



Time Matters: Time-Aware LSTMs for Predictive Business Process Monitoring

An Nguyen¹(✉), Srijeet Chatterjee¹, SvenWeinzierl², Leo Schwinn¹,
Martin Matzner², and Bjoern Eskofier¹

¹ Department of Computer Science, Friedrich-Alexander-University
Erlangen-Nürnberg (FAU), Erlangen, Germany
{an.nguyen,srijeet.chatterjee,leo.schwinn,bjoern.eskofier}@fau.de
² Institute of Information Systems, Friedrich-Alexander-University
Erlangen-Nürnberg (FAU), Nürnberg, Germany
{sven.weinzierl,martin.matzner}@fau.de

Abstract. Predictive business process monitoring (PBPM) aims to predict future process behavior during ongoing process executions based on event log data. Especially, techniques for the next activity and timestamp prediction can help to improve the performance of operational business processes. Recently, many PBPM solutions based on deep learning were proposed by researchers. Due to the sequential nature of event log data, a common choice is to apply recurrent neural networks with long short-term memory (LSTM) cells. We argue, that the elapsed time between events is informative. However, current PBPM techniques mainly use “vanilla” LSTM cells and hand-crafted time-related control flow features. To better model the time dependencies between events, we propose a new PBPM technique based on time-aware LSTM (T-LSTM) cells. T-LSTM cells incorporate the elapsed time between consecutive events inherently to adjust the cell memory. Furthermore, we introduce cost-sensitive learning to account for the common class imbalance in event logs. Our experiments on publicly available benchmark event logs indicate the effectiveness of the introduced techniques.

Keywords: Predictive business process monitoring · Deep learning · Recurrent neural network · LSTM · Time-Awareness

1 Introduction

In the last years, a variety of predictive business process monitoring (PBPM) techniques that base on machine learning (ML) were proposed by researchers [6] to improve the performance of operational business processes [4]. PBPM is a class of techniques aiming at predicting future process characteristics in running process instances [12], like next activities, next timestamps or process-related performance indicators. Such PBPM techniques produce predictions through

S. Chatterjee—Equal contribution with An Nguyen.

© Springer Nature Switzerland AG 2021

S. Leemans and H. Leopold (Eds.): ICPM 2020 Workshops, LNBIP 406, pp. 112–123, 2021.

https://doi.org/10.1007/978-3-030-72693-5_9

predictive models. These models are in turn constructed based on historical event log data.

A current trend in PBPM is to apply deep neural networks (DNNs) to learn more accurate predictive models from event log data than with “traditional” ML algorithms like probabilistic automata [7]. DNNs belong to the ML-sub-field deep learning (DL) and achieve that by identifying the intricate structures in high-dimensional data through multi-representation learning [11].

Existing DL-based PBPM techniques often rely on DNN architectures consisting of out-of-the-box constructs like layers with a “vanilla” long short-term memory (LSTM) cell [9] or state-of-the-art loss functions for parameter learning.

Event logs can be seen as sequences of events in continuous time with irregular intervals (i.e., elapsed time between consecutive events). We argue that these time intervals are informative in the case of event logs in PBPM. Intuitively, these time intervals describe human behavior of executing business processes. Thus, a time-aware PBPM technique considering information on time intervals could potentially achieve a higher predictive quality. Time information is currently only exploited via hand-crafted control-flow features as inputs to “vanilla” LSTM cells [15]. To better account for the time information in event log data, we propose a new PBPM techniques using time-aware LSTM (T-LSTM). T-LSTM extends the “vanilla” LSTM cells by incorporating the elapsed time between consecutive events in order to adjust the memory state and is inspired by work from Baytas et al. [2].

Furthermore, the problem of next activity prediction is commonly modeled as a supervised multi-class classification problem. The distribution of activities in event logs are commonly skewed. Therefore, we additionally introduce cost-sensitive learning to address the inherent class-imbalances.

