also support loops and concurrency. Furthermore, where context-aware process trees support constraints on an activity or event basis (e.g. execution of activities can tied to a particular resource or time), PMDAs express trace-based constraints (e.g. a trace corresponds to a patient of a certain age).

4.2.2. PMDA *discovery*

To automatically convert a log into a PMDA, one simply builds up a prefix tree of the behaviour in the event log using τ and α nodes.

To filter noise and to reduce the number of nodes for readability purposes, existing filter strategies could be adapted. For instance, adapting [49], a PMDA is constructed while keeping track of how often each node is visited, that is, the number of traces that visit the node. All traces are removed that traverse an arbitrarily chosen node that is visited by the least number of traces, and a new PMDA is constructed. This process is repeated until removing another node would drop the total number of supported traces below a certain user-selectable threshold.

Interactively, a user can indicate where ρ nodes should be added to distinguish between groups of traces explicitly. These nodes *increase* the total number of nodes in the tree, but allow for explicit modelling of groups of traces. The resulting tree could be further adjusted to adhere to recommendations from e.g. medical decision analysis [44], such as balancedness and symmetry.

4.2.3. PMDA *conformance checking*

In this section, we describe how conformance of a trace $\langle a_1, \ldots a_x \rangle^T$ with respect to a PMDA can be determined on two levels: on the trace level (i.e. a boolean answer indicating *whether* the trace fits the PMDA), or on the event level (i.e. a fraction answer indicating *how much* of the trace fits the PMDA).

On the trace level, we consider the nodes of the PMDA to be consumers of either events (α) or trace attributes (ρ). To support the overwriting semantics of trace attributes, we introduce a special value $\perp$, which indicates that a trace attribute has been overwritten by a higher-up node. Then, whether a trace σ adheres to a PMDA E (denoted by $\sigma \models E$) can be computed by a recursive function:

$$\begin{aligned} \tau \models \langle\rangle^T &\equiv \text{true} \\ \tau \models \langle a_1, \ldots a_y \rangle^T &\equiv \text{false} \\ \times(N_1, \ldots N_x) \models \sigma^T &\equiv \exists_{1 \leq i \leq x} N_i \models \sigma^T \\ \alpha(a, N_1, \ldots N_x) \models \langle a_1, \ldots a_y \rangle^T &\equiv a = a_1 \wedge \exists_{1 \leq i \leq y} N_i \models \langle a_2 \ldots a_y \rangle^T \\ \rho(t, v, N_1, \ldots N_x) \models \sigma^T &\equiv (T(t) = v \vee T(t) = \perp) \wedge \exists_{1 \leq i \leq x} N_i \models \sigma^{T\{t:\perp\}} \end{aligned}$$

This function is linear in the size of the PMDA, and linear in the size of the trace, however exponential in the non-determinism of the PMDA.

On the event level, the question to what degree a trace σ^T fits a PMDA E is an optimisation problem, analogous to alignments [50]. To solve this, a trace $\sigma'^{T'} \in \mathcal{L}(E)$ must be found, such that σ^T can be transformed into $\sigma'^{T'}$ using a minimal budget of the following edit operations:

	σ^T	$\sigma'^{T'}$	cost
synchronous move	a	a	0
log move	a	-	1
model move	-	a	1
attribute equal	$T(v) = t$	$T'(v) = t$	0
attribute from log	$T(v) = t$	$T'(v)$ undefined	0
attribute from model	$T(v)$ undefined	$T'(v) = t$	1
attribute revalue	$T(v) = t$	$T'(v) \neq t$	1

For instance, to transform the log trace $\langle p, q\rangle^{c:100,d:50}$ into the PMDA trace $\langle q, p\rangle^{b:100,d:70}$, we would need the following edit operations:

log σ^T	p	q	-	-	$c : 100$	$d : 50$
model $\sigma'^{T'}$	-	q	p	$b : 100$	-	$d : 70$
operation	lm	sm	mm	afm	afl	ar
cost	1	0	1	1	0	1

Using such a trace $\sigma'^{T'}$, the conformance of σ to E is then the fraction of synchronous moves and attribute equal operations over the total number of edit operations (except attribute from log). In our example, the event-level conformance would thus be $\frac{1}{5}$.

Furthermore, the edit operations provide detailed information on conformance deviations, that could be utilised to study detailed deviations and

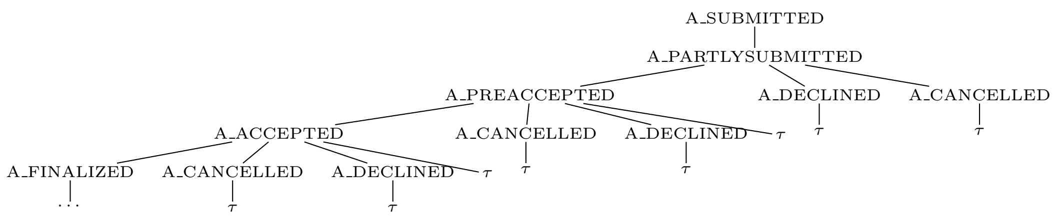

Figure 2: Example PMDA (partial).

their effects on trace attributes.

We conjecture the complexity of the event-level conformance problem to be exponential, like alignments are exponential, thus not adding theoretical complexity.

4.2.4. PMDA *enhancement*

For real-life event logs, unfiltered PMDAs may be too large to be visualised in their entirety. Visualisations of PMDA can utilise the inherent tree structure. For instance, the BPIC12 log could be visualised as shown in Figure 2, in which some of the nodes have been replaced by ... to reduce the size of the tree. Another approach, used in our implementation described in Section 5.1, is to use interaction: initially, only the root node of the PMDA is shown, and the user can expand nodes to explore the behaviour in the PMDA.

4.3. Apportionment of cumulative trace attributes

In some cases, it might be known that the execution of events contribute to a cumulative trace attribute, but in the data only the cumulative trace attribute is present. For instance, a total cost or total service time might be known per trace, while it is desirable to know how each activity contributes to the cost or service times.

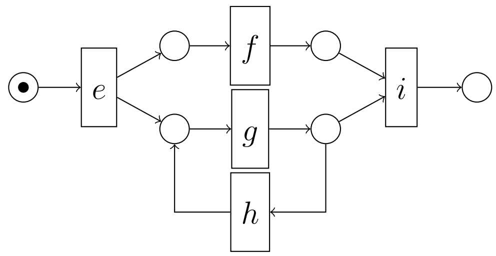

Figure 3: A Petri net.

