



HAL
open science

A Framework for Streaming Event-Log Prediction in Business Processes

Benedikt Bollig, Matthias Függer, Thomas Nowak

► **To cite this version:**

Benedikt Bollig, Matthias Függer, Thomas Nowak. A Framework for Streaming Event-Log Prediction in Business Processes. 2025. hal-04866045

HAL Id: hal-04866045

<https://hal.science/hal-04866045v1>

Preprint submitted on 6 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Framework for Streaming Event-Log Prediction in Business Processes^{*}

Benedikt Bollig¹[0000-0003-0985-6115], Matthias Függer¹[0000-0001-5765-0301],
and Thomas Nowak^{1,2}[0000-0003-1690-9342]

¹ Université Paris-Saclay, CNRS, ENS Paris-Saclay, LMF, Gif-sur-Yvette, France

² Institut Universitaire de France, Paris, France

Abstract. We present a Python-based framework for event-log prediction in streaming mode, enabling predictions while data is being generated by a business process. The framework allows for easy integration of streaming algorithms, including language models like n -grams and LSTMs, and for combining these predictors using ensemble methods.

Using our framework, we conducted experiments on various well-known process-mining data sets and compared classical batch with streaming mode. Though, in batch mode, LSTMs generally achieve the best performance, there is often an n -gram whose accuracy comes very close. Combining basic models in ensemble methods can even outperform LSTMs. The value of basic models with respect to LSTMs becomes even more apparent in streaming mode, where LSTMs generally lack accuracy in the early stages of a prediction run, while basic methods make sensible predictions immediately.

Keywords: event-log prediction · n -gram · LSTM · ensemble methods · streaming data · batch processing

1 Introduction

Digital data is omnipresent and continuously produced in today’s world. A particular type of data that is available to many organizations and institutions, either through explicit mining or implicitly within business processes, are event logs, i.e., timed traces of events recorded along the execution of a process. Examples are examinations of patients entering a hospital, server log files, and clicks of a user browsing a website.

While such event logs are often readily available at organizations, process models are typically absent. More so, quantitative models that allow one to describe, analyze, optimize, and predict possible future events are challenging to obtain.

^{*} The work was supported by the French National Research Agency (ANR) projects DREAMY (ANR-21-CE48-0003) and COSTXPRESS (ANR-23-CE45-0013), as well as the SAIF project, funded by the “France 2030” government investment plan managed by ANR, under the reference ANR-23-PEIA-0006.

The discipline of process mining [42, 43] aims at automatically discovering a model of the underlying process, reflecting its causalities and concurrency, with applications in streamlining it and finding potential bottlenecks. Classically, process mining is applied as a batch process to post-mortem data, with the discovery process being re-executed as more data is accumulated, and properties of the model being analyzed for potential optimization [35].

More recently, process mining has focused on prediction tasks [44]. In this context, a (language) model predicts the next activity of a given case based on its historical event data. Within the realm of event-log prediction, one distinguishes between two paradigms: batch learning and streaming learning, also referred to as offline and online learning. The choice of these approaches depends on the application as well as the kind of available data: In batch mode, a model is trained on post-mortem data; potentially including data that is different from what is seen in production: data may be annotated post-mortem with additional information. On the other side, streaming mode differs from batch mode in that: (i) Events are given one by one, interleaved across cases. (ii) Training is conducted online, i.e., in production. (iii) Predictions are made on pre-mortem data and can be exploited for decision-making in the running process, impacting and optimizing future trajectories. A major difference is that, in streaming mode, meaningful predictions are required early on when data is still sparse.

Activity prediction in discrete, case-based event logs. While process event logs can be diverse and include attributes of various types, in this work we consider them to consist of finitely many discrete activities, each belonging to a so called case. That is, an event consists of a case ID and an activity. An example from healthcare would be patients as cases and organizational, diagnostic, and treatment measures as activities. In the benchmarks used in this work, the number of activities ranges from 13 to 42. In the streaming case, however, the set of activities does not need to be known in advance.

Given a prefix of such an event log, we study the problem of predicting the most likely next activity for a case, including a distinguished `stop` symbol to signal the end of the case. A solution to this problem is a prediction function, also referred to as language model, which, given the prefix and the case, returns the predicted next activity. For simplicity, we only consider prediction functions whose output for a case depends solely on the historical event data for that specific case. For example, such functions do not capture global decisions within a hospital of adapting a diagnostic measure for a patient due to resource-limitations induced by other patients.

While various synthetic datasets have been used as benchmarks for learning algorithms for inferring probabilistic automata [52], it is difficult to draw conclusions from these performance metrics for real-life scenarios. Generating meaningful patterns randomly is already challenging when considering only single-case behaviors and it is not clear how the traces of several cases should be interleaved. By contrast, in real-life datasets, interleavings are often imposed by resource constraints or external factors, which may follow non-trivial patterns. We thus perform experiments based on 7 well known real-world process-mining datasets [39]

from different domains, including healthcare, finance, and IT service management. The number of events ranges from 15,214 to 2,514,266. These datasets have been previously used in [18] to establish experimental performance metrics for recurrent neural networks. Our results align with theirs in the context of neural networks, although the exact accuracies differ slightly.

Contributions. Our contributions are twofold. (i) We provide a Python framework that facilitates experiments for predictions in batch and streaming process mining. The framework allows base models to be easily assembled into composed models, which can themselves serve as components of other models, in both streaming and batch mode. (ii) Experiments conducted in our framework demonstrate the potential of simple base language models, especially when combined in ensemble methods. For all tested datasets, long short-term memory (LSTM) networks [19] outperform base models such as prefix trees and n -grams. However, n -grams with a relatively small n often comes close to the LSTM’s performance. The performance gap is further reduced (or even inverted) by combining base models using ensemble methods.

Related Work. Our work touches upon several fields including automata learning and grammatical inference [12, 40], process mining [42], and various contexts within machine learning.

Compared to classical automata learning in batch mode, grammatical inference for streaming data has received considerably less attention. Notable exceptions include [1, 2, 36], which propose efficient algorithms for probabilistic automata to address computational complexities in streaming settings. Non-quantitative incremental automata learning was explored in [15].

