



Wayeb: a Tool for Complex Event Forecasting

Elias Alevizos^{1,3}, Alexander Artikis^{2,1}, and Georgios Paliouras¹

¹ NCSR Demokritos, Athens, Greece

{alevizos.elias,a.artikis,paliourg}@iit.demokritos.gr

² University of Piraeus, Piraeus, Greece

³ National Kapodistrian University of Athens, Athens, Greece

Abstract

Complex Event Processing (CEP) systems have appeared in abundance during the last two decades. Their purpose is to detect in real-time interesting patterns upon a stream of events and to inform an analyst for the occurrence of such patterns in a timely manner. However, there is a lack of methods for forecasting when a pattern might occur before such an occurrence is actually detected by a CEP engine. We present Wayeb, a tool that attempts to address the issue of Complex Event Forecasting. Wayeb employs symbolic automata as a computational model for pattern detection and Markov chains for deriving a probabilistic description of a symbolic automaton.

1 Introduction

A Complex Event Processing (CEP) system takes as input a stream of events, along with a set of patterns, defining relations among the input events, and detects instances of pattern satisfaction, thus producing an output stream of complex events. Typically, an event has the structure of a tuple of values which might be numerical or categorical, with the *event type* and *timestamp* being the most common attributes. Since time is of critical importance for CEP, a temporal formalism is used in order to define the patterns to be detected. Such a pattern imposes temporal (and possibly atemporal) constraints on the input events, which, if satisfied, lead to the detection of a complex event. Efficient processing is of paramount importance since complex events must be detected with very strict latency requirements. Moreover, both the input events and the patterns may carry a certain degree of uncertainty, e.g., due to noisy sensors or an incomplete knowledge of the domain. These issues have been the focus of CEP research for the past two decades. See [8, 2] for surveys of (probabilistic) CEP. In this paper, we focus on another dimension of CEP research, that of Complex Event Forecasting (CEF), i.e., the ability of a CEP system to provide forecasts about the possible occurrence of complex events in the future. Although some conceptual proposals towards this direction have appeared in the literature [13, 11], there still remains a lack of concrete algorithms and systems that can actually perform CEF. In what follows, we present Wayeb, a CEP/F engine that, besides detecting CEP patterns, it can also forecast when such patterns may occur.

2 Forecasting with Classical Automata

In order to allow for a self-contained presentation, we begin by briefly describing how our proposed method works with classical automata (for details, see [1]). In this case, we assume that the stream $S = t_1, t_2, \dots$ is sequence of symbols from a finite alphabet $\Sigma = \{e_1, \dots, e_r\}$, i.e., $t_i \in \Sigma$. A pattern is defined as a (classical) regular expression R and the usual tools of standard automata theory may be employed [16]. The finite automaton used for event detection is the one corresponding to the expression $\Sigma^* \cdot R$ (\cdot denotes concatenation and $*$ Kleene-star), since it should work on streams and be able to start the detection at any point. Appending Σ^* at the beginning allows the automaton to skip any number of events. As a next step, we construct the deterministic finite automaton (DFA) for $\Sigma^* \cdot R$, $DFA_{\Sigma^* \cdot R}$. This allows us to convert $DFA_{\Sigma^* \cdot R}$ to a Markov chain. If we assume that the stream is composed of i.i.d. events from Σ , then the sequence $Y = Y_0, Y_1, \dots, Y_i, \dots$, where $Y_0 = q_0$ and $Y_i = \delta(Y_{i-1}, t_i)$, with q_0 the automaton's start state and δ its transition function, is a 1-order Markov chain [21]. Such a Markov chain, associated with a pattern R , is called a Pattern Markov Chain (PMC). The transition probability between two states connected with the symbol e_j is simply its occurrence probability, $P(X_i=e_j)$. If we assume that the process generating the stream is of a higher order $m \geq 1$, then we must first convert $DFA_{\Sigma^* \cdot R}$ to an m -unambiguous DFA, i.e., to an equivalent DFA where each state can "remember" the last m symbols. This can be achieved by iteratively duplicating those states of $DFA_{\Sigma^* \cdot R}$ for which we cannot unambiguously determine the last m symbols that can lead to them and then convert it to a PMC, denoted by PMC_R^m [20, 21].

