






# Learning Accurate Business Process Simulation Models from Event Logs via Automated Process Discovery and Deep Learning

Manuel Camargo<sup>1,2,3</sup> , Marlon Dumas<sup>1</sup> , and Oscar González-Rojas<sup>2</sup> 

<sup>1</sup> University of Tartu, Tartu, Estonia  
marlon.dumas@ut.ee

<sup>2</sup> Universidad de los Andes, Bogotá, Colombia  
o-gonza1@uniandes.edu.co

<sup>3</sup> Apromore, Tartu, Estonia  
manuel.camargo@apromore.com

**Abstract.** Business process simulation is a well-known approach to estimate the impact of changes to a process with respect to time and cost measures – a practice known as what-if process analysis. The usefulness of such estimations hinges on the accuracy of the underlying simulation model. Data-Driven Simulation (DDS) methods leverage process mining techniques to learn process simulation models from event logs. Empirical studies have shown that, while DDS models adequately capture the observed sequences of activities and their frequencies, they fail to accurately capture the temporal dynamics of real-life processes. In contrast, generative Deep Learning (DL) models are better able to capture such temporal dynamics. The drawback of DL models is that users cannot alter them for what-if analysis due to their black-box nature. This paper presents a hybrid approach to learn process simulation models from event logs wherein a (stochastic) process model is extracted via DDS techniques, and then combined with a DL model to generate timestamped event sequences. An experimental evaluation shows that the resulting hybrid simulation models match the temporal accuracy of pure DL models, while partially retaining the what-if analysis capability of DDS approaches.

**Keywords:** Process mining · Simulation · Deep learning

## 1 Introduction

Business Process Simulation (BPS) models allow analysts to estimate the impact of changes to a process with respect to temporal and cost measures – a practice known as “what-if” process analysis [1]. However, the construction and tuning

---

Work funded by European Research Council (PIX Project).

© The Author(s) 2022

X. Franch et al. (Eds.): CAiSE 2022, LNCS 13295, pp. 55–71, 2022.

[https://doi.org/10.1007/978-3-031-07472-1\\_4](https://doi.org/10.1007/978-3-031-07472-1_4)

of BPS models is error-prone, as it requires careful attention to numerous pitfalls [2]. Moreover, the accuracy of manually tuned BPS models is limited by the completeness of the process model used as a starting point, yet manually designed models often do not capture exceptional paths. Previous studies have proposed to extract BPS models from execution data (event logs) via process mining techniques [3]. While Data-Driven Simulation (DDS) models extracted in this way can be tuned to accurately capture the control-flow and temporal behavior of a process [4], they suffer from fundamental limitations stemming from the expressiveness of the modeling notation (e.g. BPMN, Petri nets) and from assumptions about the resources' behavior. One such assumption is that all waiting times are due to resource contention (i.e. a resource not starting a task because it is busy with another task). Another assumption is that resources exhibit robotic behavior: if a resource is available, and it may perform an enabled activity instance, the resource will immediately start it. In other words, these approaches do not take into account behaviors such as multitasking, batching, fatigue effects, and inter-process resource sharing, among others [5].

Other studies have shown that Deep Learning (DL) generative models trained from logs can accurately predict the next event in a case and its timestamp or the suffix of a case starting from a given prefix [6,7]. Suitably trained DL generative models can also be used to generate entire traces and even entire logs [8], which effectively allows us to use a DL generative model as a simulation model. Camargo et al. [9] empirically show that DL models are more accurate than DDS models when it comes to generating logs consisting of activity sequences with start and end timestamps. In particular, generative DL models can emulate delays between activities that DDS models do not capture. However, unlike DDS models, DL models are not suitable for what-if analysis due to their black-box nature – they do not allow to specify a change to the process and to simulate the effect of this change.

This paper presents a method, namely DeepSimulator, that combines DDS and DL methods to discover a BPS model from a log. The idea is to use an automated process discovery technique to extract a process model with branching probabilities (a.k.a. stochastic process model [10]) and to delegate the generation of activity start and end times to a DL model.

The paper is structured as follows. Section 2 discusses methods to learn generative models from logs using DDS and DL techniques. Section 3 presents the proposed method while Sect. 4 presents an empirical evaluation thereof. Finally, Sect. 5 draws conclusions and sketches future work.

## 2 Related Work

### 2.1 Data-Driven Simulation of Business Processes

Previous studies on DDS methods can be classified in two categories. A subset of previous studies have proposed conceptual frameworks and guidelines to manually derive, validate, and tune BPS parameters from event logs [3,11], without

seeking to automate the extraction process. Other studies have proposed methods that automate the extraction and/or tuning of simulation parameters from logs. In this paper, we focus on automated methods. One of the earliest such methods is that of Rozinat et al. [12], who propose a semi-automated approach to extract BPS models based on Colored Petri Nets. Later, Khodyrev et al. [13] proposed an approach to extract BPS models from data, although they leave aside the resource perspective (i.e. the discovery of resource pools). More recently, Pourbafrani et al. [14] present an approach for generating DDS models based on time-aware process trees. In all of the above studies, the responsibility for tuning these parameters is left to the user. This limitation is addressed in the Simod method [4], which automates the extraction of BPS models by employing Bayesian optimization to tune the hyperparameters used to discover the process model and resource pools as well as the statistical parameters of the BPS model (branching probabilities, activity processing times, and inter-case arrival times). This tuning phase seeks to optimize the similarity between the logs produced by the extracted BPS model and (a testing fold of) the original log.