Event-log predictions in process mining is an active area of research. Recent works include [5, 10, 18, 32, 33, 45]. Process mining in a streaming setting has been studied, e.g., in [6, 7, 22, 51]. Note that [22] suggests using ensemble methods in the presence of noise, though not resorting to n -grams.

The use of n -grams dates back to [37] and has been widely applied to a multitude of problems including business processes [5].

Apart from prediction, another interesting application domain of automata learning is verification [23, 27, 28, 40]. For example, [27] focuses on learning deterministic probabilistic automata from batches, building on variants of the Alergia algorithm [8]. However, targeting reactive systems, the focus of this work is on infinite words rather than making predictions based on finite historical data.

Finally, several general-purpose automata learning libraries have been developed, including [4, 20, 30, 53].

Outline. We provide the technical background of the framework and the language models used in this paper in Section 2: prefix trees, n -grams, probabilistic deterministic finite automata [54], and LSTMs [19]. The computational Python framework is introduced in Section 3. The experimental setup and results are presented in Section 4. We conclude in Section 5, discussing potential future directions.

2 Language Models and Ensemble Learning

We start with a unified, automata-based view of various base language models, including frequency prefix trees [12], n -grams [37], and probabilistic deterministic finite automata [54]. This not only allows one to compare the above approaches within a single setting but has also been adapted within our Python framework, where concrete language models are obtained by refining basic automata provided by the framework.

2.1 Base Language Models

We fix a nonempty finite set Σ of activities, also referred to as activities. The set is not necessarily known in advance to a language model, though an algorithm (e.g., a neural network) may require specifying an upper bound on its size (as part of the embedding dimension in case of the neural network). We write Σ^* for the set of finite sequences of activities. In particular, the set includes the empty sequence ε . Activities are referred to as a, b , etc. We denote by **stop** a symbol that is not contained in Σ and that is used by a language model to signal the end of a sequence. The set $\Sigma_{\text{stop}} = \Sigma \cup \{\text{stop}\}$ contains all activities and the stop symbol. In the following, let σ, τ be from Σ_{stop} . Activities are associated to cases. We assume that case IDs are taken from a countably infinite set *CaseID*. Specifically, in our benchmarks, case IDs will be from the set of strings or the set of natural numbers.

A language model can take several forms. Examples are models based on n -grams, bags, probabilistic automata, and recurrent neural networks. In essence, every such language model defines a *probabilistic prediction function* or, simply, prediction function $p: \Sigma^* \rightarrow \Delta(\Sigma_{\text{stop}})$, where $\Delta(\Sigma_{\text{stop}})$ is the set of probability distributions over Σ_{stop} . Given a sequence $w \in \Sigma^*$, applying p yields a probability distribution over the set of possible next activities, including **stop**. In analogy to conditional probabilities, one usually writes $p(\sigma | w)$ for the probability $p(w)(\sigma)$. The focus in this work is on predicting the most likely next activity, which is readily obtained from a probabilistic prediction function. Note that, under certain conditions, the prediction function also allows one to generate sequences according to a probability distribution over Σ^* .

As previously stated, we follow an automata-based approach to obtain prediction functions. This is motivated by two reasons: (i) It unifies widely-used approaches, and (ii) its formalism is close to streaming learning with incremental state updates, allowing us to compare batch and streaming mode within a single framework. Given that state updates due to learning will be deterministic, we will use a deterministic automaton to describe these. Following [28], we define:

Definition 1. A probabilistic deterministic finite automaton (*PDFA*) is composed of a finite set S of states with a dedicated initial state $s_0 \in S$, a partial deterministic transition function $\delta: S \times \Sigma \rightarrow S$ (also called *update function*), and a mapping $\pi: S \rightarrow \Delta(\Sigma_{\text{stop}})$ assigning a probability distribution to every state.

With a PDFA \mathcal{A} , we associate a prediction function p inductively as follows. We extend the automaton’s transition function δ to $\hat{\delta}: S \times \Sigma^* \rightarrow S$ over sequences in the usual way by letting $\hat{\delta}(s, \varepsilon) = s$ and $\hat{\delta}(s, aw) = \hat{\delta}(\delta(s, a), w)$ for all $s \in S$, $a \in \Sigma$, and $w \in \Sigma^*$. The *prediction for* $w \in \Sigma^*$ is then given by $p(w) = \pi(\hat{\delta}(s_0, w))$. Figure 1c depicts a PDFA. Its initial state is the topmost state. We have, e.g., $p(\text{stop} \mid \varepsilon) = 0$, $p(\text{stop} \mid aa) = \frac{7}{13}$, $p(a \mid aa) = \frac{5}{13}$, and $p(b \mid aa) = \frac{1}{13}$.

Towards a data structure that provides natural ways to update transition functions, we use a related concept of frequencies rather than probabilities on activities [12]: A *frequency deterministic finite automaton* (FDFA) is defined analogously to a probabilistic deterministic automaton, the only difference being that, instead of a mapping $\pi: S \rightarrow \Delta(\Sigma_{\text{stop}})$, it is equipped with a mapping $f: S \rightarrow \mathbb{N}^{\Sigma_{\text{stop}}}$. Intuitively, for state $s \in S$ and activity $\sigma \in \Sigma_{\text{stop}}$, frequency $f(s)(\sigma)$ is the number of times σ was observed while being in state s (i.e., after parsing a sequence ending up in s). In particular, $f(s)(\text{stop})$ is the number of times a sequence stopped in s . An example FDFA is depicted in Figure 1b. For the initial (topmost) state s , we have $f(s)(\text{stop}) = 0$, $f(s)(a) = 13$, and $f(s)(b) = 17$.

For each such FDFA, we obtain a corresponding PDFA with identical states and transition function and π defined by $\pi(s)(\sigma) = f(s)(\sigma) / \sum_{\tau \in \Sigma_{\text{stop}}} f(s)(\tau)$. Here, we assume $\sum_{\tau \in \Sigma_{\text{stop}}} f(s)(\tau) \geq 1$ for all $s \in S$ which is without restriction for the automata constructed via learning. Function π is usually constructed on demand for a given state when a prediction is due or an activity is to be sampled. To simplify notation, we will use an FDFA and its corresponding PDFA synonymously.