The main contributions of this work are summarized below:

- We introduce a time-aware LSTM model for the tasks of predicting next activities and timestamps in PBPM
- We tackle the problem of skewed class distributions via cost-sensitive learning

We evaluate the effectiveness of our proposed techniques by conducting experiments for the next activity and timestamp prediction on publicly available benchmark event logs commonly used for PBPM.

The remainder of the paper is structured as follows: Sect. 2 presents related work on DL-based next activity and timestamp prediction. Section 3 introduces preliminaries and the concept of a LSTM. Sections 4 and 5 describes the architecture of T-LSTM and our experimental setup respectively. Then, in Sects. 6 and 7, we present and discuss our results. Section 8 concludes our paper and discusses future research directions.

2 Related Work

Inspired by the field of natural language processing (NLP), Evermann et al. [7] applied recurrent neural network-based and LSTM-based DNN architectures for

the next activity and next sequence of activity prediction in PBPM. They made use of word embeddings to encode activities of event log’s process instances.

Navarin et al. [14] used a “vanilla” LSTM-based DNN architecture for predicting the completion time of running process instances. They one-hot encoded the activity attributes, computed temporal control-flow attributes, and considered additional real-valued or categorical context attributes.

Tax et al. [15] proposed a multitask learning approach using “vanilla” LSTM cells for next activity and timestamp prediction respectively. Like in [14], they one-hot encoded the activity and computed temporal control-flow features. However, the authors did not consider additional data attributes in their approach. This work acts as a baseline for a variety of other techniques such as [18].

Khan et al. [10] introduced memory augmented neural networks (MANNs) in PBPM. MANNs reduce the number of trainable parameters. In general, the network’s architecture consists of an externalized state memory and two “vanilla” LSTM cells manipulating the memory. One LSTM cell works as encoder and the other one as decoder. Concerning the predictive quality, their approach is comparable to the one presented in [15].

Camagro et al. [5] extended the implementation of [15] and fed the resource attribute into the DNN model. Additionally, instead of one-hot encoding, they applied embeddings, as proposed by Evermann et al. [7].

Taymouri et al. [16] introduced generative adversarial networks (GANs) for the next activity and timestamp prediction. The network’s architecture comprises two “vanilla” LSTM cells. One for the generator and the other one for the discriminator.

To date, several studies have investigated DNN-based PBPM techniques. None of the related works proposes a DL-architecture that explicitly models the elapsed time between two successive events. We address this gap by adapting time-aware LSTM cells [2]. Further, Mehdijev et al. [13] tackle the class imbalance problem in the context of the DNN-based prediction of next activities through a second neural network, namely radial basis function neural network, which generates semi-artificial data of the minority class in the pre-processing phase. In contrast, we adapt cost-sensitive learning to investigate the class-imbalance problem for DL-architectures comprising T-LSTM cells.

3 Background

3.1 Preliminaries

Definition 1 (Event, Trace, Event Log). *An event is a tuple (c, a, ts) where c is the case id, a is the activity (label) and ts is the timestamp. A trace is a non-empty sequence $\sigma = \langle e_1, \dots, e_{|\sigma|} \rangle$ of events such that $\forall i, j \in \{1, \dots, |\sigma|\} e_i.c = e_j.c$ and $e_i.ts \leq e_j.ts$, for $1 \leq i < j \leq |\sigma|$. An event log L is a set $\{\sigma_1, \dots, \sigma_{|L|}\}$ of traces. A trace can also be considered as a sequence of vectors which contain derived control flow information or features. Formally, $\sigma = \langle \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(|\sigma|)} \rangle$, where $\mathbf{x}^{(t)} \in \mathbb{R}^{n \times 1}$ is a vector, and the superscript*

indicates the time-order upon which the events happened. n is the number of features derived for each event.