## 2.2 Generative DL Models of Business Processes

A Deep Learning (DL) model is a network of interconnected layers of neurons (perceptrons) that collectively perform non-linear data transformations [15]. The objective of these transformations is to train the network to learn the patterns observed in the data. In theory, the more layers of neurons in the system, the more it will detect higher-level patterns via composition of complex functions [15]. A wide range of neural network architectures have been proposed, e.g. feed-forward networks, Convolutional Neural Networks, Variational Auto-Encoders, Generative Adversarial Networks (GAN) (often in a combination with other architectures), and Recurrent Neural Networks (RNN). The latter type of architecture is specifically designed to handle sequential data.

DL models have been widely applied in the field of predictive process monitoring. Evermann et al. [7] proposed an RNN architecture to generate the most likely remaining sequence of events (suffix) of an ongoing case. This architecture cannot handle numeric features and thus cannot generate timestamped events (timestamps are numeric). This limitation is shared by the approach in [16] and others reviewed in [17]. Tax et al. [6] use an RNN architecture known as Long-Short-Term Memory (LSTM) to predict the next event in an ongoing case and its timestamp, and to generate the remaining sequence of timestamped events from a given prefix of a case. However, this approach cannot handle high-dimensional inputs due to its reliance on one-hot encoding of categorical features. Its precision deteriorates as the number of categorical features increases. This limitation is lifted by DeepGenerator [8], which extends the approach in [6] with two mechanisms to handle high-dimensional input: n-grams and embeddings. This approach also addresses the problem of generating long suffixes (not well handled in [6]) and entire traces, by using a random next-event selection approach. It is also able to associate a resource to each event in a trace. More recently, Taymouri

et al. [18] proposed a GAN-LSTM architecture to train generative models that produce timestamped activity sequences (without associated resources).

Camargo et al. [9] compare the relative accuracy of DL models against DDS models for generating sequences of the form (activity, start timestamp, end timestamp). This comparison suggests that DL models may outperform DDS models when trained with large logs, while the opposite holds for smaller logs. Camargo et al. [9] additionally show that DL models generally outperform DDS methods when it comes to predicting activity start and end timestamps.

### 3 Hybrid Learning of BPS Models

Figure 1 depicts the architecture of the DeepSimulator approach. The architecture is a pipeline with three phases. The first phase uses PM techniques to learn a model to generate sequences of (non-timestamped) events. The second and third phases enrich these sequences with case start times and activity start and end times. Below, we discuss each phase in turn.

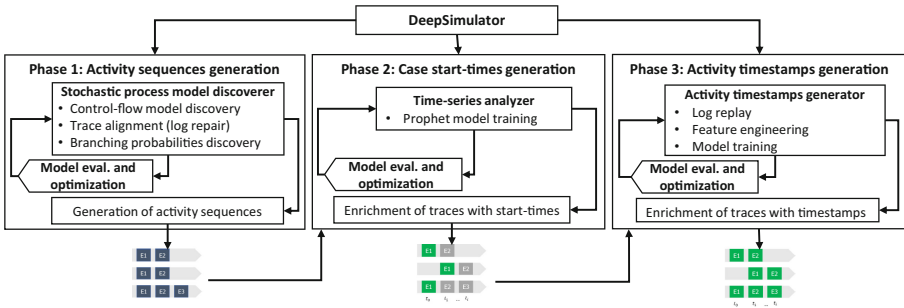


Fig. 1. Overview of the proposed BPS model discovery method

*Phase 1: Activity Sequences Generation.* The aim of this phase is to extract a stochastic process model [10] from the log and to use it to generate sequences of activities that resemble those in the log. A stochastic process model is a process model with branching probabilities assigned to each branch of a decision point. In this paper, we represent process models using the standard BPMN notation. Phase 1 starts with a *control-flow discovery* step, where we first discover a plain (non-stochastic) process model using the Split Miner algorithm [19]. This algorithm relies on two parameters: the sensitivity of the parallelism oracle ( $\eta$ ) and the level of filtering of directly-follow relations ( $\epsilon$ ). The former parameter determines how likely the algorithm will discover parallel structures, while the latter determines the percentage of directly-follows relations between activity types are captured in the resulting model. Like other automated process discovery algorithms, the Split Miner discovers a process model that does not perfectly fit the log. The discovered process model cannot parse some traces in

the log. This hinders the calculation of the branching probabilities. Accordingly, we apply the *trace alignment* algorithm in [20] to compute an alignment for each trace in the log that the model cannot parse. An alignment describes how a trace can be modified to be turned into a trace that can be parsed by the model (via “skip” operations). Based on the alignments, we either *repair* each non-conformant trace, or we *replace* it with a copy of the most similar conformant trace (w.r.t. string-edit distance). The choice between the repair and the replacement approaches is a parameter of the method.

Next, DeepSimulator uses the (conformant) event log to discover the *branching probabilities* for each branching point in the model. Here, DeepSimulator offers two options: (i) assign equal values to each conditional branch; or (ii) compute the branching probabilities by replaying the aligned event against the process model. The first approach may perform better for smaller logs, where the probabilities computed via replay are not always reliable, while the latter may be preferable for larger logs.