We can now express the classical language models within this framework. A *frequency prefix tree* (FPT) is an FDFA that organizes a finite set of sequences into a tree structure, where prefixes of arbitrary length are identified with states (Figure 1a). For $n \geq 1$, an n -gram is an FDFA whose states s are identified with the last $n - 1$ activities called the *access string* of s , i.e., the suffix of length $n - 1$ for words of length at least $n - 1$ and the word itself for shorter words (see Figure 1b for a 3-gram). For example, starting from an initial state (the root in Figure 1b), the two sequences ab and $aaab$ share the same suffix ab of length $n - 1 = 2$, and lead to the same state (the one with **stop**-frequency 1). The access string of this state is ab . A *bag* keeps track of the set of activities seen until now. It is thus an FDFA whose states are subsets of the set Σ of possible activities.

2.2 Batch Learning

We next define the prediction problem for batch learning and then adapt the solutions to the streaming case.

In batch learning, we are given an event log $L \subseteq (\text{CaseID} \times \Sigma)^*$, which is a finite set of finite sequences of case ID-activity pairs. The goal is to train a language model on L , i.e., to construct a prediction function $p: \Sigma^* \rightarrow \Delta(\Sigma_{\text{stop}})$ that generalizes well and enables informed predictions based on event data not necessarily seen before. Recall that, in our setting, a prediction for a given case ID only depends on the projection of the event log to this case. Consequently, it is sufficient

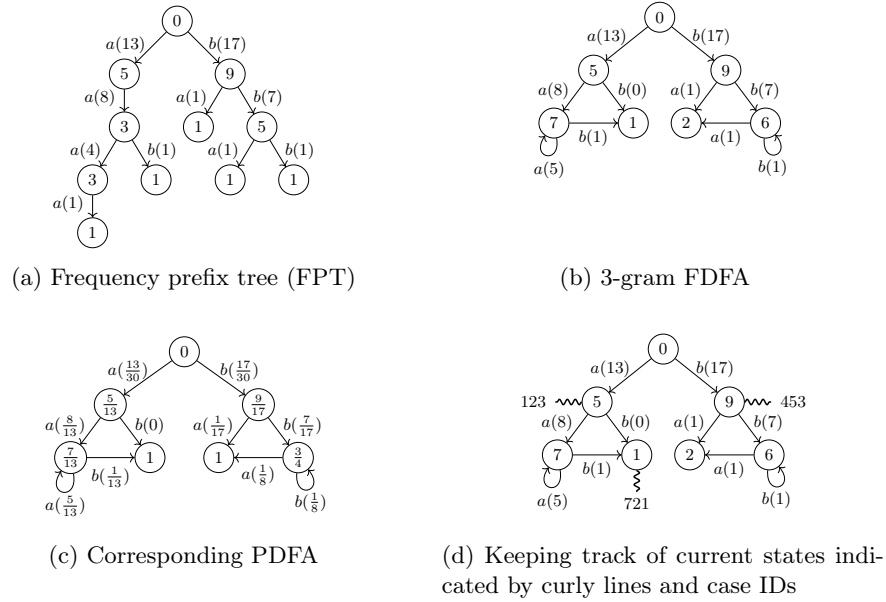


Fig. 1: Automata obtained for the event log given by the set $L = \{\langle a \rangle^5, \langle aa \rangle^3, \langle aaa \rangle^3, \langle aab \rangle^1, \langle aaaa \rangle^1, \langle b \rangle^9, \langle ba \rangle^1, \langle bb \rangle^5, \langle bba \rangle^1, \langle bbb \rangle^1\}$, with multiplicities of a sequence denoted by powers. States (nodes) and transitions upon activities (arrows) are shown. Frequencies, respectively, probabilities for activities are indicated in brackets next to activities and within nodes for the stop activity. (a) FPT (b) FDFA for the 3-gram, as well as (c) PDFFA for the 3-gram. (d) FDFA for the 3-gram enriched with the current states of cases 123, 453, and 721 as maintained during inference and streaming learning.

to maintain the event order of an event log L within a case, rather than across all cases. This allows us to represent L as a finite multiset over Σ^* , i.e., as a mapping $L: \Sigma^* \rightarrow \mathbb{N}$ with finite support. An example of L over the activity set $\Sigma = \{a, b\}$ is $L = \{\langle a \rangle^5, \langle aa \rangle^3, \langle aaa \rangle^3, \langle aab \rangle^1, \langle aaaa \rangle^1, \langle b \rangle^9, \langle ba \rangle^1, \langle bb \rangle^5, \langle bba \rangle^1, \langle bbb \rangle^1\}$ with powers denoting the multiplicity within the set. Here, $L(\varepsilon) = 0$, $L(a) = 5$, $L(bb) = 5$, etc. The multiset contains 30 sequences (cases).

Most batch automata learning algorithms conceptually proceed in two steps [12]. First, the event log is arranged as an FPT (Figure 1a). In a second step, this FPT is reduced (or, folded) by identifying, i.e., merging some of the states. The specificity of an algorithm lies in the second step:

FPT. The simplest algorithm directly uses the FPT (and the derived PDFA) as is, without any state equivalences. While this approach does not generalize well, particularly because no prediction can be made for unseen prefixes, it serves as a baseline and a potential component for ensemble methods.

Alergia. A commonly used merging algorithm is Alergia [9]. The algorithm successively merges nodes in a top-down fashion, with each merge consisting

of two stages. In the first stage, the algorithm identifies two states that can be merged, typically using a statistical test such as Hoeffding’s incompatibility test [9, 13]. Two states s and s' are considered equivalent (i.e., “mergeable”) if their probability distributions $p(s)$ and $p(s')$ do not differ significantly, and if their respective successors are equivalent, too. In the second stage, the algorithm performs the actual merging of the states and resolves any non-determinism that arises by merging subsequent states and transitions.

Gram. Finally, an FPT can be reduced to obtain an n -gram. For this purpose, two states are considered equivalent if the sequence of the last $n - 1$ activities in their corresponding access strings is identical (Figures 1b and 1c). Intuitively, the difference between state merging in Algeria and n -grams is that Algeria considers the potential future and n -grams the past of a sequence to decide if two states are equivalent.

To optimize computational complexities, implementations, e.g., of n -grams, directly construct the FDFA without detouring through an FPT. We detail this in Section 2.3 when discussing streaming learning.