In this section, we introduce a technique that estimates how activities contribute to cumulative trace attributes. We first provide two examples of such attribute models, and second describe how the parameters of these attribute models can be estimated. Notice that this method is designed for PMDAS, but can be applied, and has been implemented (see Section 5), for other process model formalisms as well.

4.3.1. Attribute model: cost

Our first example attribute model describes the cost of executing a trace in terms of the execution of activities and the time this takes. We assume an alphabet of activities Σ to be given (which are typically the labels of the transitions in a process model; alternatively, the transitions themselves can be used). Let n denote the number of activities in Σ.

Each such activity $a^i \in \Sigma$ yields two parameters: one indicating the set-up cost p_s^i of executing a^i and one indicating the cost per time unit (ms) p_u^i of executing a^i.

We denote the observations from the log L as follows: given a trace $t \in L$, let $o_s^{i,t}$ denote the number of times that activity a^i was executed in t, and let $o_u^{i,t}$ denote the total time that all executions of the activity a^i took in t

together. For instance, for this time the sojourn time can be used, which is the time between an event (of activity a^i) and the preceding sequential event (of any activity) [51].

Estimating the parameters of a cost model can then be translated to a system of linear equalities, using a per-trace error variable ϵ^t. The objective is then to find assignments for these parameters that minimise the sum of absolute values over the ϵ^t of all traces (1). Additionally, we require that all parameters are positive (3):

$$\text{Minimise} \sum_{t \in L} |\epsilon^t| \tag{1}$$

$$\text{such that } \forall_{t \in L} \sum_{1 \leq i \leq n} (o_s^{i,t} p_s^i) + \sum_{1 \leq i \leq n} (o_u^{i,t} p_u^i) + \epsilon^t = o_c^t \tag{2}$$

$$\forall_{1 \leq i \leq n} p_s^i \geq 0 \wedge p_u^i \geq 0 \tag{3}$$

4.3.2. Attribute model: performance & deviations

Given a process model, the attribute model introduced in the previous section can be extended to include conformance and performance information.

Using alignments [50], traces are adjusted to fit a model using synchronous moves, log moves and model moves (see Section 4.2.3). Intuitively, these moves correspond to deviations, and might incur additional costs. For instance, a model move means that an activity of the model is skipped, which might incur a direct or indirect cost. Similarly, a log move means that an activity was executed once more than the model prescribed, and this might incur a cost that is different from executing the activity in accordance with the model.

Thus, this attribute model has parameters for each *transition* for set-up costs of synchronous moves, for each *activity* for the cost of log moves, and for each *transition* for the cost of model moves. Additionally, the cost of service time per time unit (ms) of synchronous moves are parameters.

For instance, consider the Petri net of Figure 3 and the first trace $\langle a, e, g, h, f, h, g, i\rangle^{\text{total:cost:1 000}}$. An optimal alignment is $\langle\begin{smallmatrix} a & e & g & h & f & - & h & g & i \\ - & e & g & h & f & g & h & g & i\end{smallmatrix}\rangle$. This yields the following observations of the parameters for this trace (service time parameters have been excluded from this example for brevity):

sync move transition e	1	sync move transition f	1
sync move transition g	2	sync move transition h	2
sync move transition i	1	model move transition e	0
model move transition f	0	model move transition g	1
model move transition h	0	model move transition i	0
log move activity a	1	log move activity e	0
log move activity f	0	log move activity g	0
log move activity h	0	log move activity i	0
total cost	1 000		

4.3.3. Solving

Both example models were expressed as a system of linear equalities, with an objective to minimise the sum of absolute errors per trace, that is, the sum of $|\epsilon^t|$ for all traces t. This objective function is not linear, due to the absolute value. In this section, we transform this objective function to a linear function, such that it can be solved using linear programming in polynomial time by standard tools.

To this end, we add another variable (σ_i) for each trace t_i, which denotes the cumulative sum of absolute errors for all traces "before" t_i according to some arbitrary order of traces in the log L. For the "first" trace t_1, as we will be minimising all σ_is, we can require that $|\epsilon_1| \leq \sigma_1$. This can be translated to the equivalent $\epsilon_1 \leq \sigma_1 \wedge -\epsilon_1 \leq \sigma_1$ (5)(6). For any non-first trace t_i, the requirement for σ_i depend on the "previous" trace t_{i-1}: $|\epsilon_i| + \sigma_{i-1} = \sigma_i$; this can be translated to linear inequalities (7)(8). As an objective function, we can then simply minimise σ_1 (4):

$$\text{Minimise } \sigma_1 \tag{4}$$

$$\text{Such that (2) and (3)}$$

$$\text{and } \epsilon_1 - \sigma_1 \leq 0 \tag{5}$$

$$-\epsilon_1 - \sigma_1 \leq 0 \tag{6}$$

$$\forall_{\text{non-first traces } i} \epsilon_i - \sigma_i + \sigma_{i-1} \leq 0 \tag{7}$$

$$\forall_{\text{non-first traces } i} - \epsilon_i - \sigma_i + \sigma_{i-1} \leq 0 \tag{8}$$

This is a linear programming problem with $l*2+n$ variables and $l*3+n$ constraints, which can be solved in polynomial time by standard tools.

5. Evaluation

In this section, we show the feasibility of the approach using two implementations, the accuracy of the discovered cost models, and its applicability using semi-structured interviews with domain experts who work in financial and economic modelling within hospitals.

5.1. Implementation

The techniques introduced in Section 4 have been implemented in two contexts: as a new PMDA-based application and as part of the Visual Miner [51, 52].