The DeepSimulator combines the process model and the branching probabilities to assemble a stochastic process model. In this step, the DeepSimulator uses a Bayesian optimization technique to discover the hyperparameter settings (i.e., values of  $\epsilon$ ,  $\eta$ , replace-vs-repair, and equal-vs-computed probabilities) that maximize the similarity between the generated and the ground truth sequences in terms of activity sequences. The optimizer uses a holdout method, and as a loss function, it uses the Control-Flow Log Similarity (CFLS) metric described in [4]. The CFLS metric is the mean string-edit distance between the activity sequences generated by the stochastic process model and the traces in the ground-truth log after their optimal alignment.<sup>1</sup> Finally, in the *sequences’ generation* step, DeepSimulator uses the resulting stochastic process model to generate a bag of activity sequences without timestamps. This bag is used as the log’s base structure in Phase 3.

*Phase 2: Case Start-Times Generation.* In this phase, we generate each process instance’s start time in the output log. Traditionally, DDS models generate the start-time of cases by randomly drawing from an unimodal distribution of the interarrival times between consecutive cases. A typical BPS model captures the interarrival times using a negative exponential distribution (i.e., it models the creation of cases as a Poisson process). However, a single distribution is not realistic enough to capture real scenarios. For example, cases might be created more frequently on Mondays than on Thursdays in a claims handling process.

Instead of fitting an interarrival distribution, the DeepSimulator models the case generation as a time series prediction problem as the number of cases generated per hour of the day. This type of modeling allows us to use robust techniques such as ARIMA or ETS tested successfully in several contexts such as Stock Market Analysis or Workload Projections. DeepSimulator uses the Prophet [22]

---

<sup>1</sup> We did not use the stochastic conformance checking metrics over Petri nets of [10] since our method handles BPMN models with inclusive join gateways, which cannot be directly transformed to Petri Nets (without exponential blowout) as shown in [21].

model proposed by Facebook because it is one of the simplest but, at the same time, more accurate predictive models for this type of task. Prophet starts from the time series decomposition into four main components (i.e., trend, seasonality, holidays, and error) and applies specialized techniques to model each component.

The trend component decomposes those non-periodic changes in the time series values, which are modeled using logistic growth models or Piecewise linear models. The seasonality component decomposes the periodic changes repeated at fixed intervals (hours, weeks, months, or years), which are modeled by using the Fourier series. The holidays component represents the effects of holidays that occur on potentially irregular schedules over one or more days. This component is optionally modeled and is defined manually by a domain expert, since it is specific to each time series. The model automatically calculates the error, corresponding to all those unforeseen changes that the model cannot fit.


In the *time-series analysis* step, we use a saturated logistic growth model to fit the case generation trend. We chose this model, considering that the time series is limited by a lower and upper bound. The lower bound corresponds to 0, which is the minimum number of cases attended in the process, and the upper bound is theoretically limited by the capacity of the process. The parameter that most significantly affects data trend capture is *change-point-prior-scale*, which determines how much the trend changes at the trend change points. This parameter needs tuning, since a too low value may cause under-fitting while a too high value may cause over-fitting. Accordingly, for this parameter, DeepSimulator explores values in the interval  $[0.001, 0.5]$ . Analogously, the parameter that most directly affects the seasonality capture is the *seasonality-prior-scale*. This parameter affects the flexibility of seasonality learning. If the value is too small, the model tends to focus on small fluctuations, while a large value may cause the model to focus only on large fluctuations. For this parameter, DeepSimulator explores values in  $[0.01, 10]$ . We do not define the Holidays component in the Prophet model. The holidays component could be discovered by a calendar discovery technique such as the one proposed in [5], but discovering such calendars is orthogonal to the focus of the present paper.

We use grid search for selecting the best hyperparameters of the Prophet model, as the search space consists of only sixteen configurations (cf. Sect. 4.3). We rely on the internal mechanisms embedded in Prophet for cross-validation and selection of cutoff points. During the simulation, we use the trained Prophet model to determine the number of cases to be created at each hour of the simulation. We then generate the start-times, for each simulation hour, by modeling the intercase arrival times via a normal distribution (within the hour).

*Phase 3: Activity Timestamps Generation.* We enhance the activity sequences generated in Phase 1 to capture waiting times and processing times in this phase.

The DeepSimulator trains two LSTM<sup>2</sup> models to perform two predictive tasks: the processing time of a given activity (herein called the *current activity*) and the waiting time until the start of the next activity. This task differs in two ways from approaches used to predict the next event and its timestamp [8, 18]. First, we do not seek to predict the next event, since the sequences of activities are generated by the stochastic process model (cf. Phase 1). Second, we need to support changes in the process model (e.g., adding or removing tasks) for enabling what-if analysis.

Therefore, we train one model specialized in predicting the processing time of the current activity and another specialized in predicting the waiting time until the next activity. Both models differ in the set of features since they act at different moments in the predictive phase, as shown in Fig. 2.



**Fig. 2.** Predictive models timeline and features

The processing time predictive model uses the following features as inputs: the label and processing time of the current activity, the time of day of the current activity’s start timestamp, the day of the week, and inter-case features such as the Work-in-progress (WIP) of the process and the activity and Resources’ Occupation (RO) at the start of the activity. The waiting time predictive model uses the following features as inputs: the next activity’s label, the time of day of the current activity’s end timestamp, the day of the week, and inter-case features such as the WIP of the process and the RO at the end of the current activity.

In the *log replay* step, we calculate the waiting and processing times of each activity by replaying each trace in the input log (or in a training subset thereof) against the process model discovered in Phase 1. An activity’s processing time is the difference between its end and start timestamps. An activity’s waiting time is the difference between its start time and enablement time, i.e., when it was ready to be executed according to the process model. All waiting and processing times are scaled to the range [0...1] by dividing them by the largest values.