Further, we assumed that L is given as a multiset of cases. While this may be the case for batch learning, inference in production will likely have to cope with interleaved cases. If so, one simply tracks, for every active case ID, the current state in the FDFA (Figure 1d).

Finally, we note that an LSTM which receives as input a sequence of activities of a case and as output provides a probability distribution over the case’s next activities, can be viewed as a probabilistic automaton (with a potentially large number of states).

2.3 Streaming Learning

In streaming learning, one is given an event stream, where cases are interleaved, rather than an event log, and learning and prediction alternate. Formally, an event stream is a finite or infinite sequence of events in $CaseID \times \Sigma$. While processing the stream, after any activity, we aim to predict the next activity for each active case. For that purpose, after each activity, we perform an update of the automaton and query it to return a distribution over the next activities. For an n -gram, the steps are:

Initialization. The FDFA has a single root node ρ corresponding to the access string ε and we set the current state of each (potential) case to ρ .

Update. Upon receiving an event (id_i, a_i) , with $i \geq 1$, from the event stream, the FDFA is updated as follows. Let s be the current state of case id_i . If the case id_i has not been encountered before, i.e., the current state of the case is ρ , increment $f(s)(\text{stop})$. We then decrement $f(s)(\text{stop})$, as one fewer sequence stops in s , increment $f(s)(a_i)$, as one more sequence continues with a_i , and increment $f(\delta(s, a_i))(\text{stop})$, provided $\delta(s, a_i)$ exists. We also update the current state of case id_i to $s' = \delta(s, a_i)$. Observe that $f(s')(\text{stop}) > 0$ at the end of the update step, so that a potential subsequent decrement does not result in a negative $f(s')(\text{stop})$.

An example is given in Figure 1d: Assume a new event $(123, a)$ is received. Case 123 has been encountered before, and we let s be its current state. After the update step, $f(s)(\text{stop}) = 4$, $f(s)(a) = 9$, and $f(\delta(s, a))(\text{stop}) = 8$. If, on the other hand, $\delta(s, a)$ is not defined, a new transition is created (and possibly a new state). To illustrate this, consider again Figure 1d and suppose an event $(721, a)$ arrives. As the current state of 721 has the access string ab , the last two activities seen were ab , and so the new history of length 2 becomes ba . A new transition is created to the state with access string ba and the counts are adjusted as described above.

Query. Upon a query for $id \in \text{CaseID}$, the probability distribution induced by $f(s)$ at the current state s of case id is returned.

Both steps are local and computationally efficient on automata as a data structure. By contrast to the updates in n -gram, this locality is not ensured by Alergia.

2.4 Ensemble Methods

Ensemble methods [56] combine multiple base models, such as FPTs or n -grams, into more complex models using voting mechanisms. Voting mechanisms are distinguished based on how they handle updates and inference queries. In both *soft voting* and *hard voting*, each base model individually processes updates from incoming events. Similarly, queries are forwarded to all base models, each of which produces a prediction in the form of a probability distribution. The key difference lies in how these distributions are aggregated: In soft voting, the final probability distribution is obtained by averaging all individual distributions, while in hard voting (also referred to as *majority voting*), each probability distribution is converted into a concrete activity, and the activity chosen most frequently constitutes the final verdict (which may be represented as a Dirac probability distribution).

While hard and soft voting apply to both batch and streaming modes, *adaptive voting* is specific to streaming mode. The idea is to track the current accuracy of each base model. A query is then forwarded to the model with the highest accuracy. An update on an incoming event involves two steps. In step 1, the accuracy is adjusted: the activity a of the event is compared to the prediction of each base model. Base models that correctly predict a see their accuracy increase, while the accuracy of all other models decreases. Step 2 updates the base models in the same manner as in soft and hard voting.

Finally, we use ensemble methods as a *fallback method* to compensate for slowly-learning language models in streaming learning. For example, we apply the FPT with a minimum number of 10 visits³ whenever possible. If the minimum number has not been reached, as a fallback option, the prediction of an n -gram is returned, where, for our experiments, we use $n = 5$ or $n = 7$ depending on the dataset.

³ This is a hyperparameter that can be adjusted according to the required confidence level needed.

We will show in Section 4 that soft voting generally outperforms hard voting. Most importantly, soft voting offers considerable improvements over automata-based base models and, in streaming mode, even outperforms LSTMs on many datasets.

3 Computational Framework

A wide range of stream-processing frameworks is available, with examples being Dataflow [11], Faust [16], Flink [17], Pathway [31], as well as River [29]. Towards the goal of a combined batch and stream processing with easy general-purpose function modules, local deployment, analysis via dashboards, low overhead and latencies, and stop-resume functionalities, we developed the `logicsponge` Python library [24].

A data pipeline in `logicsponge` is made up of terms that are connected via sequential or parallel composition. We distinguish two basic types of terms: source terms and function terms. The output of a source term does not explicitly depend on any other terms in the pipeline. On the other hand, a function term receives data from other terms as input and can produce output to other downstream terms. A function term that does not produce any outputs for other terms is a sink. It may, of course, produce outputs not modeled in the pipeline, like appending data to files or user console output. In the sequential composition $t1 * t2$ of terms $t1$ and $t2$, the input of $t2$ is equal to the output of $t1$. In their parallel composition $t1 | t2$, both $t1$ and $t2$ receive the same upstream input and its downstream output is the union of outputs of $t1$ and $t2$, accessible as separate streams. Function terms that merge these streams into single streams are provided by the framework. Independently of where they appear in the pipeline, all terms are executed in separate threads.

Inputs and outputs of terms are handled via data streams. A data stream encapsulates a list of data items. A data item is a Python dictionary with some associated metadata. Each data stream is owned by a unique term to which it is associated as an output. The owner of a data stream can append data items to it. A data stream can be the input to other terms as defined by the sequential and parallel composition. Terms that have a data stream as an input do not have write access to it and can access it only via a read-only data stream view object. The underlying list data structure of a data stream is protected from concurrent access via readers-writer locks [34], minimizing conflicts. When starting a pipeline, setting an optional `persistent` function argument configures the contents of the data streams to be stored in a Python object database [57]. Restarting the pipeline after an interrupt or crash restores the state from the database and resumes computation. The library contains predefined terms to generate plots, compute basic statistics, and create interactive web dashboards.

