

Process Discovery Using Graph Neural Networks

Dominique Sommers, Vlado Menkovski, Dirk Fahland

Eindhoven University of Technology, Mathematics and Computer Science, Eindhoven, the Netherlands

{d.sommers, v.menkovski, d.fahland}@tue.nl

Abstract—Automatically discovering a process model from an event log is the prime problem in process mining. This task is so far approached as an *unsupervised* learning problem through graph synthesis algorithms. Algorithmic design decisions and heuristics allow for efficiently finding models in a reduced search space. However, design decisions and heuristics are derived from assumptions about how a given behavioral description — an event log — translates into a process model and were not learned from actual models which introduce biases in the solutions. In this paper, we explore the problem of *supervised learning* of a process discovery technique d . We introduce a technique for training an ML-based model d using graph convolutional neural networks; d translates a given input event log into a sound Petri net. We show that training d on synthetically generated pairs of input logs and output models allows d to translate previously unseen synthetic and several real-life event logs into sound, arbitrarily structured models of comparable accuracy and simplicity as existing state of the art techniques in imperative mining. We analyze the limitations of the proposed technique and outline alleys for future work.

Index Terms—Automated process discovery, machine learning, graph neural networks

I. INTRODUCTION

Automated process discovery (APD) is the problem of discovering a process model M from an event log L [1]. State-of-the-art techniques approach APD as an *unsupervised* learning problem trying to achieve pareto-optimality of M regarding *fitness* and *precision* wrt. L , *generalization* wrt. future traces not seen in L yet, and structural *simplicity* of M [2], and further to ensure *soundness* of M [3]. APD of flow-based models, such as Petri nets, is primarily approached algorithmically by synthesizing a graph from behavioral abstractions of L [4]–[6], as optimization problem over linear [7] or logical constraints [8], or genetic algorithms searching for optima in the space of models [2].

Reviews and benchmarks observe that, despite impressive progress, no *unsupervised APD technique* consistently returns fitting, precise, simple, and sound model on all problem instances in feasible time [9], [10]. Specifically, each technique is based on different algorithmic design decisions and uses different heuristics for efficiently finding models in the available search space, resulting in an inherent bias favoring some quality criteria over others that cannot be overcome [10]. These design decisions and heuristics are derived from assumptions about how a given behavioral description — an event log — translates into a process model in an ad-hoc manner and not derived systematically leading to low generalization [11] as we discuss in Sect. II.

In contrast, human modelers train their modeling skills in a *supervised* fashion by learning which model structures are adequate solutions for which behavior, and then apply these skills to build a solution piece-wise along the information provided [12]–[14]. In this paper, we study whether it is possible to design a process discovery technique that more directly emulates a human modeler, i.e. (1) *learn* how to construct models from examples of event logs and corresponding models, and then (2) *transfer* this learned knowledge to construct process models to event logs not seen previously.

We formulate the *supervised process discovery problem*. We want to develop a technique t that can train for given pairs $\langle L_i, M_i \rangle_{i=1}^k$ of (synthetically generated) event logs and corresponding process models a function $d = t(\langle L_i, M_i \rangle_{i=1}^k)$ so that d can translate an (unseen, real-life) event log L into a sound model $M = d(L)$ that has high accuracy wrt. L and is structurally simple. Where supervised learning with artificial negative events [15] learns a process model, this problem is to *learn the discovery technique* itself.

In the following, we limit ourselves to imperative models and propose a first solution to this problem for Petri nets as target modeling language where $d = (f, N_1, \dots, N_k)$ is an algorithm f constructing and updating a graph G using graph convolutional neural networks N_1, \dots, N_k . f first encodes the general translation task from given log L into the space of possible models as a graph G . G encodes the input log L with edges to a template Petri net M having one transition t_a per activity label a in L and candidate places and arcs between these transitions. f then uses N_1, \dots, N_k to select which candidate places shall remain in the model based on the information in L . The N_i update state vectors at event and transition nodes and at edges between nodes of G to propagate behavioral information from the event log to the transitions. The technique t trains N_1, \dots, N_k on $\langle L_i, M_i \rangle_{i=1}^k$ by following an iterative approach based on the process of process modeling observed in human modelers [13] to learn which places of target model M_i shall remain given the structure of the input event log L_i . Through training, the N_1, \dots, N_k learn how to update the weights in G to select places from the available candidates; f uses beam search [16] and S-coverability checks [3] to prune the search space.

We implemented the above technique using DGL <https://docs.dgl.ai/> with PyTorch and demonstrate feasibility reaching state-of-the-art performance on a variety of problem instances. We trained d on a synthetically generated training set of 2000 block-structured process models and corresponding event logs of varying size and complexity. We show that it is able to learn

to rediscover the process models from the class of problems with same representational bias as it has been trained on, achieving high accuracy both wrt. the log and wrt. the data generating process that is comparable to the state-of-the-art techniques. With further evaluation on real-life processes, we show that our technique can effectively generalize to problem instances outside the representational bias of the training data, however for some instances the resulting model may only be easy sound (contain dead parts but no deadlocks). We specifically explore the reasons for the techniques current shortcomings and outline future avenues of research.

II. RELATED WORK

We discuss literature on quality measures, biases and design decisions in APD, modeling as a human task, and graph neural networks.