In the *feature engineering* step, we compute and encode all the remaining features used by the models. We calculate the time of the day as the elapsed seconds from the closest midnight until the event timestamp; this feature is scale over 86400 s. The day of the week is modeled as a categorical attribute and encoded using one-hot encoding. We include these latter features since they provide contextual information, allowing the model to find seasonal patterns in the data that may affect waiting and processing times. In the same way,

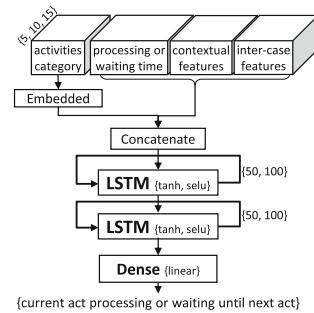
<sup>2</sup> We used LSTM networks as the core of our predictive models since they are a well-known and proven technology to handle sequences, as the nature of a business process event log [23]. The proposed models were based on previous works [6–8] that have extensively explored several architectural options of LSTMs, such as the use of stacked vs. unstacked models or the use of shared layers vs. specialized ones.

considering that the overall process performance is affected by the process' WIP and the RO [24], we use two inter-case features that measure these variations.

The WIP of the process measures the number of active tasks at each moment in the log transversally. The RO measures each resource pool's percentage occupancy in the log, implying that a new feature is created for each pool to record the occupation-specific variations. Since the information about the size and composition of the resource pools is not always included in the logs, we grouped resources into roles by using the algorithm described in [25]. This algorithm discovers resource pools based on the definition of activity execution profiles for each resource and the creation of a correlation matrix of similarity of those profiles. WIP and RO are calculated by replaying over time the log events, recording the variations in both features at every time point.

Finally, we encode the current activity's label using pre-trained embedded dimensions. We use embeddings for two reasons. First, embeddings help prevent exponential feature growth associated with one-hot encoding [8]. Second, embedded dimensions allow adding new categories (i.e., activity labels) without altering the predictive model's structure. These embedded dimensions are an n-dimensional space, where each category (each activity level) is encoded as a point in that space. An independent network fed with positive and negative examples of associations between activities is used to map the activity labels to points. The network maps activities that co-occur or occur close to each other to nearby points. This mechanism also allows adding a new point in that space by updating the encoding model without altering the predictive model's input size. Each time a new activity is added to the process model for what-if analysis, we generate examples of traces involving this new activity and use these examples to determine the coordinates of the new activity label to be encoded in the embedded space. Then, we update the predictive model's embedded layers with the new definition, and the predictive model can handle the new activity label from that point on.

Once encoded the features, we extract n-grams of fixed sizes from each trace to create the input sequences to train the model. As shown in Fig. 3, both models are composed of two stacked LSTM layers and a dense output. A model receives the sequences as inputs and the expected processing and waiting times as a target. The user can vary the number of units in the LSTM layers, the activation function, the size of the n-gram, and the use of all the RO inter-cases or just the one of the resource pool associated to the execution of the activity.



**Fig. 3.** DL models architectures

*Assembling the Output Log.* The output log is generated by assembling each generated sequence (see Phase 1), with the generated case start time (see Phase 2) and the processing and waiting times predicted iteratively (see Phase 3). In each iteration, the trained model predicts times relative to the current activity in



seconds, which are transformed into absolute times by adding them to the start time of the case. Then, the DeepSimulator generates a simulated log composed of a bag of traces, each trace consisting of a sequence of triplets (activity label, start-timestamp, end timestamp).

## 4 Evaluation

We empirically compare the DeepSimulator method vs. DDS and DL approaches in terms of the similarity of the simulated logs they generate relative to a fold of the original log. We also evaluate the accuracy of DeepSimulator for “what-if” analysis tasks of modifying the case creation intensity and adding new activities to a process.

### 4.1 Datasets

We evaluated the approaches using 9 logs that contain both start and end timestamps. We use real-life logs (R) from public and private sources and synthetic logs (S) generated from simulation models of real processes. Table 1 provides descriptive statistics of the logs. The BPI17W log have the largest number of traces and events, while CFS and P2P have fewer traces but more events/trace.

**Table 1.** Event logs description. (\*) Private logs, (\*\*) Generated from simulation models of real processes

Size	Source	Log	#Traces	#Events	#Act.	Avg. activities per trace	Avg. duration	Max. duration	Description
LARGE	R	BPI17W	30276	240854	8	7.96	12.66 days	286.07 days	Dutch financial institution updated
	R	BPI12W	8616	59302	6	6.88	8.91 days	85.87 days	Dutch financial institution
	S	CVS	10000	103906	15	10.39	7.58 days	21.0 days	CVS retail pharmacy**
	S	CFM	1670	44373	29	26.57	0.76 days	5.83 days	Anonymized confidential process**
SMALL	R	INS	1182	23141	9	19.58	70.93 days	599.9 days	Insurance claims process*
	R	ACR	954	4962	16	5.2	14.89 days	135.84 days	Academic Credential Recognition
	R	MP	225	4503	24	20.01	20.63 days	87.5 days	Manufacturing Production
	S	CFS	800	21221	29	26.53	0.83 days	4.09 days	Anonymized confidential process**
	S	P2P	608	9119	21	15	21.46 days	108.31 days	Purchase-to-Pay process

### 4.2 Evaluation Measures

To evaluate the accuracy of a model  $M$  produced by one of the methods under evaluation, we compute a distance measure between a log generated by model  $M$  and a ground-truth log (a testing subset of the original log). In all of our experiments, we use two distance measures: the Mean Absolute Error (MAE) of cycle times and the Earth-Mover’s Distance (EMD) of the normalized histograms of activity timestamps grouped by day/hour.