We implemented the process-mining algorithms used for this work as function terms in the `logicsponge-processmining` package [25]. We first define a list of language models:

```
models=[
    BasicMiner(algorithm=Bag()),
    BasicMiner(algorithm=FrequencyPrefixTree()),
    BasicMiner(algorithm=NGram(window_length=2)),
    BasicMiner(algorithm=NGram(window_length=3)),
]
```

This list is then provided to a term that performs soft voting on the models:

```
soft_voting = StreamingActivityPredictor(
    strategy=SoftVoting(models=models)
)
```

Additionally, we create a term that acts as an LSTM model:

```
lstm = StreamingActivityPredictor(
    strategy=NeuralNetworkMiner(
        model=LSTMModel(vocab_size, *dim, device=device),
        criterion=nn.CrossEntropyLoss(),
        optimizer=optim.Adam(model.parameters(), lr=0.001),
        batch_size=8,
    )
)
```

Both terms are combined in parallel with a subsequent respective evaluation term that generates statistics on their accuracy and performance. The pipeline's source term generates a data stream out of the dataset. The data stream is subsequently processed by applying a key filter and adding start symbol before it is fed into the parallel predictor terms:

```
pipeline = (
    ListStreamer(data_list=dataset)
    * ls.KeyFilter(keys=["case_id", "activity"])
    * AddStartSymbol()
    * (
        (soft_voting * Evaluation("soft_voting"))
        | (lstm * Evaluation("lstm"))
    )
)
```

Finally, the pipeline is started:

```
pipeline.start()
```

4 Experimental Evaluation

To evaluate our framework and approach, we performed experiments based on seven well-known real-world business-process datasets (Table 1). All experiments were performed on an 11th Gen Intel Core i9-11900K architecture (3.50GHz,

Table 1: Dataset statistics and experimental setup (voting uses FPT + bag + the three n -grams specified below)

	Sepsis Cases [26]	BPI 2012 [46]	BPI 2013 [38]	BPI 2014 [47]	BPI 2017 [48]	BPI 2018 [50]	BPI 2019 [49]
#Activities	16	24	13	39	26	41	42
#Cases	1,050	13,087	7,554	46,616	31,509	43,809	251,734
Avg. case length	14.49	20.04	8.68	10.01	38.16	57.39	6.34
#Events	15,214	262,200	65,533	466,737	1,202,267	2,514,266	1,595,923
Voting: (FPT+bag+...)	3, 4, 5	3, 5, 7	3, 4, 5	3, 4, 5	3, 5, 7	3, 5, 7	3, 5, 7
Fallback: FPT \rightarrow ...	5-gram	7-gram	5-gram	5-gram	7-gram	7-gram	7-gram

8 cores, 24GB RAM) with an NVIDIA GeForce RTX 3090 (24GB) running Ubuntu (24.04.1) and Python (3.12.3). For the LSTM training and inference we used PyTorch with CUDA. We measured accuracy in both batch and streaming modes. In both modes, we used incremental inference based on the current state of a case. In our comparison, we included n -grams with n ranging from 1 to 8, FPTs, bags, and LSTMs. The LSTMs consisted of two hidden layers of dimension 128, preceded by an embedding layer of dimension 50. Every voting method used five submodels, including one FPT, one bag, and three n -grams with varying combinations of n (Table 2). The fallback method is instantiated with an FPT (considering only states with at least 10 total visits) and an n -gram with $n = 5$ or $n = 7$.

4.1 Batch Mode

In batch mode, a given dataset (event log) was divided into a training set (70%), a validation set (15%), and a test set (15%), where the percentages refer to the number of sequences in the event log. To account for **stop**-predictions, we added a **stop** symbol to the end of each sequence. The validation set was used only for LSTM training to determine a stopping criterion. We averaged the results over 5 runs. We performed at most 20 epochs per run with a batch size of 8, but stopped a run as soon as a model failed to improve for 3 consecutive epochs. Per run, we then selected the LSTM model with the best validation accuracy.

Every model was evaluated on the test set. For every activity a in the dataset, a given model selected the outcome \hat{a} with the highest probability based on its current state as its prediction. Thus, the prediction was correct if $a = \hat{a}$. Accuracy measured the proportion of correct predictions.

For n -grams, we implemented backoff [21]: if the next activity cannot be parsed by an n -gram, the model attempts to parse shorter and shorter suffixes, always starting from the initial state. The test results for Alergia were obtained for significance parameter 0.5, using the automata-learning library AALpy [30].