Definition 2 (Prefix and Label). Given a trace $\sigma = \langle e_1, \dots, e_k, \dots, e_{|\sigma|} \rangle$, a prefix of length k , that is a non-empty sequence, is defined as $f_p^{(k)}(\sigma) = \langle e_1, \dots, e_k \rangle$, with $0 < k < |\sigma_c|$. A next activity label for a prefix of length k is defined as $f_{l,a}^{(k)}(\sigma) = e_{k+1}.a$, whereas a next timestamp label for a prefix of length k is defined as $f_{l,ts}^{(k)}(\sigma) = e_{k+1}.ts$. The above definition also holds for an input trace representing a sequence of vectors. For example, the tuple of all possible prefixes, the tuple of all possible next activity labels and the tuple of all possible next timestamp labels for $\sigma = \langle \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)} \rangle$ are $\langle \langle \mathbf{x}^{(1)} \rangle, \langle \mathbf{x}^{(1)}, \mathbf{x}^{(2)} \rangle, \langle e_2.a, e_3.a \rangle, \text{ and } \langle e_2.ts, e_3.ts \rangle$.

3.2 Long Short-Term Memory Cells

Most of the DNN architectures proposed for the next activity and timestamp prediction in PBPM [17] use “vanilla” LSTM cells [9]. LSTMs belong to the class of recurrent neural networks [11] and are designed to handle temporal dependencies in sequential prediction problems [3].

Given a sequence of inputs $\sigma = \langle \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}, \dots, \mathbf{x}^{(k)} \rangle$, a LSTM computes sequences of outputs $\langle \mathbf{h}^{(1)}, \mathbf{h}^{(2)}, \mathbf{h}^{(3)}, \dots, \mathbf{h}^{(k)} \rangle$ via the following recurrent equations:

$$\begin{aligned}
 \mathbf{f}_g^{(t)} &= \text{sigmoid}(\mathbf{U}_f \mathbf{h}^{(t-1)} + \mathbf{W}_f \mathbf{x}^{(t)} + \mathbf{b}_f) && \text{(forget gate),} \\
 \mathbf{i}_g^{(t)} &= \text{sigmoid}(\mathbf{U}_i \mathbf{h}^{(t-1)} + \mathbf{W}_i \mathbf{x}^{(t)} + \mathbf{b}_i) && \text{(input gate),} \\
 \tilde{\mathbf{c}}^{(t)} &= \text{tanh}(\mathbf{U}_g \mathbf{h}^{(t-1)} + \mathbf{W}_g \mathbf{x}^{(t)} + \mathbf{b}_g) && \text{(candidate memory),} \\
 \mathbf{c}^{(t)} &= \mathbf{f}_g^{(t)} \circ \mathbf{c}^{(t-1)} + \mathbf{i}_g^{(t)} \circ \tilde{\mathbf{c}}^{(t)} && \text{(current memory),} \\
 \mathbf{o}_g^{(t)} &= \text{sigmoid}(\mathbf{U}_o \mathbf{h}^{(t-1)} + \mathbf{W}_o \mathbf{x}^{(t)} + \mathbf{b}_o) && \text{(output gate),} \\
 \mathbf{h}^{(t)} &= \mathbf{o}_g^{(t)} \circ \text{tanh}(\mathbf{c}^{(t)}) && \text{(current hidden state),} \\
 \forall t \in \{1, 2, \dots, k\}.
 \end{aligned} \tag{1}$$

$\{\mathbf{U}_{f,i,g,o}, \mathbf{W}_{f,i,g,o}, \mathbf{b}_{f,i,g,o}\}$ are trainable parameters, \circ denotes the Hadamard product (element-wise product), $\mathbf{h}^{(t)}$ and $\mathbf{c}^{(t)}$ are the hidden state and cell memory of a LSTM cell. Additionally, a LSTM cell uses four gates to manage its states over time to avoid the problem of exploding/vanishing gradients in the case of longer sequences [3]. $\mathbf{f}_g^{(t)}$ (forget gate) determines how much of the previous memory is kept, $\mathbf{i}_g^{(t)}$ (input gate) controls the amount new information is stored into memory, $\tilde{\mathbf{c}}^{(t)}$ (candidate memory) defines how much information is stored into memory and $\mathbf{o}_g^{(t)}$ (output gate) determines how much information is read out of the memory. The hidden state $\mathbf{h}^{(t)}$ is commonly forwarded to a successive layer.