The quality of a process model M to an event log L is assessed along 5 criteria. *Fitness* is the share of process executions (*traces*) of L that is described or accepted by M ; *precision* is the share of process executions described by M that is not in L . *Alignment-based fitness* and estimating precision via *escaping edges* (only the first steps of M not in L) [17] are most widely used [10]. The monotone precision/recall measures [18] rank models more consistently but are slower to compute. *Generalization* is the likelihood that M accepts another unseen trace from the process that generated L and can only be estimated [17]. *Simplicity* states how clearly M describes the logic or cause-effect relations of the process that generated L and is estimated through graph size and complexity [10]. M is *sound* if every partial execution from the initial state can be extended to a terminating execution in a designated final state [3]. A sound workflow net N can be decomposed into so-called *S-components*, conversely if N cannot be decomposed into *S-components*, i.e., is *not S-coverable* it is not sound. We exploit this property to exclude non-sound solutions.

The fundamental challenge in APD is that the target search space of models is too large to exhaust [2], [7]. Genetic algorithms [2] effectively explore the search space but take too long to find a satisfactory model for practical applications [10]. Time-efficient algorithmic solutions to APD synthesize a graph from behavioral abstractions of the log [4]–[6]. Thereby design decisions and heuristics bias the algorithm regarding fitness, precision, simplicity, and soundness. Enforcing a specific *representational bias* [19] in the problem formulation, e.g., limiting the search space to block-structured models [2], [4], ensures simplicity and soundness, although sound models with high fitness and precision lie outside the chosen representational bias [10]. Algorithmic design decisions can favor or even guarantee solutions of a specific quality criterion at the cost of loss in another criterion, e.g., ensuring fitness [4], [7], [8] lowers precision [8], [10]. Heuristics-based filtering and pattern detection on behavioral abstractions of event data results in models of high fitness and precision [5], [6] which in turn may be unsound and have high complexity [10]; the heuristics may not generalize to new data [20] or larger




Fig. 1. High level overview of the approach.

samples [11]. Techniques relying on behavioral abstractions of event logs [4]–[6], [21] fail when the log contains behaviors not preserved by the abstraction [21], [22]. Techniques avoiding behavioral abstractions solve an optimization problem over linear [7], [23] or logical constraints [8] over the event log that ensures fitness at the cost of precision and soundness [8], [10] or prohibitively high running times [23].

The cognitive process of humans creating models has been studied empirically. Humans create models by iterating three phases: *comprehend* (a chunk of the information about the process), *model* (by adding or removing formal model structures related to the comprehended chunk), and *reconcile* (by reorganizing model layout to better comprehend the created model structures) [13]. Empirical evaluations have shown that breadth-first modeling strategies along the structure of the target model lead to highest precision and recall [12].

Neural networks have not been exploited for APD as of yet. However, they have been developed for various graph problems [24]. Such networks learn representations of graphs, node and edges based on an information propagation process to perform tasks like node classification or link prediction. Generative graph neural networks are developed that learn and represent conditional structures on graphs [25]. The attention mechanism have proven useful in generative tasks that can redistribute the weights of the different inputs [26] and has been exploited in graph neural networks as well [27].

III. DEFINING THE LEARNING TASK

The fundamental task of process discovery is to translate an event log L into a Petri net N (see Fig. 1). In this section, we show how to formally encode this discovery problem in a graph G (Sect. III-B) which contains both the event log L and a candidate Petri net N with candidate places P . We then define the learning task for solving the discovery problem (Sect III-C) which is to select the candidate places from P that most likely explain the behavioral information in L ; the subgraph with the selected candidate places is the resulting model N . In Sect. IV, we explain how to train a function d to learn how to solve this discovery problem in an iterative manner by propagating behavioral information from L to model structures in the candidate Petri net N and select candidate places one-by-one. Fig. 1 summarizes this approach.

A. Preliminaries

We write A for the (finite) set of *activity names* extended with $> \in A$ and $| \in A$ for artificial start and end. A log L

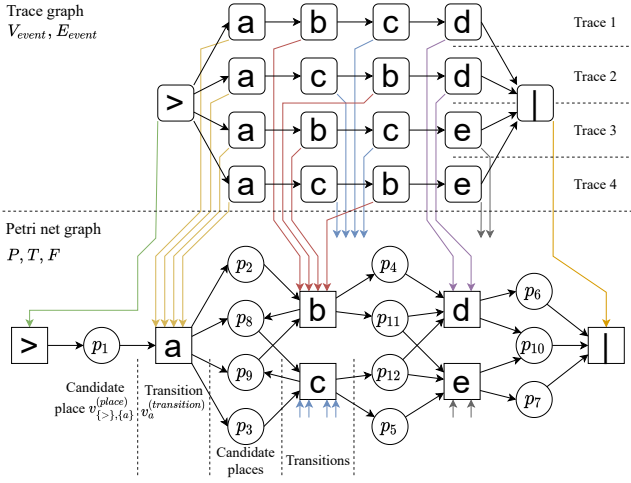


Fig. 2. Encoding of process discovery problem.

(over A) is a finite multiset of traces where a *trace* is a finite sequence $\sigma = \langle \rangle, a_1, \dots, a_n, | \rangle \in A^*$. Each occurrence of an activity $a_i \in \sigma$ is called an *event*.