Table 2: Comparison of batch and streaming learning across datasets

Batch	Sepsis Cases		BPI 2012		BPI 2013		BPI 2014		BPI 2017		BPI 2018		BPI 2019	
	Acc. (%)	#States	Acc. (%)	#States	Acc. (%)	#States	Acc. (%)	#States	Acc. (%)	#States	Acc. (%)	#States	Acc. (%)	#States
FPT	43.26	4935	68.63	45277	54.01	13861	33.01	152402	70.32	206498	37.71	1038724	65.75	187421
bag	56.61	157	73.89	319	55.79	299	43.09	8994	64.25	445	41.21	2598	71.85	2545
1-gram	20.86	1	20.15	1	41.50	1	17.31	1	17.18	1	18.26	1	17.02	1
2-gram	55.90	17	65.08	25	60.58	14	38.44	40	66.87	27	57.00	42	60.10	43
3-gram	57.78	127	81.70	147	70.42	92	47.52	768	74.23	199	63.03	621	72.33	644
4-gram	59.76	472	83.66	420	70.71	410	49.10	6003	83.71	684	68.63	3664	74.00	3681
5-gram	61.81	1250	84.82	876	70.39	1267	48.50	25516	87.04	1682	71.90	13673	74.46	12577
6-gram	59.85	2625	85.57	1552	69.79	2987	46.10	69870	87.37	3422	74.86	39192	74.46	29742
7-gram	56.20	4680	85.66	2408	68.62	5778	42.43	139750	87.46	6116	75.87	93049	74.36	54556
8-gram	52.80	7278	85.63	3576	67.16	9662	38.73	225371	87.49	10003	76.34	190524	74.08	83467
fallback	61.82	N/A	85.68	N/A	70.55	N/A	48.57	N/A	87.59	N/A	76.59	N/A	74.37	N/A
hard voting	64.09	N/A	85.43	N/A	70.81	N/A	49.75	N/A	86.32	N/A	74.52	N/A	74.76	N/A
soft voting	64.20	N/A	85.55	N/A	70.99	N/A	49.98	N/A	87.42	N/A	76.36	N/A	74.76	N/A
Alergia	50.95	47	75.25	136	55.72	32	31.54	74	73.48	480	N/A	N/A	70.76	373
LSTM	64.00	N/A	85.71	N/A	70.86	N/A	51.90	N/A	88.49	N/A	82.14	N/A	75.27 ^a	N/A
Streaming	Acc. (%)	Latency (ms)	Acc. (%)	Latency (ms)	Acc. (%)	Latency (ms)	Acc. (%)	Latency (ms)	Acc. (%)	Latency (ms)	Acc. (%)	Latency (ms)	Acc. (%)	Latency (ms)
FPT	42.84	0.04	66.39	0.06	51.09	0.04	30.97	0.09	68.92	0.04	35.90	0.04	63.90	0.14
bag	57.95	0.09	75.86	0.06	52.68	0.04	41.67	0.07	58.43	0.05	41.69	0.04	71.58	0.05
1-gram	22.21	0.13	20.91	0.06	46.14	0.06	18.96	0.30	17.42	0.05	18.44	0.05	15.70	0.07
2-gram	56.57	0.09	67.78	0.06	56.34	0.06	39.97	0.06	68.66	0.06	56.80	0.05	57.92	0.05
3-gram	57.76	0.08	81.37	0.06	67.17	0.04	44.97	0.07	74.19	0.04	62.94	0.04	72.00	0.05
4-gram	58.98	0.08	83.34	0.07	67.19	0.04	45.85	0.06	83.79	0.06	68.55	0.04	73.81	0.05
5-gram	60.37	0.13	84.48	0.06	66.49	0.04	44.46	0.07	87.12	0.05	71.73	0.04	74.08	0.08
6-gram	57.80	0.12	85.15	0.06	65.22	0.05	41.55	0.06	87.44	0.05	74.45	0.04	73.89	0.06
7-gram	53.82	0.16	85.15	0.06	63.47	0.04	38.03	0.11	87.49	0.05	75.12	0.05	73.54	0.08
8-gram	50.18	0.08	85.05	0.07	61.46	0.05	34.99	0.06	87.46	0.05	75.15	0.05	73.16	0.05
fallback	60.40	0.05	85.13	0.06	66.66	0.05	44.48	0.07	87.58	0.06	76.06	0.05	73.55	0.08
hard voting	63.68	0.28	84.97	0.58	67.23	0.17	46.05	0.19	84.91	0.15	74.22	0.16	74.26	0.20
soft voting	64.80	0.17	85.20	0.11	67.46	0.07	46.84	0.13	87.50	0.09	74.39	0.07	74.56	0.12
adaptive voting	60.37	0.28	85.15	0.23	67.19	0.15	45.84	0.22	87.49	0.17	75.12	0.15	74.08	0.23
LSTM	61.66	6.41	84.34	4.46	66.33	7.00	44.39	7.83	87.27	7.21	80.15 ^b	10.03	75.75 ^c	30.14

^a only one run and no patience counter^b stopped after 8 hours and 2,351,500 events^c stopped after 14 hours and 1,471,201 events

4.2 Streaming Mode

In streaming mode, we did not add `stop`-symbols to mark the end of a sequence, as adding them to pre-mortem data is unsuitable unless explicitly present in the dataset. On the other hand, to handle the first event of a case in a data stream, we insert an “init” symbol to mark the start of a sequence allowing an LSTM to make a prediction from the very beginning.

For querying an LSTM, we parse the current sequence through the model and retrieve the corresponding outcome. LSTMs are not inherently designed for streaming learning. We update a model efficiently with each incoming event as follows: After each new event, we perform a training iteration. Each training step uses a batch size of 8, always including the updated case sequence. The remaining sequences are selected in a round-robin fashion to ensure all sequences are considered repeatedly.

Note that, unlike in batch learning, there is no separation of the dataset in training, validation, and test set.

4.3 Results


Experimental results are depicted in Table 2. In batch mode, LSTMs perform best across most datasets. However, in most cases, an n -gram model comes very close. Combining these models, along with frequency prefix trees and bags, into ensemble methods further boosts the performance of simple base models: in two cases, LSTMs are even outperformed (by soft voting).

In streaming mode, LSTMs appear to struggle during the initial stage of a stream. As illustrated in Figure 2a for the sepsis dataset, which depicts the accuracy evolution for various models, LSTM performance is initially well below that of other base models, catching up only later (Figure 2b). Ensemble methods can outperform LSTMs on several datasets. Notably, the average latency (the total time required for making one prediction and one update) of base models is significantly lower than that of LSTMs.


5 Conclusion and Future Work

We presented a computational framework for event-log predictions in streaming mode. Experimental results on business processes carried out in this framework suggest that process data often exhibits behavior that is consistent with computationally efficient n -gram predictions. Moreover, ensemble methods considerably boost smaller base models, at the expense of only little computational overhead.

One direction for future research is refining the methodology. Fallback algorithms show promising potential, but more detailed experiments are required to explore more sophisticated fallback strategies that rely, for example, on statistical tests. A potential starting point is the work in automata learning [12]) to compare states as well as the work on probabilistic bisimilarity distances [41]. It would be interesting to investigate whether ensemble methods involving nested



(a) First 550 predictions



(b) Predictions for the complete dataset

Fig. 2: Prediction accuracy for different language models in streaming mode over the number of activities in the sepsis dataset from [26].

models (such as voting among other voting models) could further improve performance. In streaming mode, more fine-grained adaptive methods, particularly concerning LSTM training, may help reduce latency and avoid overfitting.

From a theoretical perspective, it would be worthwhile to explain the n -gram phenomenon: Can a random automaton or process often be approximated by an n -gram? What is the optimal n , or what are the best combinations of n 's in ensemble methods? Can these be determined online? Is it worthwhile to distribute them sparsely, or are clusters of n preferable? In other words, what is the best way to determine the optimal hyperparameters?