5.1.1. PMDA *View*

The new PMDA View software leverages the structure and hierarchy of PMDAs to mimic economic decision tree models. As shown in Figure 4, on

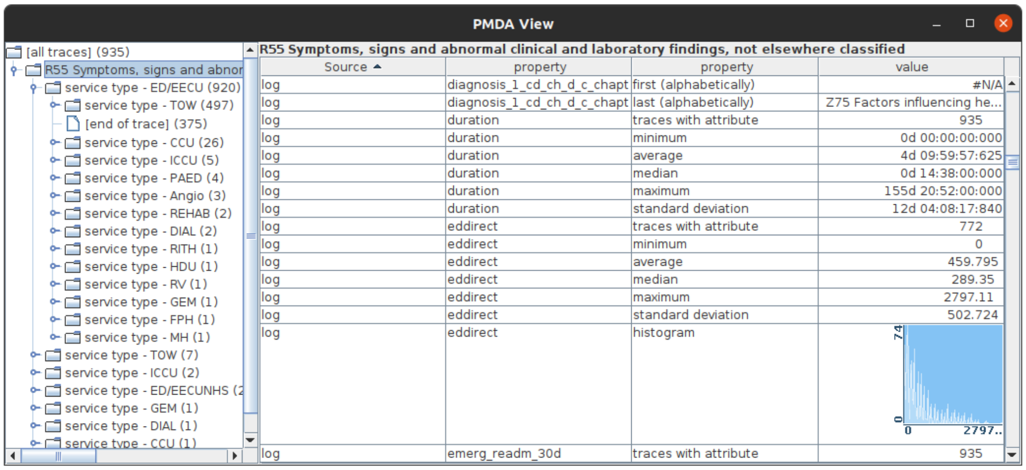

Figure 4: PMDA View: select a PMDA node on the left and see all its trace attributes summarised on the right.

the left there is an automatically discovered unfiltered PMDA. The user can expand each node and dig deeper into the behaviour of the log as required. Each node of the PMDA represents a collection of traces, thus, when the user selects a node on the left, the right side of the PMDA View shows a summary of the trace attributes present in that collection.

Additionally, the trace attribute model described in Section 4.3.1 has been integrated: a cost model on the activities is computed automatically (for the entire log), and the modelled cost is shown (for the selected PMDA node) as well as the actual cost (for the selected PMDA node).

The PMDA View software allows for a quick starter in decision-analytic modelling: given a PMDA, the distributions and other summative information on trace attributes can be quickly inspected and compared between sub-groups of traces. PMDA View and its source code is available from `https://svn.win.tue.nl/repos/ProM/Packages/SanderLeemans`.

5.1.2. Visual Miner

The Visual Miner is an open-source plug-in of the ProM framework [53], which provides an end-user friendly tool to analyse event logs using several process mining techniques, which are applied in sequence without any required user intervention: first, a process model (either a process tree or a Directly Follows Model [49]) is discovered from the log. Second, the model is aligned with the log [50]. Third, the model is annotated with performance or conformance information, after which users can inspect data attributes or drill down into the process in the log by means of filters.

The Visual Miner has been extended with several variants of the trace attribute model described in Section 4.3.2. That is, without necessary interaction of users, the Visual Miner will construct a trace attribute model (where a user can choose one of three options, including conformance and performance), estimate its parameters, and visualise the results in tabular form (Figure 5a) or on the shown model (Figure 5b).

5.2. Cost Accuracy

In this section, we evaluate the accuracy of the cost prediction on synthetic data and our case context.

5.2.1. Synthetic Data

For the synthetic data experiment, we first construct a set of 10 process models, consisting of activities, organised using sequential, exclusive choice, loop and concurrent relations. For each activity, we create a distribution using a randomly generated parameter and a distribution type of {constant, normal, triangular, gamma, log-normal ($\sigma = 0.5$) and Weibull ($k = 5$)}.

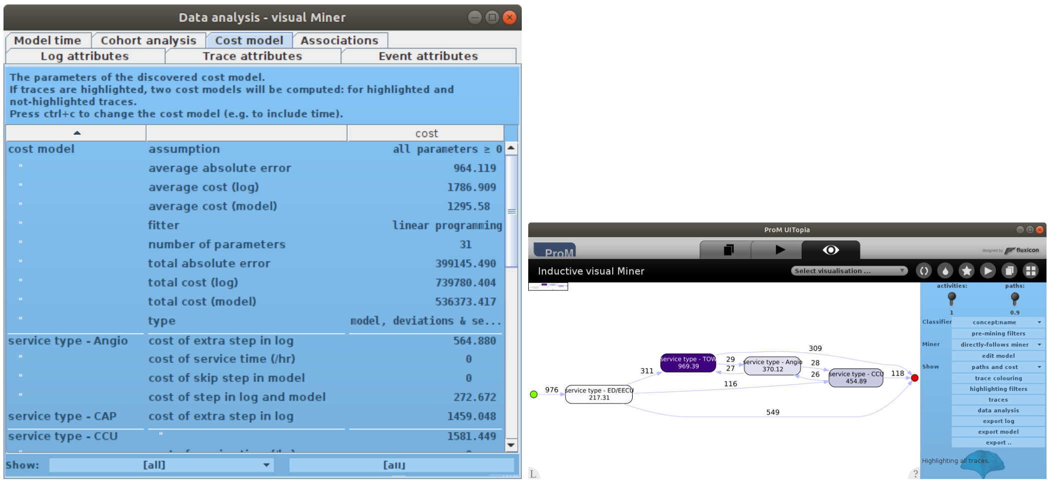

(a) Model in tabular form. Generic info and the models' parameters are shown.

(b) Parameters visualised. The numbers indicate the set-up cost of each activity.

Figure 5: Attribute models in the Visual Miner.

Then, we generate a training log with 1 000 traces, where every event incurs a cost, randomly drawn from the created distribution. The traces are annotated with the sum of the event-based costs. For each of these logs, we obtain a cost model as described in Section 4.3. To measure the predictive quality, we generate a test log of 1 000 traces – with a different random seed as the training log. Next, we let the discovered cost model predict the cost of each trace in the training log; we report on the median cost in log and model (Figure 6a), and average absolute error made (Figure 6b). The entire procedure is repeated 10 times to nullify random effects; the experiment code is available at `https://svn.win.tue.nl/repos/prom/Packages/SanderLeemans/Trunk/src/svn48healthcare/`.

The results shown in Figure 6 indicate that the constant and normal distributions pose little challenge for our method, whereas using the log-normal distribution involves a roughly 20% median error (indicated by the dotted line); the gamma distribution, the Weibull distribution (with $k = 5$) and the triangular distributions again pose fewer challenges. In practice, an additive cost distribution must be non-negative [54], however the results of our method are perfectly interpretable with negative costs as well: the cost can be attributed to a activity executions; and some activities may make traces *cheaper* overall. The gamma and log-normal distributions are commonly used to represent stochastic variability in healthcare costing data. Both of these distributions can be highly positively skewed, reflective of what is commonly found in costing data [54]. Some control-flow constructs are challenging for our cost approximation approach. For instance, if two activities always occur together, then there is no evidence available on their individual contribution to cost. Furthermore, other similar model structures and dependencies may pose similar challenges.