We recall notations for Petri nets with τ -transitions and refer to [3] for definitions. A Petri net $N = (P, T_A \cup T_\tau, F)$ has places P , transitions $T = T_A \cup T_\tau$ where T_A and T_τ are *visible* and *invisible* transitions (i.e., the T_τ do not occur in firing sequences of N), respectively, and arcs F . In a *workflow net* N , every node of N is on a path (along F) from the unique source place $i \in P$ (no incoming arcs) to the unique sink place o (no outgoing arcs).

B. Encoding the discovery problem in a graph

We encode the problem of translating a given input event log L into a Petri net N as a graph G of 3 parts as illustrated in Fig. 2. (1) We encode L as a *trace graph*. (2) We encode the solution space of all possible models for N as a *candidate Petri net* (over-approximating the required places and arcs). (3) *Links* from event nodes in the trace graph to transitions in the candidate Petri net encode which transition shall describe which event.

To encode L , the *trace graph* (V_{event}, E_{event}) contains one start and one end event node $v_{>}, v_{|} \in V_{event}$, and for each trace $\sigma = \langle \rangle, a_1, \dots, a_n, | \rangle \in L$ the chain of unique event nodes $v_{a_i}^{\sigma, i}, 1 \leq i \leq n$ with directed edges $(v_{a_i}^{\sigma, i}, v_{a_{i+1}}^{\sigma, i+1}), (v_{a_{i+1}}^{\sigma, i+1}, v_{a_i}^{\sigma, i}) \in E_{event}, 1 \leq i \leq n$ and directed edges between $v_{>}$ and $v_{a_1}^{\sigma, 1}$ and between $v_{a_n}^{\sigma, n}$ and $v_{|}$.

The *candidate Petri net graph* (P, T, F) defines a superset of the candidate places P and arcs F needed to describe the behavior in L ; within this superset we search for the target model N . T contains one transition $t_a \in T$ for each activity $a \in A$ in L ($\rangle, | \in A$). Similar to the α -algorithm [21], we fully characterize each candidate place p and its incoming and outgoing arcs by a pair $(X, Y), X, Y \subseteq T$ of input and output transitions. The pair (X, Y) defines place $p_{X, Y}$ with arcs $(t, p_{X, Y}), t \in X, (p_{X, Y}, t), t \in Y$. Within the set of all possible candidate places (and arcs) $P = \{(X, Y) \mid X, Y \subseteq$

$T\}$ we search the Petri net (P', T, F') , with places $P' \subseteq P$ and corresponding arcs $F' \subseteq F$ that best describe L . Pre-processing can reduce this search space, see [28].

Links are edges from event to transition nodes. For each event node $v_{a_i}^{\sigma, i}$ we add a directed edge $(v_{a_i}^{\sigma, i}, t_{a_i})$ to t_{a_i} . Thus, the graph structure encodes which transition nodes shall model which events.

C. Defining the learning task

Graph G essentially defines the learning task of identifying only the subset of P that restrict the transitions T to the behavioral patterns present in L , e.g., choices, parallelism.

Where the α -algorithm constructs the subset of P from L , we train a function d on examples $\langle L_i, N_i \rangle$ to (learn to) estimate the likelihood that a specific place $v_i \in P$ is best suited to be included in N given that places v_1, \dots, v_{i-1} were chosen before, i.e., $p(v_i | v_1, \dots, v_{i-1})$. However, we have to make sure that the entire set P' describes the most likely solution for a specific L , i.e. the joint probability $p(P', \pi | L) = \prod_{i=1}^{|P'|} p(v_i | v_1, \dots, v_{i-1}, L)$ with ordering π . Note that the actual ordering of P' is irrelevant in a graph and therefore ideally $p(P', \pi | L)$ is independent of π . *Marginal joint probability* models this probability over all permutations $\mathcal{P}(P')$ of P' : $p(P' | L) = \sum_{\pi \in \mathcal{P}(P')} p(P', \pi | L)$. $p(P' | L)$ can be used to find $\arg \max P'$ for event log L by sampling those resulting in the highest probability.

IV. APPROACH

Our approach for defining d consists of a set of (graph) neural networks, each taking as input (a part of) G and emulates a specific step in sequentially creating a model from an event log by propagating behavioral information and selecting candidate places (see Fig. 3). These networks together model the joint probability $p(P' | L)$ of selecting only the subset $P' \subseteq P$ needed to describe the behavioral relations between transition nodes available from an event log L .

We do not directly translate from L to $p(P' | L)$ but introduce a latent parameter space to encode and propagate behavioral relations from L to the nodes in N . In Sect. IV-A, we introduce a feature vector h_i for each node $v_i \in V$ of G ; h_i allows to encode how v_i is related to each activity $a \in A$ recorded in L . In IV-A, we also explain how each of the (graph) neural networks either updates h_i to propagate information in G (comprehension/reconciliation) or uses h_i to learn how to maximize $p(P' | L)$ (i.e., select the most likely places P'). Sect IV-B discusses how to train the NNs for maximizing $p(P' | L)$ together based on training data $\langle L_i, N_i \rangle$. Sect IV-C introduces extensions to f to ensure S-coverability and the possibility of incorporating invisible τ -transitions. A more extensive explanation of the approach can be found in [29].