As object-centric process mining attracts increasing attention [55], incorporating global behaviors into predictions gains in importance. In terms of automata, a natural model that distinguishes between global and local behaviors is data automata [3], which have already been studied in a classification learning setting [14]. Frequency-enriched and probabilistic versions are hypothesized to be good starting points to pursue this direction beyond the local predictors considered in this paper. Another extension involves incorporating more attributes or activities, including continuous ones such as energy consumption, clinical parameters, or timestamps and activity durations [18, 45].

References

1. B. Balle, J. Castro, and R. Gavaldà. Adaptively learning probabilistic deterministic automata from data streams. *Mach. Learn.*, 96(1-2):99–127, 2014.
2. R. Baumgartner and S. Verwer. Learning state machines from data streams: A generic strategy and an improved heuristic. In F. Coste, F. Ouardi, and G. Rabusseau, editors, *International Conference on Grammatical Inference, ICGI 2023, 10-13 July 2023, Rabat, Morocco*, volume 217 of *Proceedings of Machine Learning Research*, pages 117–141. PMLR, 2023.
3. M. Bojanczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on data words. *ACM Trans. Comput. Log.*, 12(4):27:1–27:26, 2011.
4. B. Bollig, J. Katoen, C. Kern, M. Leucker, D. Neider, and D. R. Piegdon. libalf: The automata learning framework. In T. Touili, B. Cook, and P. B. Jackson, editors, *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, volume 6174 of *Lecture Notes in Computer Science*, pages 360–364. Springer, 2010.
5. D. Breuker, M. Matzner, P. Delfmann, and J. Becker. Comprehensible predictive models for business processes. *MIS Q.*, 40(4):1009–1034, 2016.
6. A. Burattin. Streaming process mining. In W. M. P. van der Aalst and J. Carmona, editors, *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*, pages 349–372. Springer, 2022.
7. A. Burattin, A. Sperduti, and W. M. P. van der Aalst. Control-flow discovery from event streams. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2014, Beijing, China, July 6-11, 2014*, pages 2420–2427. IEEE, 2014.
8. R. C. Carrasco and J. Oncina. Learning stochastic regular grammars by means of a state merging method. In *International Colloquium on Grammatical Inference*, pages 139–152. Springer, 1994.
9. R. C. Carrasco and J. Oncina. Learning stochastic regular grammars by means of a state merging method. In R. C. Carrasco and J. Oncina, editors, *Grammatical Inference and Applications, Second International Colloquium, ICGI-94, Alicante, Spain, September 21-23, 1994, Proceedings*, volume 862 of *Lecture Notes in Computer Science*, pages 139–152. Springer, 1994.
10. M. Ceci, P. F. Lanotte, F. Fumarola, D. P. Cavallo, and D. Malerba. Completion time and next activity prediction of processes using sequential pattern mining. In S. Dzeroski, P. Panov, D. Kocev, and L. Todorovski, editors, *Discovery Science - 17th International Conference, DS 2014, Bled, Slovenia, October 8-10, 2014. Proceedings*, volume 8777 of *Lecture Notes in Computer Science*, pages 49–61. Springer, 2014.
11. Google Cloud Dataflow. <https://cloud.google.com/products/dataflow>. Accessed: 2024-12-04.
12. C. de la Higuera. *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press, USA, 2010.
13. C. de la Higuera and F. Thollard. Identification in the limit with probability one of stochastic deterministic finite automata. In A. L. Oliveira, editor, *Grammatical Inference: Algorithms and Applications, 5th International Colloquium, ICGI 2000, Lisbon, Portugal, September 11-13, 2000, Proceedings*, volume 1891 of *Lecture Notes in Computer Science*, pages 141–156. Springer, 2000.
14. N. Decker, P. Habermehl, M. Leucker, and D. Thoma. Learning transparent data automata. In G. Ciardo and E. Kindler, editors, *Application and Theory of Petri Nets and Concurrency - 35th International Conference, PETRI NETS 2014, Tunis*,

- Tunisia, June 23-27, 2014. Proceedings*, volume 8489 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2014.
15. P. Dupont. Incremental regular inference. In L. Miclet and C. de la Higuera, editors, *Grammatical Inference: Learning Syntax from Sentences, 3rd International Colloquium, ICGI-96, Montpellier, France, September 25-27, 1996, Proceedings*, volume 1147 of *Lecture Notes in Computer Science*, pages 222–237. Springer, 1996.
 16. Faust. <https://github.com/faust-streaming/faust>. Accessed: 2024-12-04.
 17. Apache Flink. <https://flink.apache.org/>. Accessed: 2024-12-04.
 18. M. Hinkka, T. Lehto, and K. Heljanko. Exploiting event log event attributes in RNN based prediction. In P. Ceravolo, M. van Keulen, and M. T. Gómez-López, editors, *Data-Driven Process Discovery and Analysis - 8th IFIP WG 2.6 International Symposium, SIMPDA 2018, Seville, Spain, December 13-14, 2018, and 9th International Symposium, SIMPDA 2019, Bled, Slovenia, September 8, 2019, Revised Selected Papers*, volume 379 of *Lecture Notes in Business Information Processing*, pages 67–85. Springer, 2019.
 19. S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
 20. M. Isberner, F. Howar, and B. Steffen. The open-source learnlib - A framework for active automata learning. In D. Kroening and C. S. Pasareanu, editors, *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, volume 9206 of *Lecture Notes in Computer Science*, pages 487–495. Springer, 2015.
 21. S. M. Katz. Estimation of probabilities from sparse data for the language model component of a speech recognizer. *IEEE Trans. Acoust. Speech Signal Process.*, 35(3):400–401, 1987.
 22. B. Krawczyk and A. Cano. Online ensemble learning with abstaining classifiers for drifting and noisy data streams. *Appl. Soft Comput.*, 68:677–692, 2018.
 23. M. Leucker. Learning meets verification. In F. S. de Boer, M. M. Bonsangue, S. Graf, and W. P. de Roever, editors, *Formal Methods for Components and Objects, 5th International Symposium, FMCO 2006, Amsterdam, The Netherlands, November 7-10, 2006, Revised Lectures*, volume 4709 of *Lecture Notes in Computer Science*, pages 127–151. Springer, 2006.
 24. logicsponge. <https://github.com/innatelogic/logicsponge>. Accessed: 2024-12-04.
 25. logicsponge-processmining. <https://github.com/innatelogic/logicsponge-processmining>. Accessed: 2024-12-04.
 26. F. Mannhardt. Sepsis Cases - Event Log. <https://doi.org/10.4121/uuid:915d2bfb-7e84-49ad-a286-dc35f063a460>, 2016.
 27. H. Mao, Y. Chen, M. Jaeger, T. D. Nielsen, K. G. Larsen, and B. Nielsen. Learning deterministic probabilistic automata from a model checking perspective. *Mach. Learn.*, 105(2):255–299, 2016.
 28. F. Mayr, S. Yovine, M. Carrasco, F. Pan, and F. Vilensky. A congruence-based approach to active automata learning from neural language models. In F. Coste, F. Ouardi, and G. Rabusseau, editors, *International Conference on Grammatical Inference, ICGI 2023, 10-13 July 2023, Rabat, Morocco*, volume 217 of *Proceedings of Machine Learning Research*, pages 250–264. PMLR, 2023.
 29. J. Montiel, M. Halford, S. M. Mastelini, G. Bolmier, R. Sourty, R. Vaysse, A. Zouitine, H. M. Gomes, J. Read, T. Abdessalem, and A. Bifet. River: machine learning for streaming data in python. *J. Mach. Learn. Res.*, 22(1), Jan. 2021.
 30. E. Muskardin, B. K. Aichernig, I. Pill, A. Pferscher, and M. Tappler. AALpy: an active automata learning library. *Innov. Syst. Softw. Eng.*, 18(3):417–426, 2022.