4 Methodology

4.1 Time-Aware Long Short-Term Memory Cells

“Vanilla” LSTM cells, as described in Sect. 3.2, assume a uniform distribution of the elapsed time between events ($\Delta^{(t)} := x_{ts}^{(t)} - x_{ts}^{(t-1)}$). This assumption does not hold for most event logs analyzed in PBPM though (see Fig. 4). The elapsed time between consecutive events might have an impact on the next activity and timestamp prediction. Hence, a LSTM cell should be able to take irregular elapsed times into account when processing event logs.

Time-aware long short-term memory (T-LSTM) cells are an extension of the LSTM. Figure 1 depicts the T-LSTM cell and highlights its differences with regard to the LSTM cell.




Fig. 1. Illustration of a T-LSTM cell with its computational components at time step t . The dashed and blue components indicate the extensions to the “vanilla” LSTM cell. The previous cell memory $\mathbf{c}_S^{(t-1)}$ is adjusted to $\mathbf{c}_*^{(t-1)}$ (see Eq. (2)) and is then processed together with $\mathbf{h}^{(t-1)}$ and $\mathbf{x}^{(t)}$ via the LSTM computations, as formalized in Eq. (1).

The main idea behind T-LSTM is to perform a subspace decomposition of the previous cell memory $\mathbf{c}^{(t-1)}$. First, a short term memory component $\mathbf{c}_S^{(t-1)}$ is extracted via a network. Next, the short term memory is discounted via a decay function of the elapsed time and yields $\hat{\mathbf{c}}_S^{(t-1)}$. Then, the long term memory ($\mathbf{c}_T^{(t-1)} = \mathbf{c}^{(t-1)} - \mathbf{c}_S^{(t-1)}$) is calculated. Finally, the previous cell memory is adjusted $\mathbf{c}_*^{(t-1)} = \mathbf{c}_T^{(t-1)} + \hat{\mathbf{c}}_S^{(t-1)}$. The adjusted previous memory $\mathbf{c}_*^{(t-1)}$ is then, together with $\mathbf{h}^{(t-1)}$ and $\mathbf{x}^{(t)}$, further processed as in LSTM cells by substituting $\mathbf{c}^{(t-1)}$ with $\mathbf{c}_*^{(t-1)}$ in Eq. (1). The following equations summarize the T-LSTM specific computations for the subspace decomposition and adjustment of the previous memory.

$$\begin{aligned}
\mathbf{c}_S^{(t-1)} &= \tanh(\mathbf{W}_d \mathbf{c}^{(t-1)} + \mathbf{b}_d) && \text{(short term memory),} \\
\hat{\mathbf{c}}_S^{(t-1)} &= \mathbf{c}_S^{(t-1)} * \text{decay}(\Delta^{(t)}) && \text{(discounted short term memory),} \\
\mathbf{c}_T^{(t-1)} &= \mathbf{c}^{(t-1)} - \mathbf{c}_S^{(t-1)} && \text{(long term memory),} \\
\mathbf{c}_*^{(t-1)} &= \mathbf{c}_T^{(t-1)} + \hat{\mathbf{c}}_S^{(t-1)} && \text{(adjusted previous memory),} \\
&\dots && \text{(LSTM computations as in Eq. (1)),} \\
&\forall t \in \{1, 2, \dots, k\}.
\end{aligned} \tag{2}$$