A. Sequential candidate selection

Function d uses 2 decision making neural networks (NNs), “Select candidate” NN, and “Stop” NN, and 2 graph convolutional networks (GCNs), propagation networks PN1 and PN2 as follows. d solves the task of sequentially identifying P' in

two steps: deciding which candidate place to select next (Sect. IV-A2), and deciding when to stop selecting more places (Sect. IV-A3). For both, a regular “decision making” NN is sufficient to model the probability distribution on nodes for selection; see Fig. 3.

However as is stated in above, these NNs have to make their decisions based on behavioral information from L encoded in the trace graph in G . Graph convolutional neural networks (GCNs) can propagate and process information through G to encode this behavioral information in a node’s feature vector h_i (also called *node embedding*). Initially, each node v_i is given a one-hot encoded feature vector $h_i^{(0)}$ of length $|A|$ denoting its activity label concatenated with its frequency in the event log, if available. $h_i^{(0)}$ for candidate places are zero vectors.

GCNs allow to update a node’s feature vector h_i based on the feature vectors of its (indirect) neighbors, taking the surrounding graph structure into account. This allows us to train a GCN PN1 to process and propagate the behavioral information from the event log nodes in G to the place nodes in G such that the NNs can make their decision for selection (Sect. IV-A1).

As the method is sequential, more propagation is required by a second GCN PN2 through G to aggregate the NN’s past decisions as future decisions are conditionally dependent on them (Sect. IV-A4). These four NNs now each have their own smaller learning task which align with the steps in the process of process modeling (see Section II). Fig. 3 shows how they are connected and modify G , which is similar to the graph generation process as proposed in [25]. The following sections go into detail of each NN’s learning task.

1) *Propagation network 1*: We propagate behavioral information from the trace graph to the places by a GCN with K -headed attention mechanism as described in [27]. The number of layers l in a GCN controls how many propagation steps are performed. Each place p in G has at least input transition t_a and one output transition t_b ; to at least propagate behavioral information from the predecessor and successor events of both a and b , PN1 should contain at least three layers (as illustrated in Fig. 4). The attention mechanism can be used to give unequal weights to nodes and is necessary because of the different node classes in G which are not equally important for this initial step.

The result of this step is the same graph with updated node embeddings as computed by the internal weights of the GCNs as follows:

For the first $(l - 1)$ layers of the network the nodes’ embeddings are updated by:

$$h_i^{(l+1)} = \parallel_{k=1}^K \left(\sum_{j \in \mathcal{N}(i)} d_{ij} \begin{bmatrix} \alpha_{ij}^{(l)k} h_j^{(l)} W^{(l)k} \\ \alpha_{ij_r}^{(l)k} h_j^{(l)} W_r^{(l)k} \end{bmatrix} \right) \quad (1)$$

where h_i^l denotes the embedding of node v_i at layer l , $\mathcal{N}(i)$ denotes the *outgoing* neighbors of node v_i and $W^{(l)k}$ is the weight matrix of the GCN at layer l and attention head k . To also aggregate information from *incoming* neighbors, all arcs

in G , except for the links, are made bi-directional, where the vector d_{ij} encodes the direction of the arc: $[1, 0]$ if from v_i to v_j , and $[0, 1]$ if from v_j to v_i . Furthermore, self-loops are added to retain a node’s own information during aggregation. $W_r^{(l)k}$ processes information on these *reversed* arcs.

The update function of the last layer, having a single attention head, is:

$$h_i^{(l+1)} = \text{RELU} \left(\sum_{j \in \mathcal{N}(i)} d_{ij} \begin{bmatrix} \alpha_{ij}^{(l)} h_j^{(l)} W^{(l)} \\ \alpha_{ij_r}^{(l)} h_j^{(l)} W_r^{(l)} \end{bmatrix} \right) \quad (2)$$

with the RELU activation function.

The nodes’ embeddings h_i are used by the GCN to encode the behavioral information and are only interpretable by the “Select candidate” and “Stop” NNs. to be later used for classification of the places. Such information is similar to how a human modeler would aggregate information from the event log and the place’s structural properties to decide whether it fits the process.

2) *“Select candidate” network*: We now determine for each candidate place $v_i \in P \setminus P'$ that was not selected yet, the probability p_i that v_i should be added next, and pick $\arg \max_{v_i \in P \setminus P'} p_i$. v_i is marked as selected by adding a feature value to the embedding.

To compute p_i , we define CN as a regular NN with input vector h_i (as returned by the preceding step) and output a single value p_i . CN learns the weight matrix W to map h_i to p_i as a single fully-connected layer.

$$s_v = h_v W \quad (3)$$

By normalizing all the candidates’ scores using the SOFTMAX function we get a probability p_v for each node:

$$p_v = \text{SOFTMAX}(s_v) = \frac{\exp(s_v)}{\sum_{u \in (P \setminus P')} \exp(s_u)} \quad (4)$$

with $\exp(x) = e^x$.