31. Pathway. <https://github.com/pathwaycom/pathway>. Accessed: 2024-12-04.
32. M. Pegoraro, M. S. Uysal, D. B. Georgi, and W. M. P. van der Aalst. Text-aware predictive monitoring of business processes. In W. Abramowicz, S. Auer, and E. Lewanska, editors, *24th International Conference on Business Information Systems, BIS 2021, Hannover, Germany, June 15-17, 2021*, pages 221–232, 2021.
33. M. Polato, A. Sperduti, A. Burattin, and M. de Leoni. Time and activity sequence prediction of business process instances. *Computing*, 100(9):1005–1031, 2018.
34. M. Raynal. *Concurrent Programming: Algorithms, Principles, and Foundations*. Springer, Heidelberg, 2012.
35. L. Reinkemeyer. Process mining in action. *Process mining in action principles, use cases and outlook*, 2020.
36. J. Schmidt and S. Kramer. Online induction of probabilistic real-time automata. *J. Comput. Sci. Technol.*, 29(3):345–360, 2014.
37. C. E. Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
38. W. Steeman. BPI Challenge 2013, incidents. <https://doi.org/10.4121/uuid:500573e6-acc-4b0c-9576-aa5468b10cee>, 2013.
39. TF-PM. BPI Challenges: 10 Years of Real-Life Datasets. <https://www.tf-pm.org/newsletter/newsletter-stream-2-05-2020/bpi-challenges-10-years-of-real-life-datasets>, May 2020. Accessed: 2024-12-04.
40. F. W. Vaandrager. Model learning. *Commun. ACM*, 60(2):86–95, 2017.
41. F. van Breugel. Probabilistic bisimilarity distances. *ACM SIGLOG News*, 4(4):33–51, 2017.
42. W. M. P. van der Aalst. *Process Mining - Data Science in Action, Second Edition*. Springer, 2016.
43. W. M. P. van der Aalst and J. Carmona, editors. *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*. Springer, 2022.
44. W. M. P. Van der Aalst, M. Pesic, and M. Song. Beyond process mining: From the past to present and future. In *Advanced Information Systems Engineering: 22nd International Conference, CAiSE 2010, Hammamet, Tunisia, June 7-9, 2010. Proceedings 22*, pages 38–52. Springer, 2010.
45. W. M. P. van der Aalst, M. H. Schonenberg, and M. Song. Time prediction based on process mining. *Inf. Syst.*, 36(2):450–475, 2011.
46. B. van Dongen. BPI Challenge 2012. <https://doi.org/10.4121/uuid:3926db30-f712-4394-aebc-75976070e91f>, 2012.
47. B. van Dongen. BPI Challenge 2014: Activity log for incidents. <https://doi.org/10.4121/uuid:86977bac-f874-49cf-8337-80f26bf5d2ef>, 2014.
48. B. van Dongen. BPI Challenge 2017. <https://doi.org/10.4121/uuid:5f3067df-f10b-45da-b98b-86ae4c7a310b>, 2017.
49. B. van Dongen. BPI Challenge 2019. <https://doi.org/10.4121/uuid:d06aff4b-79f0-45e6-8ec8-e19730c248f1>, 2019.
50. B. van Dongen and F. Borchert. BPI Challenge 2018. <https://doi.org/10.4121/uuid:3301445f-95e8-4ff0-98a4-901f1f204972>, 2018.
51. S. J. van Zelst, B. F. van Dongen, and W. M. P. van der Aalst. Event stream-based process discovery using abstract representations. *Knowl. Inf. Syst.*, 54(2):407–435, 2018.
52. S. Verwer, R. Eyraud, and C. de la Higuera. Pautomac: a probabilistic automata and hidden markov models learning competition. *Mach. Learn.*, 96(1-2):129–154, 2014.

53. S. Verwer and C. A. Hammerschmidt. flexfringe: A passive automaton learning package. In *2017 IEEE International Conference on Software Maintenance and Evolution, ICSME 2017, Shanghai, China, September 17-22, 2017*, pages 638–642. IEEE Computer Society, 2017.
54. E. Vidal, F. Thollard, C. de la Higuera, F. Casacuberta, and R. C. Carrasco. Probabilistic finite-state machines-part I. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(7):1013–1025, 2005.
55. W. Zhou, A. Polyvyanyy, and J. Bailey. PELP: pioneer event log prediction using sequence-to-sequence neural networks. *CoRR*, abs/2312.09741, 2023.
56. Z.-H. Zhou. *Ensemble Methods: Foundations and Algorithms*. Chapman & Hall/CRC, 1st edition, 2012.
57. ZODB. <https://github.com/zopectfoundation/ZODB>. Accessed: 2024-12-04.