Note, that we only add $\{\mathbf{W}_d, \mathbf{b}_d\}$ as trainable parameters compared to the LSTM cell. As recommended in Baytas et al. [2], we chose $\text{decay}(\Delta^{(t)}) = 1/\log(e + \Delta^{(t)})$ since we input the elapsed times in seconds and therefore have large values for $\Delta^{(t)}$. Any other monotonic decreasing function and scale for $\Delta^{(t)}$ would be valid as well, but our initial choice proved to be effective. The intuition behind the subspace decomposition is that the short term memory should be discounted if the elapsed time is very large, while the long term memory should be maintained in the adjusted previous cell memory $\mathbf{c}_*^{(t-1)}$. Similar as for LSTMs, the hidden state $\mathbf{h}^{(t)}$ is forwarded to successive layer for further processing. Hence, it is straightforward to substitute LSTM with T-LSTM cells in a given DNN architecture.

4.2 Network Architecture

We adapted the multitask architecture proposed by Tax et al. [15] as a baseline (see Fig. 2). The predicted next activity $\hat{e}_{k+1}.a$ is the output of a softmax activation after the last dense layer, where the output dimension is equal to the number of unique activity labels. $\hat{e}_{k+1}.a$ is evaluated against the one-hot encoded ground truth label $e_{k+1}.a$ by using the Cross-Entropy (CE) loss. The predicted next timestamp $\hat{e}_{k+1}.ts$ is a scalar output of a dense layer. We do not apply any additional activation after the time specific dense layer to be consistent with the implementation¹ of Tax et al. [15]. $\hat{e}_{k+1}.ts$ is compared with the ground truth timestamp $e_{k+1}.ts$ using the Mean Absolute Error (MAE). The total loss is the sum of both losses, as implemented in Tax et al. [15]. Further, they applied one-hot encoding for the activities and compute time-related control-flow features, which we also used in our experiments. We refer to the baseline architecture as “**Tax**”. We performed an ablation study and made three modifications to the baseline DNN architecture:

- We weighted the CE loss function based on the distribution of activity labels in the training set. Hence, the classification of under-represented event classes had larger influence during training. We refer to this model as “**Tax+CS**”.
- We replaced all LSTM layers with T-LSTM layers and refer to this model as “**Tax+T-LSTM**”.
- We added cost-sensitive learning and replaced all LSTM layers with T-LSTM layers. We call this model “**Tax+CS+T-LSTM**”.

¹ <https://github.com/verenich/ProcessSequencePrediction>.




Fig. 2. Network architecture for this work based on the multitask learning approach proposed by Tax et al. [15]. The dashed components are either LSTM or T-LSTM layers. The input is of the network is a sequence of vectors representing a prefix $\langle e_1, \dots, e_k \rangle$ as in Tax et al. [15]. For the baseline architecture we applied one-hot encoding and LSTM layers as in [15]. The outputs of the model are the predicted next activity ($\hat{e}_{k+1}.a$) and timestamp ($\hat{e}_{k+1}.ts$). Each of the LSTM layers is followed by a batch normalization layer (BN) to speed up training, as used in Tax et al. [15].

5 Experiments

5.1 Datasets

We performed our experiments on the same publicly available datasets as Tax et al. [15] to validate the effectiveness of our proposed techniques. Figure 3 shows the distribution of the activities (labels) for the different datasets. It is evident that the distributions of activities are skewed for both event logs. Table 1 presents descriptive statistics of the datasets used in this work.

Helpdesk²: This event log originates from a ticket management process of an Italian software company.

BPI’12 W Subprocess³ (**BPI12W**): The Business Process Intelligence (BPI) 2012 challenge provided this event log from a German financial institution. The data come from a loan application process. The ‘W’ indicates state of the work item for the application.