The *cycle time MAE* measures the temporal similarity between two logs at the *trace level*. The absolute error of a pair of traces T1 and T2 is the absolute value of the difference between their cycle times. The cycle time MAE is the mean of the absolute errors over a collection of paired traces. Given this trace distance notion, we pair each trace in the generated log with a trace in the original log using the Hungarian algorithm [26] so that the sum of the trace errors between the paired traces is minimal.

The cycle time MAE is a rough measure of the temporal similarity between the ground-truth and the simulated traces. But it does not consider the start time of each case, nor the start and end timestamps of each activity. To complement MAE, we use the *Earth Mover’s Distance (EMD)* between the normalized histograms of the timestamps grouped by day/hour in the ground-truth and the generated logs. The EMD between two histograms, H1 and H2, is the minimum number of units that need to be added, removed, or transferred across columns in H1 to transform it into H2. The EMD is zero if the observed distributions in the two logs are identical, and it tends to one the more they differ.

### 4.3 Experiment 1: AS-IS Accuracy of Generated Models

**Setup.** This experiment aims to compare the accuracy of DeepSimulator models (herein called DSIM models) vs. DDS and DL models. We use SIMOD [4] as a baseline DDS approach since it is fully automated both w.r.t. parameter discovery and tuning. As DL baselines, we use an adaptation of the LSTM approach proposed by Camargo et al. [8] (herein labeled the LSTM method) as well as the GAN-LSTM approach by Taymouri et al. [18] (herein labeled GAN). Both of these DL approaches have been shown to achieve high accuracy w.r.t. the task of generating timestamped trace suffixes [23]. Figure 4 summarizes the experimental setup. We use the hold-out method with a temporal split criterion to divide the logs into two main folds: 80% for training-validation and 20% for testing. From the first fold, we took the first 80% for training and 20% for validation. We use temporal splits to prevent information leakage [8, 18].




Fig. 4. Setup of experiment 1

The DDS technique (SIMOD) is set to explore 15 parameter configurations to tune the stochastic process model. For each configuration, we execute five simulation runs and compute the CFLS measure (cf. Sect. 3) between each simulated

log and the validation fold. We select the stochastic model that gives the lowest average CFLS w.r.t. the validation fold. The optimizer is set to explore 20 simulation parameter configurations (i.e. the parameters that Simod uses to model resources and processing times), again using five simulation runs per configuration. We select the configuration with the lowest average EMD (cf. Sect. 4.2) between the simulated log and the validation fold. We used the parameter ranges given in Table 2 for tuning.

The LSTM technique is hyperparameter-optimized using grid search over a space of 48 possible configurations (see Table 2) For LSTM model training, we use 200 epochs, the cycle time MAE as the model’s loss function, Nadam as the optimizer, and early stopping and dropout to avoid model over-training. The GAN technique is configured to dynamically adjust the size of the hidden units in each layer so that their size is twice the input’s size, as proposed by the authors [18]. We use 25 training epochs, a batch of size five, and a prefix size of five. DSIM is tuned by randomly exploring 15 parameter configurations with

**Table 2.** Hyperparameters used by optimization techniques

Model	Stage	Parameter	Distribution	Values
SIMOD	Structure discovery	Parallelism threshold ( $\epsilon$ )	Uniform	[0...1]
		Percentile for frequency threshold ( $\eta$ )	Uniform	[0...1]
		Conditional branching probabilities	Categorical	{Equiprobable, Discovered}
	Time-related parameters discovery	Log repair technique	Categorical	{Repair, Removal, Replace}
		Resource pools similarity threshold	Uniform	[0...1]
		Resource availability calendar support	Uniform	[0...1]
		Resource availability calendar confidence	Uniform	[0...1]
		Instances creation calendar support	Uniform	[0...1]
Instances creation calendars confidence	Uniform	[0...1]		
LSTM	Training	N-gram size	Categorical	{5, 10, 15}
		Input scaling method	Categorical	{Max, Lognormal}
		# units in hidden layer	Categorical	{50, 100}
		Activation function for hidden layers	Categorical	{selu, tanh}
		Model type	Categorical	{shared_cat, concatenated}
DSIM	Structure generation	Parallelism threshold ( $\epsilon$ )	Uniform	[0...1]
		Percentile for frequency threshold ( $\eta$ )	Uniform	[0...1]
		Conditional branching probabilities	Categorical	{Equiprobable, Discovered}
	Cases start-times generation	Changepoint-prior-scale	Categorical	{0.001, 0.01, 0.1, 0.5}
		Seasonality-prior-scale	Categorical	{0.01, 0.1, 1.0, 10.0}
	Timestamps generation	N-gram size	Categorical	{5, 10, 15}
		# units in hidden layer	Categorical	{50, 100}
		Activation function for hidden layers	Categorical	{selu, tanh}
Single resource-pool intercase feature		Boolean	{True, False}	

five simulation runs per configuration in the stochastic model discovery phase (cf. Sect. 3, Phase 1). In Phases 2 and 3, we use grid search to explore the space of hyperparameter configurations specified in Table 2

We generate four models per log: one SIMOD, one LSTM, one GAN, and one DSIM. We then generate five logs per retained model, each with the same number of traces as the original log’s testing fold to ensure the comparability. Each generated log is compared with the testing fold using the MAE and EMD measures. We report the mean of each of these measures across 5 runs.