5.2.2. Real-Life Data Sets from our Context

Next, we evaluate the accuracy of the cost estimations. We do this on five logs of our case context with the most-occurring diagnoses: R07 - pain in throat and chest; R55 - syncope and collapse; R10 - abdominal and pelvic pain; L03 - cellulitis and acute lymphangitis; and I50 - heart failure.

All logs are summarised in Table 1. Note that the logs of our case context do not have uniform distributions: most of the traces are short, but extreme cases of over 50 events also occur. Diagnoses may have different levels of granularity, due to their reliance on bio-chemical or psycho-social indicators. For instance, the base-line clinical condition and sub-processes for cohorts such as R07 are more homogeneous than R55. To this end, we assess the

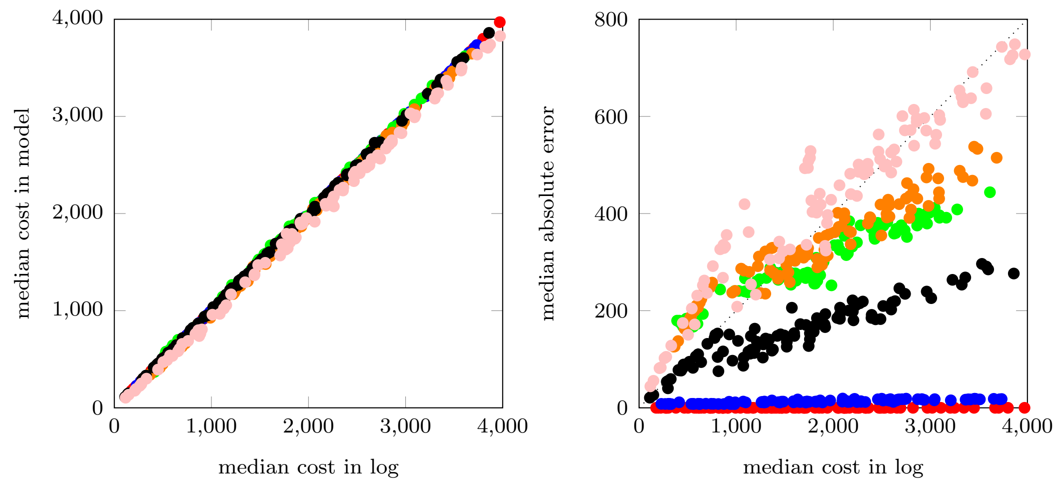

(a) Log vs. model.

(b) Log cost vs. median absolute error.

Figure 6: Synthetic experimental data with constant, normal, triangular, gamma, log-normal and Weibull ($k = 5$) distributions.

accuracy of cost models using a 5-fold cross validation: the data is trace-wise split into 2 buckets, and on each bucket a cost model is learned that is validated on the other bucket. To avoid randomised effects, the procedure is repeated 50 times, for a total of 500 combinations of training and test log.

For each such training log, a cost model was obtained as described in Section 4.3; the accuracy was measured on the corresponding test log as (i) the median cost according to the test log, (ii) the median cost according to the discovered cost model, and (iii) the median absolute error in cost per trace of the model vs. the log.

Figure 7 shows the results; each dot denotes a single cost model. Figure 7a shows that, *on average*, the cost estimations can get close: R07, R55 and R10 can provide pretty accurate cost modelling, while for L03 there is a tendency to under-estimate the costs. The median absolute error *per-trace*, divided by

Table 1: Log characteristics.

Diagnosis	traces	events	activities
R07	976	1 648	19
R55	935	1 735	20
R10	799	1 488	17
L03	363	853	17
I50	368	922	15

the median log cost, shown in Figure 7b, hovers around 0.13, as indicated by the dotted line in Figure 7b. This means that per dollar cost in the log, the modelled cost is inaccurate by approximately 13 cents. The relative differences with respect to this line furthermore allow for interesting insights into the differences in costing between these diagnoses. As observed with R55 (blue) in Figure 7, the absolute differences are relatively higher than those of R07 and L03, perhaps due to the greater cohort heterogeneity and variability in terms of e.g. specific tests, treatments or consultations, and their total duration.

All computations were performed once, so time measures can only indicate generic trends. The measures were taken on a 10-year old laptop running an up-to-date Ubuntu installation. Still, the maximum time for discovery of a cost model in this experiment was 601ms.

5.3. Interview feedback from healthcare decision-modellers

In order to test and explore the face-validity and usability of the approach and mined results, we conducted semi-structured interviews with those who

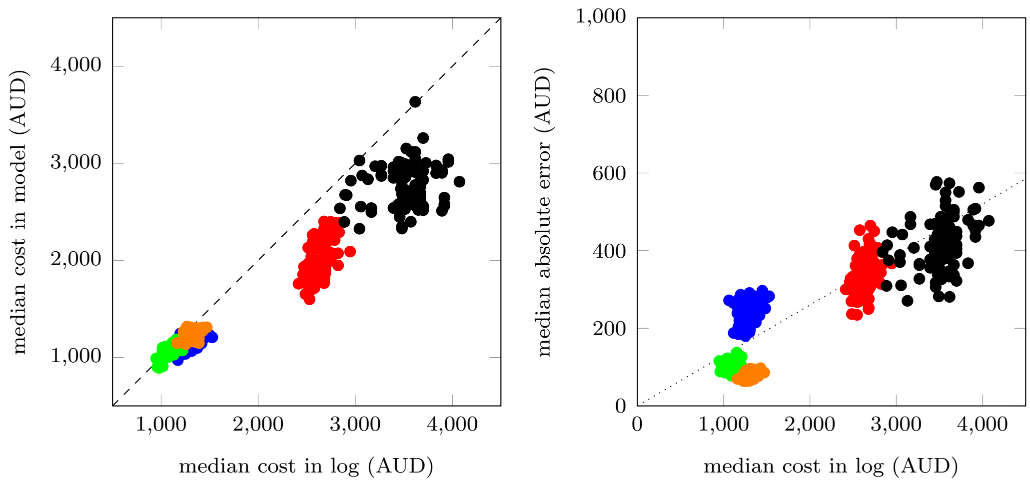

(a) Log vs. model. Dashed line indicates equality.