5.2 Preprocessing

We used the cleaned and prepared datasets as in Tax et al. [15]. The datasets can be found on the corresponding GitHub repository⁴. The preprocessing steps include splitting the data into training and test set, calculating time divisors, and ASCII encoding activities and sequence generation. Datasets were split into 2/3rd and 1/3rd for training and testing respectively and preserve the temporal order of cases. We additionally used the last 20% of the training data as a validation set in order to tune the hyperparameters. We adapted the sequence

² <https://doi.org/10.4121/uuid:0c60edf1-6f83-4e75-9367-4c63b3e9d5bb>.

³ <https://doi.org/10.4121/uuid:3926db30-f712-4394-aebc-75976070e91f>.

⁴ <https://github.com/verenich/ProcessSequencePrediction/tree/master/data>.




Fig. 3. Activity distribution in training and test set for Helpdesk and BPI12W datasets. It is evident that the distributions of the activity labels are skewed.

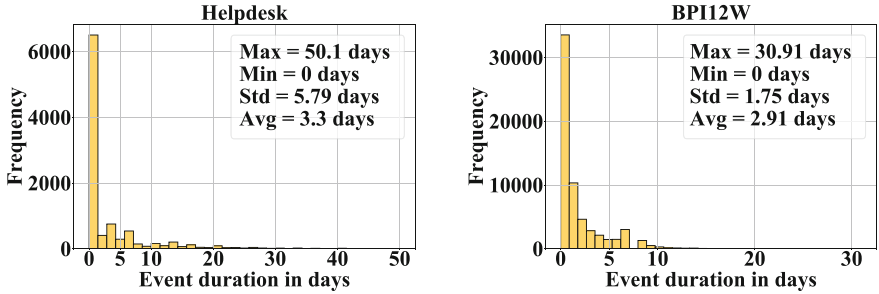


Fig. 4. Event duration distribution for the complete Helpdesk and BPI12W datasets. It can be observed that the majority of the events are completed within one day. However, there are many events with longer duration. Note that we input the elapsed time between events (Δ^t) in seconds for T-LSTM.

and feature generation methods by Tax et al. [15]. The features include the activity of the event, position of the event in the case, time since the last event, time from the starting event of the case, time from midnight, and day of the week. We create one-hot encoded versions of the ground truth labels $e_{k+1}.a$ for the next activity prediction in order to compare them with the predicted next activity labels $\hat{e}_{k+1}.a$.

5.3 Training Setup

For hyperparameter tuning, we performed a grid search on the training set and chose the model with the lowest validation loss. The validation loss is the sum of activity-related validation loss and time-related validation loss. The number of LSTM or T-LSTM units were set to 64 or 100. For the dropout rate (for both dense layers), we tried the values 0.0 and 0.2. We choose Nadam as an optimization algorithm, as used in [15]. Nesterov accelerated gradient (NAG) calculates the step using the ‘lookahead’ algorithm, which approximates the next parameters. Adam optimizer estimates learning rates based on ini-

Table 1. Descriptive statistics of the datasets used in this study.

Characteristic	Helpdesk	BPI12W
Number of instances	3,804	9,658
Case variants	154	2,263
Unique activities	9	6
Events	13,710	72,413
Max # events per case	14	74
Min # events per case	1	1
Avg # events per case	3.604	7.497

tial moments of the gradients. Nadam is a combination of both and is robust in noisy datasets. Furthermore, we tested a range of different learning rates $\{0.0001, 0.0002, 0.001, 0.002, 0.01\}$ since this is known to have a large impact on LSTMs [8]. We trained each model for 150 epochs, with a batch size of 64 and apply early stopping with patience 25 for regularization.

5.4 Evaluation

We applied the same evaluation metrics as in [15]. We used the *Accuracy* metric to evaluate the next activity prediction. For the next timestamp prediction, we used the *Mean Absolute Error (MAE)* to evaluate our models.

5.5 Implementation

We conducted all experiments on a workstation with 24 CPU cores, 748 GB RAM and a single GPU NVIDIA QUADRO RTX6000. We implemented the experiments in Python 3.7. We used the DL framework TensorFlow 2.1⁵. The source code is available on GitHub⁶.