**Results.** Table 3 show the results grouped by metrics, log size and source type. Note that MAE and EMD are error/distance measures (lower is better). In 3 out of 4 large logs, DSIM outperforms LSTM and SIMOD w.r.t the MAE measure. In the small logs, DSIM attains lower MAE in 3 out of 5 logs and has similar MAE w.r.t. SIMOD in one other log. Similarly, DSIM outperforms SIMOD in 6 of 9 logs w.r.t. EMD, and achieves results similar to SIMOD in two others.<sup>3</sup>

The results suggest that DSIM is often able to outperform the baselines when it comes to replicating the as-is behavior recorded in an event log. This conclusion should be tempered in light of two threats to validity: (i) an external threat to validity stemming from the limited number of events logs in the experiment; and (ii) a threat to construct validity created by the fact that the accuracy measures do not necessarily capture all the nuances in the control-flow and temporal behavior captured in the original and simulated event logs.

#### 4.4 Experiment 2: What-if Analysis

In this experiment, we compare DSIM’s ability to simulate a process after a change (what-if analysis). We consider two scenarios. In the first one, we assess DSIM’s ability to capture variations in the inter-arrival time between cases (a.k.a. *arrival intensity*), specifically alternations between periods of lower arrival

**Table 3.** Evaluation results (lower values are better)

Size	Type	Log	MAE				EMD	
			GAN	LSTM	SIMOD	DSIM	SIMOD	DSIM
LARGE	R	BPI17W	828165	603688	961727	<u>418422</u>	<u>0.016873</u>	0.034584
		BPI12W	653656	<u>327350</u>	662333	548813	0.056789	<u>0.020098</u>
	S	CVS	952004	667715	1067258	<u>158902</u>	0.018955	<u>0.000004</u>
		CFM	956289	15078	252458	<u>8441</u>	0.240567	<u>0.239087</u>
SMALL	R	INS	1302337	1516368	<u>1090179</u>	1190019	0.143675	<u>0.142097</u>
		ACR	296094	341694	230363	<u>165411</u>	0.207050	<u>0.205106</u>
		MP	210714	321147	298641	<u>157453</u>	0.062227	<u>0.050479</u>
	S	CFS	717266	33016	<u>15297</u>	24326	<u>0.222515</u>	0.266749
		P2P	2347070	2495593	1892415	<u>1836863</u>	<u>0.130655</u>	0.132898

<sup>3</sup> We cannot measure EMD for the LSTM and GAN models because EMD requires that the timestamps in the log are absolute timestamps, while the LSTM and GAN approaches produce relative timestamps (w.r.t. an unknown case start-time). It would be possible to extend the LSTM/GAN approaches to generate logs with absolute timestamps by coupling them with a model to generate start-times (e.g. based on Prophet) but this extension is outside the scope of this paper.

intensity and periods of higher intensity. In the second experiment, we evaluate the ability of DSIM (vs baselines) to estimate the impact of adding a never-before-observed activity to a process. This scenario is challenging for the DSIM and the DL models because these models need to infer the temporal behavior of the new activity using their embedding layers.

**Setup Scenario 1.** We create two modified versions of the three largest records (BPI12W, BPI17W, and CVS). First, to capture a periodic reduction in the arrival intensity of cases, we divide each log into six batches of the same number of cases, and then we create two alternating groups of three batches each.

The first group consists of batches 1, 3, and 5. The batches in this group are left unaltered. These groups represent periods of high arrival intensity. The second group comprises batches 2, 4, and 6. This group is used to emulate periods of low arrival intensity. To do so, we reduce the case arrival rate in these batches by 1/3, by

randomly eliminating two out of three cases. The above altered version of the log capture a situation where the arrival intensity varies, but the waiting times within the cases remain the same. In general, when the arrival intensity goes down, the waiting times should go down. To capture this latter scenario, we create a second altered version of the log to capture decreases in waiting times associated with decreases in arrival rates. Accordingly, we take the first altered log, and we reduce the waiting times in group 2 (batches 2, 4, and 6) by 30%. For illustration, Fig. 5 sketches the modifications made to the BPI17W log.

After altering the logs as above, we train and evaluate the DSIM and SIMOD models as in Experiment 1 (see Fig. 4). Since we are particularly interested here in assessing the ability of the evaluated techniques to model the temporal dynamics of case arrivals, we also report the Dynamic Time Warping (DTW) distance between the time series of the number of cases generated per hour of the day, between the testing partition and the logs generated.




Fig. 5. Original vs modified case creations

**Setup Scenario 2.** For each of the synthetic logs (CVS and CFM), we select a random activity A and eliminate all its occurrences from the log. We then train a DSIM simulation model using this modified log (cf. left-hand side of Fig. 6b). Next, we generate synthetic data consisting of positive and negative samples of pairs composed by the




Fig. 6. Pipeline of scenario 2

activity label and associated resource (cf. Sect. 3). Using these synthetic samples, we update the embedded dimensions to include activity A (without modifying the embedding of the remaining activities). We then plug the updated embedding into the previously trained DSIM model (cf. right-hand side of Fig. 6). We calculate the errors of the DSIM model of the “as-is” process (before a change) and the DSIM model of the “what-if” process (after adding an activity). We measure the error using MAE. Additionally, we report the RMSE and SMAPE metrics to confirm that the results do not depend on the chosen metric.