(b) Log cost vs. median absolute error.

Figure 7: Accuracy of apportionment. Colours indicate logs: R07, R55, R10, L03, I50.

would most likely utilise the information for decision analysis.

Four individual interviews were conducted online during the month of February 2022 with domain experts experienced in financial and decision-analytic modelling. Interview participants were industry consultants, healthcare managers and/or academic researchers with experience in conducting analyses within hospitals, of their services. Participants were recruited using a purposive, maximum variation sampling approach [55], which also included peer-referral to gain different perspectives. Those identified as potential participants were approach via email and provided with participation information and consent forms seeking written, informed consent prior to any data collection.

The interviewees were shown an introduction to the approach and a demonstration of PMDA View, interleaved with questions grouped around three topics: (i) How the outputs could be used, given the context of decision-analyses; (ii) how similar information presented within the PMDA would otherwise be compiled by a decision-modeller; and (iii) thoughts on potential improvements to the methods and graphical presentation/interface. The audio files were transcribed verbatim, and the transcripts thematically analysed. Key feedback from interview participants is summarised below.

5.3.1. How the outputs could be used

Firstly, *"a lot of our research is clinician driven"*[ID2], rather than systemmatically driven by the LHN or broader system, by using performance data. In this context, the emphasis of healthcare modellers is on *"... getting the clinicians to be able to articulate what a relevant comparator or comparators [to a proposed new intervention], would be"*[ID1]. This was clearly a way that the PMDA View was seen as potentially useful.

All of the participants were quick to reflect on their roles and that the existing models used for decision analysis are *"pretty unsophisticated"*[ID1] and sometimes *"back of the envelope kind of stuff"*[ID2], and that that when it comes to using software *"... there's a preference for Excel because it's transparent"*[ID1]. One participant noted *"... transparency is king and if I can see what's going on and and see the logic of that"*[ID1], and that because they require to *"take the decision makers on those journeys with you"*[ID1], that their experience is that they are *"... finding that we lose our audience"*[ID1] when employing more complicated methods. This feedback lent some face-validity to some of our earlier thinking and the premise to the developed PMDA methods. Some direct feedback was that *"I like the transparency of that, so you can sense check [the service model] quite easily"*[ID1].

The overarching feedback was that the presented PMDA and cost estimates would be great for exploring potential relationships and understanding particular measures. It was noted that it could be particularly helpful to engage clinicians to articulate what would be a relevant 'existing care' comparator. Too often clinical and executive stakeholders moved on from understanding their usual care too quickly, because they're trying to push *"a shiny new thing"* that they want.

In addition to setting up for decision-analytic modelling, one participant also raised the before-mentioned idea of conformance to 'value-based reference process'. They noted that *"... in some instances there are going to be clinical guidelines that dictate what should be happening, and so I spend a lot of work in the implementation space and see big differences between what does happen and we ought to be happening. But you can use this to kind of go alright, let's look at that ... see where you need to put more effort in adherence"*[ID1].

5.3.2. How information would otherwise be compiled

Feedback focused on how routine data is often inaccessible. Information on processes and costs metrics are usually translated from other settings or studies; elicited from stakeholders; and synthesised from subjective and anecdotal evidence. Patient pathways may be inferred from data, but not in an automated way, which is readily visualised in PMDA View. This being data-driven using a local source, provides an objective account of existing care pathways, and the statistics for process, cost and outcome metrics associated with these pathways.

Participants described how they usually map out existing care processes and patient pathways, and then *"... it's a bit of an audit against known data* *systems"*[ID1], pointing out that healthcare information systems are *"... not designed to have information easily extacted out of them"*[ID1], and that *"... the information is unidirectional for central reporting and [LHNs] don't get that information back in house to be able to inform decision making"*[ID1]. Further, that *"... we'd have to do a lot of work to kind of get those like those are those model inputs from our from our administrative data sets ... it's a really painful process of getting that information"*[ID2]. One participant also outlined how *"... we obtain [input estimates] from annual reports and published figures"*[ID2], rather than local data sources.

It was noted, particularly by [ID3] that not all of the relevant and necessary information is captured within routinely-collected data. There may be a number of infrequent but impactful and meaningful events that are not observable. So while being data-driven helps to rise above anecdotal evidence, there may be other types of evidence which are still required to form a position on the expected relationship between patient pathways (i.e., processes) and their effects at both individual and aggregate-levels, and which should determine the structure of a decision-tree model. Sometimes the wish of service-planners or decision-makers is to change components of care that are not observable in routine data or the existing process.

5.3.3. Improvements to methods and presentation

The idea of having this information readily available, and as part of a globally accessible tool was very appealing to all of the interview participants. While the detail on descriptive statistics on previous process behaviour, and the ability to interrogate the tree is important for an analyst, the feedback was that the flow-diagram visualisations are helpful and should be retained.

Likewise, the distributions are particularly helpful for quantities such as costs, which can be bi-modal.

The general concern was expressed that linear modelling for the cost estimates may be inaccurate, because there are many non-process attributes such as patient disposition and severity of illness, that may be influential.

Further, the risk was raised that a data-derived model might miss or mis-classify important information, and lead to an observational bias and misunderstanding of the causal logic. Participants reflected on how decision-analyses usually start with a programme-logic or conceptual model and the elicitation of unobserved or poorly-evidenced quantities, which is probably still necessary, but which these process mining outputs could helpfully inform, stating that *"... this is going to get you so far, and then we're going to need to do a bit more of a deep dive into unpacking some of those boxes, in the traditional sense"*[ID1].

5.4. Discussion

The feedback gleaned from the stakeholder interviews has outlined how the PMDA approach may be used by decision analysts working within health, and supports the micro-costing apportionment of trace-level costs to specific events within patient pathways that are represented within routinely collected data. However, the relevance of different discovery, conformance and enhancement methods are in part defined by their contribution towards how value is created within and between organisations [56].