3) *“Stop” network*: We define the “Stop” network SN for deciding when to stop adding candidate places. SN takes as input all Petri net node embeddings $h_i, v_i \in V \setminus V_{event}$ and has as output a probability p_{add} that another place should be added; we then make the binary decision to stop adding by comparing to a chosen threshold. SN has two layers. Layer 1 aggregates $h_i, v_i \in V \setminus V_{event}$ into a graph embedding h_G by the following equation

$$h_G = \sum_{v \in V \setminus V_{event}} \text{SIGMOID}(h_v W_a) (h_v W_g) \quad (5)$$

with learnable weight matrices W_a and W_g . $\text{SIGMOID}(h_v W_a)$ serves as a “gating function”, how much each node should contribute to h_G by the weights in W_a [25], where $\text{SIGMOID}(x) = \frac{1}{1+e^{-x}}$ is used to map a value from \mathbb{R} to a value between 0 and 1.

Layer 2 computes the probability p_{add} to add more nodes from h_G by:

$$p_{add} = \text{SIGMOID}(h_G W_d) \quad (6)$$




Fig. 3. Detailed overview of the candidate selection.

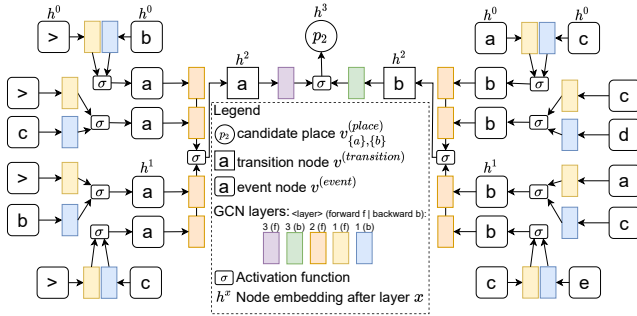


Fig. 4. Partial illustration of information propagation through the graph for node $v_{\{a\},\{b\}}^{(place)}$ from Fig. 2.

with a learnable weight matrix W_d . ($h_C W_d$) is a score which is converted to a probability using the logistic sigmoid function.

If the network's decision is to stop, we are done. Otherwise, we continue at Section IV-A4.

4) *Propagation network 2*: This step in the process is similar as the one in Section IV-A1 where information is propagated through the graph by a GCN with multi-headed attention mechanism and is used to process the previously made decision. It is necessary for the GCN to have at least two layers to ensure that information about the previous decision reaches other candidates. Recall that the decisions are encoded as features in the embeddings. By propagating this information two steps in the graph, d can encode information whether other candidates fit the already selected candidates. The nodes' embeddings are updated in a similar fashion as is described in Section IV-A1, after which we loop back to Section IV-A2 to make the next decision.

B. Training and inference

During training on instances $\langle L_j, N_j \rangle$ with known solutions $N_j = (P_j, T_j, F_j)$, the weight matrices in the NNs are optimized such that a specified loss function is minimized. The loss function to be minimized during training is as follows:

$$l = - \sum_{i=1}^{|P'|} \log p(v_i | v_1, \dots, v_{i-1}) \quad (7)$$

with known places $P_j = P' \subseteq P$ to be selected and $p(v_i | v_1, \dots, v_{i-1})$ the probability of choosing candidate place v_i after having selected candidate places v_1 to v_{i-1} . It is intractable to learn the marginal joint probability because many permutations may exist. A canonical ordering can give a lower bound. For learning, we choose π based on breadth-first search on N_j ; this order is close to canonical. With π and Eq. 7, the model learns the joint distribution $q = p(P', \pi | \mathcal{L})$ of the data by maximizing the expected joint log-likelihood:

$$\mathbb{E}_{p_{data}(P', \pi | \mathcal{L})} [\log q] = \mathbb{E}_{p_{data}(P' | \mathcal{L})} \mathbb{E}_{p_{data}(\pi | P')} [\log q] \quad (8)$$

with $\mathcal{P}' = \{N_j\}_{j=1}^k$ and $\mathcal{L} = \{L_j\}_{j=1}^k$ from a dataset $\langle L_j, N_j \rangle_{j=1}^k$. With a sufficiently large and varying dataset $\mathbb{E}_{p_{data}(P', \pi | \mathcal{L})} [\log q]$ approaches $\mathbb{E}_{p(P' | \mathcal{L})} [\log p(P' | \mathcal{L})]$ which is the true probability distribution of the correct selection of candidates P' from P given an event log L , solving the learning task.

Teacher forcing [30] is used in training to enforce the breadth-first ordering π during the generation process and to correct the prediction after every candidate choice, which addresses slow convergence and instability. Note that this requires a predefined ordering.

During inference—the counterpart of training where d is used on unseen data—teacher forcing is inaccessible and

therefore the sampled candidates are selected by the NNs without correction. A suboptimal choice in the beginning of the generation process can lead to a low joint probability at the end, which is exactly what we want to maximize. To mitigate this, beam search is used; a heuristic search algorithm to find the highest joint probability. Using beam search with a beam width of b , d selects the b candidates with highest conditional probability at each step from b unfinished runs, resulting in a set of b^2 (un)finished runs of which the b runs with highest joint probabilities are taken for the next step. Afterwards, the b Petri nets with highest joint probabilities are returned.