6 Results

6.1 Next Activity Prediction

Table 2 shows the results for the next activity prediction in terms of Accuracy. For Helpdesk and BPI12W, the approach Tax+CS+T-LSTM achieved the highest Accuracy (0.724 and 0.778) among all approaches. The approach’s improvement compared to the baseline is 0.012 and 0.018. While the two approaches, Tax+CS and Tax+T-LSTM, outperformed the baseline for Helpdesk, these approaches achieved a lower Accuracy for BPI12W than the baseline.

⁵ <https://www.tensorflow.org>.

⁶ <https://github.com/annguy/time-aware-pbpm>.

Table 2. Results for the next activity prediction in terms of Accuracy. The best result for each dataset is highlighted (larger is better).

Approach	Helpdesk	BPI12W
Tax (baseline)	0.712	0.760
Tax+CS	0.713	0.757
Tax+T-LSTM	0.718	0.693
Tax+CS+T-LSTM	0.724	0.778

6.2 Next Timestamp Prediction

Table 3 shows the results for the next timestamp prediction task in terms of MAE in days. All approaches with a T-LSTM cell, clearly outperformed the baseline for both event logs. Thereby, the approach Tax+CS achieved the lowest MAE of 2.87 days and 0.88 days for Helpdesk and BPI12W respectively. Compared to the baseline, this approach reduced the MAE by 0.88 days (Helpdesk) and 0.68 days (BPI12W). The other two approaches, Tax+T-LSTM and Tax+CS+T-LSTM, achieved a slightly worse MAE values compared to Tax+CS for both event logs. It is worth noticing that for Helpdesk Tax+CS+T-LSTM and for BPI12W Tax+T-LSTM yielded the second best results with MAE close to Tax+CS.

Table 3. Results for next step time prediction in terms of MAE in days. The best result for each dataset is highlighted (lower is better).

Approach	Helpdesk	BPI12W
Tax (baseline)	3.75	1.56
Tax+CS	2.87	0.88
Tax+T-LSTM	3.01	0.88
Tax+CS+T-LSTM	2.94	0.90

7 Discussion

In this paper, we argued that the elapsed time between consecutive events carries valuable information on human behavior in running business processes. Therefore, we introduced T-LSTM cells for PBPM which inherently model the elapsed time between consecutive events. Further, we introduced of cost-sensitive learning to better cope with the problem of imbalanced data.

The obtained results indicate that the elapsed time between consecutive events is informative and that a DNN architecture relying on T-LSTM cells can yield more accurate models for PBPM. Especially, with the approach Tax+CS+T-LSTM, we could outperform the baseline (Tax) for both datasets

(i.e., Helpdesk and BPI12W) and both prediction tasks (i.e., next activity prediction and next timestamp prediction). Thereby, we could observe that cost-sensitive learning plays a crucial role for the predictive quality of a DNN architecture using T-LSTM cells instead of “vanilla” LSTM cells. Interestingly, the effectiveness of the introduced techniques is more evident for the next timestamp prediction compared to the next activity prediction

Even though our presented results on DNN architectures using T-LSTMs seem promising, there are a few limitations to our work. First, we need to verify our findings by performing experiments on more datasets. Second, a better hyperparameter tuning approach like Bayesian optimization [1] could be applied for all configurations to get a better estimate of their effectiveness. Further, several runs with random initialization should be performed to estimate the stability of the models.

8 Conclusion and Future Work

We propose T-LSTM as an alternative to the commonly used “vanilla” LSTM cell to better exploit information on the elapsed time between consecutive events. Furthermore, we introduced the concept of cost-sensitive learning to account for the common class-imbalance in event log data. Our results indicate the effectiveness of the introduced techniques for the next activity and timestamp prediction. This suggests that integrating specific mechanisms into neural network layers to incorporate event log specific characteristics might be an interesting direction for future research. Here, we mainly demonstrated the benefit of replacing a normal LSTM with a time-aware LSTM cell for a given baseline approach [15].