The process mining approach presented within this study uses as an input the routinely-collected administrative, clinical and costing data from episode logs used in activity-based funding and transforms the data into an event log that depicts patient pathways. The presented approach enables the automation, visualisation and description of data within an event log, so that information can be easily abstracted into simple forward-looking models, such as tree-based decision analytic models and Markov models. Furthermore, it would be interesting to extend the PMDA approach to include more data available, such as demographic data or diagnoses made. It is possible that apportionment could be applied to other types of quantities of Patient Reported Outcome and Experience Measures (PROMs and PREMs, respectively) that are accrued along a patient journey. However, significantly greater consideration of confounding and omitted variables bias would be required as these types of outcomes are driven by structural and exogenous factors beyond an observed process within a hospital.

Decision analytics within healthcare is an important end-use focus for process mining, which could help 'close the loop' between applied process mining research and the practice of making decisions about the organisation and the delivery of care. We have illustrated an approach that may readily be taken up by analysts who inform decisions using economic evidence. Whilst there have been efforts previously to link process mining and event-driven simulation modelling methods, the approach presented here represents a first-foray into situating process mining outputs as inputs for decision tree and other simple 'forward looking' decision-analytic approaches. With this applied use case in mind, the lessons learnt can inform future development of further techniques.

The cost apportionment is only as good as the process or PMDA model, and more heterogeneity within cohorts will require more granular event logs that not only track patient movements, as presented in this study, but also the incidence of a range of diagnostic tests, treatments and consultations. While the take-up and quality of electronic records has improved within healthcare, they are often disjointed, and sometimes legacy systems that do not capture timestamps or other features to assist with an integrated process and outcomes perspective of patient pathways [57].

6. Conclusion

This paper presents a suite of process mining techniques that could inform the structure and parameter estimation of decision-analytic models in healthcare. In particular, we have introduced a new hybrid model, *process models for decision analytics* (PMDA), that express processes and emitted trace attributes. While simple, PMDAs enable the study of trace attribute data in combination with processes. We have shown how PMDAs can be discovered from event logs automatically, how their conformance can be assessed and how they can be enhanced with summative supplementary data of event logs. Furthermore, we have shown how cumulative event data attributes (such as cost) can be apportioned to events if only summative trace-level data is available. The PMDA formalism has been implemented in a new tool (PMDA View) that performs the discovery and enhancement steps; the apportionment approach has been implemented as part of the Visual Miner and is part of PMDA View.

Not all changes result in improvements, and not all improvements are worthwhile. To inform the allocation of resources when designing services, outcomes and their opportunity costs (i.e., whether a greater scale of forgone outcomes may have been achieved through a different option using the same resources) must be included within the efficiency gains equation. While we have focused on apportioning trace attributes for costs to events within a process model, the main aim is not to *"solve this for the minimal cost solution"*, as put by one interview participant. Future work could adapt the presented methods to take account of triple [58] or quadruple aims [59] of healthcare to similarly look at apportioning other cumulative numerical trace attributes, such as patient reported outcomes and experience, and potentially even workforce experience and other metrics related to patient pathways.

There will likely always be a need for interpretation of process and outcomes information mined from activity data, which can be quantified through structured elicitation exercises. Elicitation is the method by which opinions are quantified, so as to minimise cognitive biases and enable statistical analyses [60]. Elicitation informed by process mining outputs will be particularly relevant in the case of building upon mined models to explore expected causal effects of changes to patient pathways. Where there are no ledgers of specific unit costs against which to validate the accuracy of cost models, quantitative elicitation exercises could also be used.

From a techniques development perspective, it would be interesting to expand process discovery of PMDAs using machine learning techniques. For instance, GINI-like measures could be leveraged to decide which trace attributes lowers variability of the process "enough" and an ρ should be introduced. Furthermore, precision measures could be defined for PMDAs.

Acknowledgements. This work was partly supported by QUT's Centre for Data Science, and the National Health and Medical Research Council

(NHMRC) Partnership Centre for Health System Sustainability (grant ID: 9100002) administered by the Australian Institute of Health Innovation, Macquarie University.

References

[1] N. Martin, J. D. Weerdt, C. Fernández-Llatas, A. Gal, R. Gatta, G. Ibáñez, O. A. Johnson, F. Mannhardt, L. Marco-Ruiz, S. Mertens, J. Munoz-Gama, F. Seoane, J. Vanthienen, M. T. Wynn, D. B. Boilève, J. Bergs, M. Joosten-Melis, S. Schretlen, B. B. V. Acker, Recommendations for enhancing the usability and understandability of process mining in healthcare, Artif. Intell. Medicine 109 (2020) 101962.

[2] S. J. Leemans, N. Tax, Causal reasoning over control-flow decisions in process models, in: CAiSE, LNCS, Springer, 2021.

[3] M. Shoush, M. Dumas, Prescriptive process monitoring under resource constraints: A causal inference approach, in: J. Munoz-Gama, X. Lu (Eds.), ICPM, LNBIP, Springer, 2022.

[4] A. Rozinat, R. S. Mans, M. Song, W. M. P. van der Aalst, Discovering simulation models, Information Systems 34 (3) (2009) 305–327. doi:10.1016/j.is.2008.09.002.

[5] B. Estrada-Torres, M. Camargo, M. Dumas, L. García-Bañuelos, I. Mahdy, M. Yerokhin, Discovering business process simulation models in the presence of multitasking and availability constraints, Data & Knowledge Engineering 134 (2021) 101897.

[6] M. Wolpert, H. Rutter, Using flawed, uncertain, proximate and sparse (fups) data in the context of complexity: learning from the case of child mental health, BMC Med 16 (1) (2018) 82.

[7] J. Karnon, Alternative decision modelling techniques for the evaluation of health care technologies: Markov processes versus discrete event simulation, Health Economics 12 (10) (2003) 837–848.

[8] A. Wagstaff, A. J. Culyer, Four decades of health economics through a bibliometric lens, J Health Econ 31 (2) (2012) 406–39.

[9] A. Baghbanian, I. Hughes, A. Kebriaei, F. A. Khavarpour, Adaptive decision-making: how australian healthcare managers decide, Aust Health Rev 36 (1) (2012) 49–56.

[10] J. Braithwaite, Changing how we think about healthcare improvement, BMJ 361 (2018) k2014. `doi:10.1136/bmj.k2014`.