C. Extensions

As illustrated in Fig. 3, extensions on the technique as described above are made to ensure certain properties of the Petri net, also on real-life data. Firstly, we do not add a candidate place v_i chosen by SCN if the net N with v_i is not S-coverable (a polynomial check on the structure of N) as then the resulting net would not be sound [3]; in this case, we exclude v_i and let SCN return the next most probable candidate (see Fig. 3). This ensures reachability of the final marking but cannot avoid dead transitions. For this, a global check for a workflow net structure could be introduced.

Secondly, SN (Sect. IV-A3) may decide to stop while some transitions do not have an incoming or outgoing place (not a workflow net); we check for non-connectedness of N and override the stop decision in this case.

Invisible τ -transitions are needed to model silent skips in real-life event logs while they are not recorded in L . We add τ -transitions T_τ to the candidate nodes (previously P) for the SCN to select. In principle, the set of all possible candidate τ -transitions is $T_\tau = \{(X, Y) | X, Y \subseteq P\}$, however, for feasibility, we only add candidate τ -transitions between already selected candidate places: $T_\tau = \{(X, Y) | X, Y \subseteq P'\}$. Candidate τ -transitions are added to G after selecting a new place and its feature vector is defined as the sum of the feature vectors of its neighboring nodes.

V. EVALUATION

We assess the feasibility of our approach for solving the supervised process discovery problem, i.e., training d on $\langle L_i, N_i \rangle$ synthetic process instances to apply d for discovering models on unseen synthetic and real-life data.

A. Experimental setup

We evaluate wrt. 3 objectives: (O1) Is the loss function of Eq. 7 able to optimize the NNs in d to achieve high precision/recall in selecting the ground truth places on synthetic test data? (O2) Does the loss function also optimize fitness and precision of the discovered model wrt. the input model on synthetic test data compared to state-of-the-art methods? (O3) How does the quality of models compare to the other methods on real-life data, i.e., does the ML program generalize beyond the task it was trained on? For training and testing, we used the PTAndLogGenerator plugin in ProM to generate $\langle L_i, N_i \rangle_{i=1}^{2663}$ problem instances: we generated 2663 random

process trees, converted to block-structured Petri net N_i , and generated corresponding event log L_i by simulation with the following parameters:

- number of nodes: mode=8, min=4, max=15;
- structure probabilities: sequence=0.4, choice=0.32, parallel=0.2, loop=0.08, or=0.0;
- number of silent transitions: 4;
- number of simulated traces per Petri net: 1000.

The hyperparameters for our model d are:

- PN1: 4 layers with 20;32;64;32 neurons, output size 16;
- PN2: 2 layers with 17;32 neurons, output size 16;
- SCN: 1 layers with 17 neurons, output size 1;
- SN: 1 layers with 16 neurons, output size 1.

For (O1), d is trained and tested on the dataset for 100 epochs, with a 75/25 (2000/663) train/test split, all while keeping track of the loss and percentages of true/false positive places.

For (O2), inference is performed using the trained d on the test dataset from (O1), with a beam width $b = 10$. For the produced process models, we measure alignment-based fitness and precision [17] to obtain F-scores and simplicity scores based on the inverse arc degree [31]. We compared to Inductive Miner (IM) [4], Split Miner (SM) [6], and Heuristics Miner (HM) [32].

For (O3), inference is performed on eleven real-life datasets that are selected from the BPI challenge, with a beam width $b = 50$. Because of the input size of 20 and the artificial start and end transition, only the 18 most frequent occurring activities are taken from each dataset. For both O2 and O3, we lowered beam width after each selected place. Furthermore, we only used a sample of size between 8 and 50 (depending on the size of G) from each log to keep the number of candidate places tractable. Note that conformance checking is always done on the complete event log.

All experiments are performed on an Intel i7-8705G CPU (3.10GHz) with 16 GB RAM and no GPU support; the source is available at gitlab.com/dominiquesommers/apd-ml. See [29] for a more extensive evaluation.

B. Results

(O1) the loss converges towards zero after the 100 epochs, but does not reach zero, which is expected since the ordering of candidates is not completely canonical (c.f Sect. IV-B). Inference during training shows a similar trend for selecting true positive candidate places. The number of false positives keeps decreasing after the loss is converged, due to the use of teacher forcing causing a mismatch in training and inference. Note that the loss is directly related to the number of true positives, but not to the false positives as teacher forcing corrects d after every selection during training.

(O2) the scatter plot in Fig. 5 shows the F-scores and simplicities of all produced models for the test dataset. Only for a subset of around 65% did our approach find a valid process model, i.e. being at least easy sound and only their values are shown in Fig. 5. The medians and means for each method and the ground truth are shown as well indicating

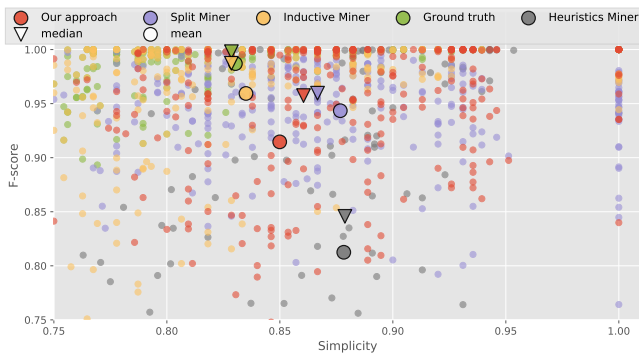


Fig. 5. Conformance results on the synthetic test dataset with mean and median centroids.