An avenue for future research is to investigate if T-LSTM cells might also improve other LSTM-based PBPM approaches such as Camargo et al. [5] involving resource attributes or Taymouri et al. [16] generating fake event logs. Another direction for future research is to further customize an LSTM cell in terms specifically for PBPM. For example, a process-aware LSTM cell could not only deal with time information but also with resource information.

Author contributions. Srijeet Chatterjee: Equal contribution with An Nguyen

References

1. Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M.: Optuna: a next-generation hyperparameter optimization framework. In: Proceedings of the 25rd International Conference on Knowledge Discovery and Data Mining (KDD) (2019)
2. Baytas, I.M., Xiao, C., Zhang, X., Wang, F., Jain, A.K., Zhou, J.: Patient subtyping via time-aware LSTM networks. In: Proceedings of the 23rd International Conference on Knowledge Discovery and Data Mining (KDD), pp. 65–74 (2017)
3. Bengio, Y., Simard, P., Frasconi, P., et al.: Learning long-term dependencies with gradient descent is difficult. *Trans. Neural Networks* **5**(2), 157–166 (1994)
4. Breuker, D., Matzner, M., Delfmann, P., Becker, J.: Comprehensible predictive models for business processes. *MIS Q.* **40**(4), 1009–1034 (2016)

5. 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
6. Di Francescomarino, C., Ghidini, C., Maggi, F.M., Milani, F.: Predictive process monitoring methods: which one suits me best? In: Weske, M., Montali, M., Weber, I., vom Brocke, J. (eds.) BPM 2018. LNCS, vol. 11080, pp. 462–479. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98648-7_27
7. Evermann, J., Rehse, J.R., Fettke, P.: Predicting process behaviour using deep learning. *Decis. Support Syst.* **100**, 129–140 (2017). <https://www.evermann2017predicting>
8. Greff, K., Srivastava, R.K., Koutník, J., Steunebrink, B.R., Schmidhuber, J.: LSTM: a search space odyssey. *IEEE Trans. Neural Networks Learn. Syst.* **28**(10), 2222–2232 (2017)
9. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
10. Khan, A., et al.: Memory-augmented neural networks for predictive process analytics. arXiv preprint [arXiv:1802.00938](https://arxiv.org/abs/1802.00938) (2018)
11. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436 (2015)
12. Maggi, F.M., Di Francescomarino, C., Dumas, M., Ghidini, C.: Predictive monitoring of business processes. In: Jarke, M., et al. (eds.) CAiSE 2014. LNCS, vol. 8484, pp. 457–472. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07881-6_31
13. Mehdiyev, N., Evermann, J., Fettke, P.: A novel business process prediction model using a deep learning method. *Bus. Inf. Syst. Eng.* **62**(2), 143–157 (2018). <https://doi.org/10.1007/s12599-018-0551-3>
14. Navarin, N., Vincenzi, B., Polato, M., Sperduti, A.: LSTM networks for data-aware remaining time prediction of business process instances. In: IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1–7. IEEE (2017)
15. 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
16. Taymouri, F., La Rosa, M., Erfani, S., Bozorgi, Z.D., Verenich, I.: Predictive business process monitoring via generative adversarial nets: the case of next event prediction. In: Proceedings of the 18th International Conference on Business Process Management (BPM) (2020)
17. Weinzierl, S., et al.: An empirical comparison of deep-neural-network architectures for next activity prediction using context-enriched process event logs. [arXiv:2005.01194](https://arxiv.org/abs/2005.01194) (2020)
18. Weinzierl, S., Dunzer, S., Zilker, S., Matzner, M.: Prescriptive business process monitoring for recommending next best actions. In: Fahland, D., Ghidini, C., Becker, J., Dumas, M. (eds.) BPM 2020. LNBIP, vol. 392, pp. 193–209. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58638-6_12