**Results Scenario 1.** Table 4 presents the MAE, EMD and DTW results. DSIM has a lower error in cycle times in all cases. In version 2 the MAE logs are considerably reduced compared to those of the version 1. We can explain this result because the DL models in charge of predicting waiting and processing times consider inter-case attributes that capture the workload of the process, allowing their adjustment to workload variations. Regarding the results of EMD and DTW, both metrics follow the same trends. In most cases, DSIM obtains the best results. This trend is more evident in version 2, in which DSIM gets better results in all cases. These results indicate that Prophet and DL models are more effective than the baselines at capturing the temporal variations in waiting times due to a decrease in the arrival intensity. This observation should again be tempered by the threats to validity acknowledged above.

**Results Scenario 2.** Table 4 presents the MAE, RSME, and SMAPE grouped for each log, both for the simulation model of the as-is process vs. the model derived after the addition of the activity (what-if model). The what-if model has higher MAE than the baseline models in both event logs. The higher error values are evident in the CVS log, where the SMAPE of the updated model is 184% compared to 31.97% for the baseline model. These results suggest that embedded dimensions incorporated in DSIM can predict the presence of activities that

**Table 4.** Results of scenarios 1 and 2

	Log	MAE		EMD		DTW	
		SIMOD	DSIM	SIMOD	DSIM	SIMOD	DSIM
Scenario 1	<i>Version 1</i>						
	BPI17W	971151	<u>417572</u>	<u>0.02222</u>	0.03593	<u>3185</u>	3647
	BPI12W	660211	<u>534341</u>	0.11295	<u>0.04853</u>	515	<u>458</u>
	CVS	1489252	<u>467572</u>	0.03213	<u>0.00001</u>	3380	<u>849</u>
	<i>Version 2</i>						
	BPI17W	895524	<u>290980</u>	0.06438	<u>0.03218</u>	4528	<u>3431</u>
BPI12W	550266	<u>524995</u>	0.25888	<u>0.22003</u>	726	<u>507</u>	
CVS	540112	<u>246159</u>	0.15674	<u>0.05708</u>	2453	<u>1967</u>	
Scenario 2	Log	MAE		RMSE		SMAPE	
		AS-IS	WHAT-IF	AS-IS	WHAT-IF	AS-IS	WHAT-IF
	CFM	<u>7155</u>	17546	<u>22006</u>	33137	<u>0.15629</u>	0.28762
CVS	<u>283061</u>	1040344	<u>357717</u>	1052255	<u>0.31972</u>	1.84601	

were not present in the training set, but it is unable to adequately estimate their temporal behavior. This observation suggests that DSIM could be extended with more sophisticated embedding techniques (e.g. word2vec or transformer models) to better capture the temporal dynamics of previously unobserved activities (by analogy to activities that have been observed in similar contexts).

## 5 Conclusion

This paper presented a method, namely DeepSimulator, to learn BPS models from event logs based on process mining and DL techniques. The design of DeepSimulator draws upon the observation that DDS methods (based on process mining) do not capture delays between activities caused by factors other than resource contention (e.g. fatigue, batching, inter-process dependencies). In contrast, DL techniques can learn temporal patterns without assuming these patterns stem only from resource contention. Accordingly, DeepSimulator discovers a stochastic process model from a log using process mining, and then uses a DL model to add timestamps to the events produced by the stochastic model. The stochastic model can be modified (activities may be added/removed, branching probabilities may be altered), thus enabling some forms of what-if analysis.

The paper reported on an empirical comparison of the proposed technique with respect to: (i) its ability to replicate the observed as-is behavior; and (ii) its ability to estimate the impact of changes (what-if settings). The evaluation in the “as-is” setting shows that the DeepSimulator method outperforms the baselines (one DDS and two DL methods). The evaluation in the what-if analysis setting shows that DeepSimulator can better estimate the impact of changes in the arrival rate of new cases (the demand) in settings where such changes have been previously observed in the data. However, the accuracy of DeepSimulator degraded when evaluated in a previously unobserved scenario, specifically a scenario where a completely new task is added to the process. We foresee that this drawback could be at least partially addressed by adapting more sophisticated embedding techniques, such as word2vec or transformer models.

Another avenue for future work is to extend the approach to generate events with a “resource” attribute. A related avenue is to extend the approach to support a broader range of changes, such as adding or removing resources. Yet another avenue is to validate the proposed method via case studies, to complement the post-mortem evaluation reported in this paper.

**Reproducibility.** The source code is available at <https://github.com/AdaptiveBProcess/DeepSimulator.git>. The datasets, models, and evaluation results can be found at <https://doi.org/10.5281/zenodo.5734443>.