Fig. 6. Histogram of F-score ratios comparing our approach to the ground truth and other methods.

the average model qualities across the methods. Looking at the median centroids, IM clearly comes closest to the ground truth (of block-structured models). Our approach and the SM discover simpler models with a small trade-off in F-score. HM lacks in F-score while achieving high simplicity. The large gap between the median and the mean for our approach shows that for a lot of samples, the conformance is very high and that for few samples the conformance is very low. Looking further into these, most of the lower scoring samples are not sound (only easy sound) and the processes are complex in structure: parallelism in the upper part of the process tree with multiple parallelism and loops in the lower parts, causing a large number of candidate places.

To evaluate relative performance per log, we compute the F-score ratios of our approach over each other approach: a ratio > 1 means our approach performed better on a log. The histogram in Fig. 6 shows that our approach outperforms SM on many logs while our approach struggles more significantly on a smaller subset of logs.

(O3) Fig. 7 plots F-score and simplicity for each method and real-life dataset separately. The IM is known to problem to have low precision causing low F-scores. The SM generally achieves similar F-scores as the HM, but has higher simplicity. Our approach competes with the best scoring methods on all datasets in terms of F-scores and has the highest simplicity (except for BPIC'12); the average score plot in Fig. 7 reinforces this observation.

Fig. 8 shows the model discovered by our approach in the Road Traffic Fine dataset. The Petri net is not block structured but sound proving generalization beyond the block-structured

training data. Furthermore, the model captures a complex synchronization of 2 parallel branches (repeated payments, sending fine) by an optional appeals procedure (with penalty in parallel) in an unstructured loop: payments can resume after conclusion of appeals.

Regarding complexity and running times, creating graph G from L is linear in the number of unique trace variants and polynomial in the number of activities by pruning candidate places [28] (we observed 1-10 secs). Information propagation and the computation of the embeddings is independent of the size of G . Place selection is linear in the number of selected candidates (on avg. 1 sec per place) multiplied by the beam width. When d has a hard time selecting places that keep the S-coverability, substantially more candidates have to be considered. We observed running times of several minutes while the algorithmic methods complete in the matter of seconds [10].

VI. CONCLUSION AND FUTURE WORK

Our approach proved to perform well on the task it was trained on and generalizes beyond that exceeding the state-of-the-art methods for imperative models on real-life datasets. This research laid down the foundation of introducing machine learning to APD and provides a framework in which various components can be optimized and replaced. Its design being based on a human modeler, which is still superior to automated methods, could enable new lines of APD research; especially to consider declarative models as target language [33].

A limiting factor for this method to be scalable to large processes is the choice of the initial features, being one-hot encodings, allowing only a predefined maximum number of distinct activities. Future research is needed on a different label encoder that is scalable to the cardinality to counter this limitation.

Robustness has been a focus in automated process discovery since a wide variety of processes exist in terms of structure and algorithmic approaches often tackle only a specific kind of process. This has been a problem for our method as well where no sound model was discovered at all in some cases, which could be caused by the lack of generalization from the dataset it has been trained on, having a very specific representational bias. In our experiments, we reduced the number of traces to reduce the size of G until a sound model was discovered. Although this is justified by the fact that human modelers also only look at the top most frequent/important traces and discover a model based on that, it is not ideal since possibly valuable information is lost. A first improvement here would be to increase the variety in the training data, but a proper fallback method should be in place to ensure robustness.

REFERENCES

- [1] S. J. J. Leemans, Automated process discovery, in: Encyclopedia of Big Data Technologies, Springer, 2019.
- [2] J. C. A. M. Buijs, B. F. van Dongen, W. M. P. van der Aalst, Quality dimensions in process discovery: The importance of fitness, precision, generalization and simplicity, Int. J. Cooperative Inf. Syst. 23 (1) (2014).
- [3] H. M. W. Verbeek, T. Basten, W. M. P. van der Aalst, Diagnosing workflow processes using woflan, Comput. J. 44 (4) (2001) 246–279.




Fig. 7. Conformance results on real-life datasets.




Fig. 8. Discovered Petri net for the Road Traffic Fine dataset.