[11] K. Skivington, L. Matthews, S. A. Simpson, P. Craig, J. Baird, J. M. Blazeby, K. A. Boyd, N. Craig, D. P. French, E. McIntosh, M. Petticrew, J. Rycroft-Malone, M. White, L. Moore, A new framework for developing and evaluating complex interventions: update of medical research council guidance, BMJ 374 (2021) n2061.

[12] G. F. Moore, S. Audrey, M. Barker, L. Bond, C. Bonell, W. Hardeman, L. Moore, A. O'Cathain, T. Tinati, D. Wight, J. Baird, Process evaluation of complex interventions: Medical research council guidance, BMJ : British Medical Journal 350 (2015).

[13] R. Urquhart, C. Kendell, A. Folkes, T. Reiman, E. Grunfeld, G. Porter, Factors influencing middle managers' commitment to the implementation of innovations in cancer care, J Health Serv Res Policy (2018) 1355819618804842.

[14] G. Harvey, A. Kitson, Parihs revisited: from heuristic to integrated framework for the successful implementation of knowledge into practice, Implement Sci 11 (2016) 33.

[15] A. Partington, J. Karnon, It's not the model, it's the way you use it: Exploratory early health economics amid complexity; comment on "problems and promises of health technologies: The role of early health economic modelling", International Journal of Health Policy and Management (2020) –.

[16] G. Merlo, K. Page, J. Ratcliffe, K. Halton, N. Graves, Bridging the gap: Exploring the barriers to using economic evidence in healthcare decision making and strategies for improving uptake, Applied Health Economics and Health Policy 13 (3) (2015) 303–309.

[17] A. Brennan, S. E. Chick, R. Davies, A taxonomy of model structures for economic evaluation of health technologies, Health Economics 15 (12) (2006) 1295–1310. doi:10.1002/hec.1148.

[18] D. Ben-Tovim, J. Filar, P. Hakendorf, S. Qin, C. Thompson, D. Ward, Hospital event simulation model: Arrivals to discharge–design, development and application, Simulation Modelling Practice and Theory 68 (2016) 80–94.

[19] S. Qin, Using a virtual hospital for piloting patient flow decongestion interventions, in: J. Li, S. Wang, S. Qin, X. Li, S. Wang (Eds.), Advanced Data Mining and Applications, Springer International Publishing, pp. 605–616.

[20] J. Lord, S. Willis, J. Eatock, P. Tappenden, M. Trapero-Bertran, A. Miners, C. Crossan, M. Westby, A. Anagnostou, S. Taylor, I. Mavranezouli, D. Wonderling, P. Alderson, F. Ruiz, Economic modelling of diagnostic and treatment pathways in national institute for health and care excellence clinical guidelines: the modelling algorithm pathways in guidelines (mapguide) project, Health Technol Assess 17 (58) (2013) v–vi, 1–192.

[21] J. Karnon, H. Haji Ali Afzali, When to use discrete event simulation (des) for the economic evaluation of health technologies? a review and critique of the costs and benefits of des, Pharmacoeconomics 32 (6) (2014) 547–58.

[22] L. D. Phillips, A theory of requisite decision-models, Acta Psychologica 56 (1-3) (1984) 29–48.

[23] Z. Špacírová, D. Epstein, L. García-Mochón, J. Rovira, A. Olry de Labry Lima, J. Espín, A general framework for classifying costing methods for economic evaluation of health care, The European Journal of Health Economics 21 (4) (2020) 529–542.

[24] X. Xu, C. M. Lazar, J. P. Ruger, Micro-costing in health and medicine: a critical appraisal, Health Economics Review 11 (1) (2021) 1.

[25] D. Husereau, M. Drummond, F. Augustovski, E. de Bekker-Grob, A. H. Briggs, C. Carswell, L. Caulley, N. Chaiyakunapruk, D. Greenberg, E. Loder, J. Mauskopf, C. D. Mullins, S. Petrou, R. F. Pwu, S. Staniszewska, Consolidated health economic evaluation reporting standards (cheers) 2022 explanation and elaboration: A report of the ispor cheers ii good practices task force, Value Health 25 (1) (2022) 10–31.

[26] T. M. Olsson, Comparing top-down and bottom-up costing approaches for economic evaluation within social welfare, The European Journal of Health Economics 12 (5) (2011) 445–453.

[27] I. H. P. Authority, Australian hospital patient costing standards – part 1: Standards – version 4.1, Report (2021).
URL https://www.ihpa.gov.au/publications/australian-hospital-patient-costing-standards-version-41

[28] E. Rojas, J. Munoz-Gama, M. Sepúlveda, D. Capurro, Process mining in healthcare: A literature review, Journal of Biomedical Informatics 61 (2016) 224–236.

[29] P. Homayounfar, Process mining challenges in hospital information systems, in: 2012 Federated Conference on Computer Science and Information Systems (FedCSIS), IEEE, 2012, pp. 1135–1140.

[30] A. Partington, M. T. Wynn, S. Suriadi, C. Ouyang, J. Karnon, Process mining for clinical processes: A comparative analysis of four Aus-

tralian hospitals, ACM Transactions on Management Information Systems (TMIS) 5 (4) (2015) 1–18.

[31] R. Andrews, S. Suriadi, M. T. Wynn, A. H. M. ter Hofstede, S. Rothwell, Improving patient flows at St. Andrew's War Memorial Hospital's emergency department through process mining, in: BPM Cases, Digital Innovation and Business Transformation in Practice., Springer, 2018, pp. 311–333.

[32] R. Andrews, M. T. Wynn, K. Vallmuur, A. H. Ter Hofstede, E. Bosley, A comparative process mining analysis of road trauma patient pathways, International journal of environmental research and public health 17 (10) (2020) 3426.

[33] R. S. Mans, W. M. van der Aalst, R. J. Vanwersch, A. J. Moleman, Process mining in healthcare: Data challenges when answering frequently posed questions, in: Process Support and Knowledge Representation in Health Care, Springer, 2012, pp. 140–153.

[34] W. Yang, Q. Su, Process mining for clinical pathway: Literature review and future directions, in: ICSSSM 2014, IEEE, 2014, pp. 1–5.

[35] M. Wynn, W. Z. Low, A. ter Hofstede, W. Nauta, A framework for cost-aware process management: cost reporting and cost prediction, Journal of Universal Computer Science 20 (3) (2014) 406–430.