## References

1. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A.: Fundamentals of Business Process Management. Springer, Heidelberg (2018). [https://doi.org/10.1007/978-3-662-56509-4\\_10](https://doi.org/10.1007/978-3-662-56509-4_10)
2. Aalst, W.M.P.: Business process simulation survival guide. In: vom Brocke, J., Rosemann, M. (eds.) Handbook on Business Process Management 1. IHIS, pp. 337–370. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-642-45100-3\\_15](https://doi.org/10.1007/978-3-642-45100-3_15)
3. Martin, N., Depaire, B., Caris, A.: The use of process mining in business process simulation model construction. *Bus. Inf. Syst. Eng.* **58**(1), 73–87 (2015). <https://doi.org/10.1007/s12599-015-0410-4>
4. Camargo, M., Dumas, M., González-Rojas, O.: Automated discovery of business process simulation models from event logs. *Decis. Support Syst.* **134**, 113284 (2020)
5. Estrada-Torres, B., Camargo, M., Dumas, M., García-Bañuelos, L., Mahdy, I., Yerokhin, M.: Discovering business process simulation models in the presence of multitasking and availability constraints. *Data Knowl. Eng.* **134**, 101897 (2021)
6. Tax, N., Verenich, I., La Rosa, M., Dumas, M.: Predictive business process monitoring with LSTM neural networks. In: Dubois, E., Pohl, K. (eds.) CAiSE 2017. LNCS, vol. 10253, pp. 477–492. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-59536-8\\_30](https://doi.org/10.1007/978-3-319-59536-8_30)
7. Evermann, J., Rehse, J.R., Fettke, P.: Predicting process behaviour using deep learning. *Decis. Support Syst.* **100**, 129–140 (2017)
8. Camargo, M., Dumas, M., González-Rojas, O.: Learning accurate LSTM models of business processes. In: Hildebrandt, T., van Dongen, B.F., Röglinger, M., Mendling, J. (eds.) BPM 2019. LNCS, vol. 11675, pp. 286–302. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-26619-6\\_19](https://doi.org/10.1007/978-3-030-26619-6_19)
9. Camargo, M., Dumas, M., González-Rojas, O.: Discovering generative models from event logs: data-driven simulation vs deep learning. *PeerJ. Comput. Sci.* **7**, e577 (2021)
10. Leemans, S.J.J., van der Aalst, W.M.P., Brockhoff, T., Polyvyanyy, A.: Stochastic process mining: earth movers’ stochastic conformance. *Inform. Syst.* **102**, 101724 (2021)
11. Wynn, M.T., Dumas, M., Fidge, C.J., ter Hofstede, A.H.M., van der Aalst, W.M.P.: Business process simulation for operational decision support. In: ter Hofstede, A., Benatallah, B., Paik, H.-Y. (eds.) BPM 2007. LNCS, vol. 4928, pp. 66–77. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-78238-4\\_8](https://doi.org/10.1007/978-3-540-78238-4_8)
12. Rozinat, A., Mans, R.S., van der Aalst, W.M.P.: Discovering simulation models. *Inform. Syst.* **34**(3), 305–327 (2009)
13. Khodyrev, I., Popova, S.: Discrete modeling and simulation of business processes using event logs. *Procedia Comput. Sci.* **29**, 322–331 (2014)
14. Pourbafrani, M., van Zelst, S.J., van der Aalst, W.M.P.: Supporting automatic system dynamics model generation for simulation in the context of process mining. In: Abramowicz, W., Klein, G. (eds.) BIS 2020. LNBIP, vol. 389, pp. 249–263. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-53337-3\\_19](https://doi.org/10.1007/978-3-030-53337-3_19)
15. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
16. Lin, L., Wen, L., Wang, J.: MM-Pred: a deep predictive model for multi-attribute event sequence. In: Proceedings of SIAM 2019. Society for Industrial and Applied Mathematics, pp. 118–126 (2019)



17. Tax, N., Teinemaa, I., van Zelst, S.J.: An interdisciplinary comparison of sequence modeling methods for next-element prediction. *Softw. Syst. Model.* **19**(6), 1345–1365 (2020). <https://doi.org/10.1007/s10270-020-00789-3>
18. Taymouri, F., Rosa, M.L., Erfani, S., Bozorgi, Z.D., Verenich, I.: Predictive business process monitoring via generative adversarial nets: the case of next event prediction. In: Fahland, D., Ghidini, C., Becker, J., Dumas, M. (eds.) *BPM 2020*. LNCS, vol. 12168, pp. 237–256. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58666-9\\_14](https://doi.org/10.1007/978-3-030-58666-9_14)
19. Augusto, A., Conforti, R., Dumas, M., La Rosa, M., Polyvyanyy, A.: Split miner: automated discovery of accurate and simple business process models from event logs. *Knowl. Inf. Syst.* **59**(2), 251–284 (2018). <https://doi.org/10.1007/s10115-018-1214-x>
20. Reißner, D., Armas-Cervantes, A., Conforti, R., Dumas, M., Fahland, D., La Rosa, M.: Scalable alignment of process models and event logs: an approach based on automata and S-components. *Inform Syst* **94**, 101561 (2020)
21. Favre, C., Völzer, H.: The difficulty of replacing an inclusive OR-join. In: Barros, A., Gal, A., Kindler, E. (eds.) *BPM 2012*. LNCS, vol. 7481, pp. 156–171. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-32885-5\\_12](https://doi.org/10.1007/978-3-642-32885-5_12)
22. Taylor, S.J., Letham, B.: Forecasting at scale. *Am. Stat.* **72**(1), 37–45 (2018)
23. Rama-Maneiro, E., Vidal, J.C., Lama, M.: Deep learning for predictive business process monitoring: review and benchmark (2021). <https://arxiv.org/abs/2009.13251>
24. Laguna, M., Marklund, J.: *Business Process Modeling, Simulation and Design*. CRC Press, New York (2018)
25. Song, M., van der Aalst, W.M.P.: Towards comprehensive support for organizational mining. *Decis. Support Syst.* **46**(1), 300–317 (2008)
26. Kuhn, H.W.: The Hungarian Method for the assignment problem. *Nav. Res. Logist. Q.* **2**, 83–97 (1955)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