- [4] S. J. J. Leemans, D. Fahland, W. M. P. van der Aalst, Discovering block-structured process models from event logs containing infrequent behaviour, in: *BPM 2013 Workshops*, Vol. 171 of LNBIP, 2013, pp. 66–78.
- [5] S. K. L. M. vanden Broucke, J. D. Weerd, Fodina: A robust and flexible heuristic process discovery technique, *Decis. Support Syst.* 100 (2017) 109–118.
- [6] A. Augusto, R. Conforti, M. Dumas, M. La Rosa, G. Bruno, Automated discovery of structured process models from event logs: The discover-and-structure approach, *Data and Knowledge Engineering* 117 (April) (2018) 373–392.
- [7] J. M. E. M. van der Werf, B. F. van Dongen, C. A. J. Hurkens, A. Serebrenik, Process discovery using integer linear programming, *Fundam. Informaticae* 94 (3-4) (2009) 387–412.
- [8] M. Solé, J. Carmona, Encoding process discovery problems in SMT, *Softw. Syst. Model.* 17 (4) (2018) 1055–1078.
- [9] J. D. Weerd, M. D. Backer, J. Vanthienen, B. Baesens, A multi-dimensional quality assessment of state-of-the-art process discovery algorithms using real-life event logs, *Inf. Syst.* 37 (7) (2012) 654–676. doi:10.1016/j.is.2012.02.004. URL <https://doi.org/10.1016/j.is.2012.02.004>
- [10] A. Augusto, R. Conforti, M. Dumas, M. La Rosa, F. M. Maggi, A. Marrella, M. Mecella, A. Soo, Automated Discovery of Process Models from Event Logs: Review and Benchmark, *IEEE Transactions on Knowledge and Data Engineering* 31 (4) (2019) 686–705.
- [11] J. M. E. M. van der Werf, A. Polyvyanyy, B. R. van Wensveen, M. Brinkhuis, H. A. Reijers, All that glitters is not gold - towards process discovery techniques with guarantees, in: *CAiSE 2021*, Vol. 12751 of LNCS, Springer, 2021, pp. 141–157.
- [12] J. Pinggera, S. Zugal, B. Weber, D. Fahland, M. Weidlich, J. Mendling, H. A. Reijers, How the structuring of domain knowledge helps casual process modelers, in: *Conceptual Modeling - ER 2010*, Vol. 6412 of LNCS, Springer, 2010, pp. 445–451.
- [13] J. Pinggera, S. Zugal, M. Weidlich, D. Fahland, B. Weber, J. Mendling, H. A. Reijers, Tracing the process of process modeling with modeling phase diagrams, in: *BPM Workshops*, Vol. 99 of LNBIP, Springer, 2011, pp. 370–382.
- [14] J. Pinggera, P. Soffer, D. Fahland, M. Weidlich, S. Zugal, B. Weber, H. Reijers, J. Mendling, Styles in business process modeling: an exploration and a model, *Software & Systems Modeling* 14 (2013) 1055–1080.
- [15] S. Goedertier, D. Martens, J. Vanthienen, B. Baesens, Robust process discovery with artificial negative events, *J. Mach. Learn. Res.* 10 (2009) 1305–1340.
- [16] A. Graves, Sequence transduction with recurrent neural networks, arXiv preprint arXiv:1211.3711 (2012).
- [17] J. Carmona, B. F. van Dongen, A. Solti, M. Weidlich, *Conformance Checking - Relating Processes and Models*, Springer, 2018.
- [18] A. Polyvyanyy, A. Solti, M. Weidlich, C. D. Ciccio, J. Mendling, Monotone precision and recall measures for comparing executions and specifications of dynamic systems, *ACM Trans. Softw. Eng. Methodol.* 29 (3) (2020) 17:1–17:41.
- [19] W. M. P. van der Aalst, et al., Process mining manifesto, in: *BPM 2011 Workshops*, Vol. 99 of LNBIP, 2011, pp. 169–194.
- [20] S. J. J. Leemans, N. Tax, A. H. M. ter Hofstede, Indulpet miner: Combining discovery algorithms, in: *OTM 2018*, Vol. 11229 of LNCS, Springer, 2018, pp. 97–115.
- [21] W. M. P. van der Aalst, T. Weijters, L. Maruster, Workflow mining: Discovering process models from event logs, *IEEE Trans. Knowl. Data Eng.* 16 (9) (2004) 1128–1142.
- [22] S. J. J. Leemans, D. Fahland, Information-preserving abstractions of event data in process mining, *Knowl. Inf. Syst.* 62 (3) (2020) 1143–1197.
- [23] R. Bergenthum, Prime miner - process discovery using prime event structures, in: *International Conference on Process Mining, ICPM 2019, Aachen, Germany, June 24-26, 2019*, IEEE, 2019, pp. 41–48.
- [24] T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, arXiv preprint arXiv:1609.02907 (2016).
- [25] Y. Li, O. Vinyals, C. Dyer, R. Pascanu, P. Battaglia, Learning deep generative models of graphs, arXiv preprint arXiv:1803.03324 (2018).
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, arXiv preprint arXiv:1706.03762 (2017).
- [27] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, Graph attention networks, arXiv preprint arXiv:1710.10903 (2017).
- [28] D. Sommers, V. Menkovski, D. Fahland, Process discovery using graph neural networks (2021). arXiv:2109.05835.
- [29] D. Sommers, D. Fahland, V. Menkovski, Process Discovery using Machine Learning, Master's thesis, Eindhoven University of Technology, the Netherlands (2020).
- [30] R. J. Williams, D. Zipser, A learning algorithm for continually running fully recurrent neural networks, *Neural computation* 1 (2) (1989) 270–280.
- [31] F. R. Blum, Metrics in process discovery, Tech. rep., Tech. Rep. TR/DCC. 1–21 (2015).
- [32] A. Weijters, W. M. van Der Aalst, A. A. De Medeiros, Process mining with the heuristics miner-algorithm, Technische Universiteit Eindhoven, Tech. Rep. WP 166 (2006) 1–34.
- [33] S. Schöning, C. D. Ciccio, F. Maggi, J. Mendling, Discovery of multi-perspective declarative process models, in: *ICSOC*, 2016.