[36] W. Z. Low, M. Wynn, A. ter Hofstede, J. De Weerdt, W. van der Aalst, et al., Revising history for cost-informed process improvement, Computing 98 (9) (2016) 895–921.

[37] T. B. Hong Tu, M. Song, Analysis and prediction cost of manufacturing process based on process mining, in: 2016 International Conference on Industrial Engineering, Management Science and Application (ICIMSA), 2016, pp. 1–5. doi:10.1109/ICIMSA.2016.7503993.

[38] D. Thabet, S. A. Ghannouchi, H. H. B. Ghezala, A general solution for business process model extension with cost perspective based on process mining, ICSEA 2016 (2016) 251.

[39] D. Thabet, S. Ayachi Ghannouchi, H. Hajjami Ben Ghezala, A process mining-based solution for business process model extension with cost perspective context-based cost data analysis and case study, in: IFIP International Conference on Computer Information Systems and Industrial Management, Springer, 2018, pp. 434–446.

[40] D. Thabet, S. Ghannouchi, H. Ben Ghezala, Towards context-aware business process cost data analysis including the control-flow perspective-a process mining-based approach, in: International Conference on Intelligent Systems Design and Applications, AISC, Vol. 1181, 2021, pp. 193–204.

[41] M. T. Wynn, H. A. Reijers, M. Adams, C. Ouyang, A. H. ter Hofstede, W. M. van der Aalst, M. Rosemann, Z. Hoque, Cost-informed operational process support, in: International Conference on Conceptual Modeling, Springer, 2013, pp. 174–181.

[42] Z. D. Bozorgi, I. Teinemaa, M. Dumas, M. La Rosa, A. Polyvyanyy, Prescriptive process monitoring for cost-aware cycle time reduction, in:

2021 3rd International Conference on Process Mining (ICPM), IEEE, 2021, pp. 96–103.

[43] M. Drummond, M. J. Sculpher, K. Claxton, G. L. Stoddart, G. W. Torrance, Methods for the economic evaluation of health care programmes, fourth edition. Edition, Oxford medical publications, Oxford University Press, Oxford, United Kingdom ; New York, NY, USA, 2015.

[44] A. S. Detsky, G. Naglie, M. D. Krahn, D. A. Redelmeier, D. Naimark, Primer on medical decision analysis: part 2—building a tree, Medical Decision Making 17 (2) (1997) 126–135.

[45] S. J. J. Leemans, D. Fahland, Information-preserving abstractions of event data in process mining, Knowl. Inf. Syst. 62 (3) (2020) 1143–1197.

[46] W. Reisig, A primer in Petri net design, Springer Compass International, Springer, 1992.

[47] OMG, Business Process Model and Notation (BPMN), Version 2.0.2 (Dec. 2013).
URL http://www.omg.org/spec/BPMN/2.0.2

[48] R. Shraga, A. Gal, D. Schumacher, A. Senderovich, M. Weidlich, Process discovery with context-aware process trees, Inf. Syst. 106 (2022) 101533. doi:10.1016/j.is.2020.101533.

[49] S. J. J. Leemans, E. Poppe, M. T. Wynn, Directly follows-based process mining: Exploration & a case study, in: International Conference on Process Mining, ICPM 2019, Aachen, Germany, June 24-26, 2019, IEEE, 2019, pp. 25–32.

[50] W. M. P. van der Aalst, A. Adriansyah, B. F. van Dongen, Replaying history on process models for conformance checking and performance analysis, WIREs Data Mining Knowl. Discov. 2 (2) (2012) 182–192. doi:10.1002/widm.1045.

[51] S. J. J. Leemans, D. Fahland, W. M. P. van der Aalst, Process and deviation exploration with inductive visual miner, in: L. Limonad, B. Weber (Eds.), Proceedings of the BPM Demo Sessions 2014 Co-located with the 12th International Conference on Business Process Management (BPM 2014), Eindhoven, The Netherlands, September 10, 2014, Vol. 1295 of CEUR Workshop Proceedings, CEUR-WS.org, 2014, p. 46.

[52] S. Leemans, E. Poppe, M. Wynn, Directly follows-based process mining: a tool, in: Proceedings of the ICPM Demo Track 2019 (CEUR Workshop Proceedings, Volume 2374). Vol. 2374., Sun SITE Central Europe, 2019, pp. 9–12.

[53] B. F. van Dongen, A. K. A. de Medeiros, H. M. W. Verbeek, A. J. M. M. Weijters, W. M. P. van der Aalst, The prom framework: A new era in process mining tool support, in: G. Ciardo, P. Darondeau (Eds.), Applications and Theory of Petri Nets 2005, 26th International Conference, ICATPN 2005, Miami, USA, June 20-25, 2005, Proceedings, Vol. 3536 of Lecture Notes in Computer Science, Springer, 2005, pp. 444–454.

[54] A. Briggs, K. Claxton, M. Sculpher, Decision modelling for health economic evaluation, Handbooks in health economic evaluation series, Oxford University Press, Oxford, 2006.

[55] S. Husbands, S. Jowett, P. Barton, J. Coast, How qualitative methods can be used to inform model development, Pharmacoeconomics 35 (6) (2017) 607–612.

[56] N. Martin, D. A. Fischer, G. D. Kerpedzhiev, K. Goel, S. J. J. Leemans, M. Röglinger, W. M. P. van der Aalst, M. Dumas, M. La Rosa, M. T. Wynn, Opportunities and challenges for process mining in organizations: Results of a delphi study, Business & Information Systems Engineering (2021). doi:10.1007/s12599-021-00720-0.

[57] T. Botsis, G. Hartvigsen, F. Chen, C. Weng, Secondary use of ehr: Data quality issues and informatics opportunities, Summit Transl Bioinform 2010 (2010) 1–5.

[58] D. M. Berwick, T. W. Nolan, J. Whittington, The triple aim: care, health, and cost, Health Aff (Millwood) 27 (3) (2008) 759–69.

[59] T. Bodenheimer, C. Sinsky, From triple to quadruple aim: care of the patient requires care of the provider, Ann Fam Med 12 (6) (2014) 573–6.

[60] L. Bojke, M. O. Soares, K. Claxton, A. Colson, A. Fox, C. Jackson, D. Jankovic, A. Morton, L. D. Sharples, A. Taylor, Reference case methods for expert elicitation in health care decision making, Medical Decision Making 0 (0) (2021) 0272989X211028236.