After constructing PMC_R^m , we can estimate its waiting-time distributions. The waiting-time $W_R(q)$ for each non-final state q of $DFA_{\Sigma^* \cdot R}$ is a random variable, defined as the number of transitions until $DFA_{\Sigma^* \cdot R}$ visits for the first time one of its final states: $W_R(q) = \inf\{n: Y_0, Y_1, \dots, Y_n, Y_0=q, q \in Q \setminus F, Y_n \in F\}$. We compute the distribution of $W_R(q)$ by converting each state of PMC_R^m that corresponds to a final state of the DFA into an absorbing state and then re-organizing the transition matrix as follows: $\mathbf{\Pi} = \begin{pmatrix} \mathbf{N}^{(l-k) \times (l-k)} & \mathbf{C}^{(l-k) \times k} \\ \mathbf{0}_{k \times (l-k)} & \mathbf{I}_{k \times k} \end{pmatrix}$,

assuming PMC_R^m has a total of l states, k of which are final. \mathbf{N} holds the probabilities for all the possible transitions between (and only between) the non-final states, whereas \mathbf{C} holds the transition probabilities from non-final to final states. Then, the probability for the time index n when the system first enters the set of absorbing states is given by [12]: $P(Y_n \in A, Y_{n-1} \notin A, \dots, Y_1 \notin A \mid \boldsymbol{\xi}_{init}) = \boldsymbol{\xi}^T \mathbf{N}^{n-1} (\mathbf{I} - \mathbf{N}) \mathbf{1}$, where A denotes the set of absorbing states, $\boldsymbol{\xi}_{init}$ is the initial distribution on the states and $\boldsymbol{\xi}$ consists of the $l - k$ elements of $\boldsymbol{\xi}_{init}$ corresponding to non-absorbing states. Since, in our case, the current state of $DFA_{\Sigma^* \cdot R}$ is known, the vector $\boldsymbol{\xi}_{init}^T$ would have 1.0 as the value for the element corresponding to the current state (and 0 elsewhere). $\boldsymbol{\xi}$ changes dynamically as the DFA/PMC moves among its various states and every state has its own $\boldsymbol{\xi}$, denoted by $\boldsymbol{\xi}_q$:

$$\boldsymbol{\xi}_q(i) = \begin{cases} 1.0 & \text{if row } i \text{ of } \mathbf{N} \text{ corresponds to state } q \\ 0 & \text{otherwise} \end{cases}$$

The probability of $W_R(q)$ is then given by: $P(W_R(q) = n) = \boldsymbol{\xi}_q^T \mathbf{N}^{n-1} (\mathbf{I} - \mathbf{N}) \mathbf{1}$

The transition matrix $\mathbf{\Pi}$ is learnt using the maximum-likelihood estimators for its elements [3]: $\hat{\pi}_{i,j} = \frac{n_{i,j}}{\sum_{k \in Q} n_{i,k}} = \frac{n_{i,j}}{n_i}$ where n_i denotes the number of visits to state i , $n_{i,j}$ the number of transitions from state i to state j and $\pi_{i,j}$ the transition probability from state i to state j . After estimating the transition matrix, we compute the waiting-time distributions and then build the forecasts associated with each state. Each forecast is an interval $I = (start, end)$ and

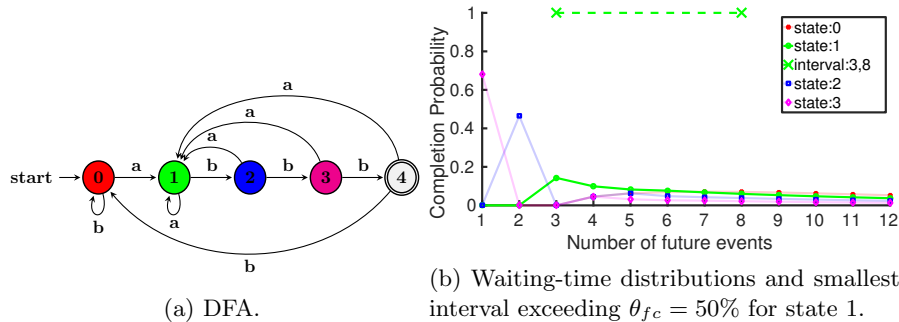


Figure 1: DFA and waiting-time distributions for $R = a \cdot b \cdot b \cdot b$, $\Sigma = \{a, b\}$, $m = 0$.

its meaning is the following: given that the DFA is in a certain state, we forecast that it will have reached one of its final states at some future point between *start* and *end*, with probability at least θ_{fc} . The calculation of this interval is done by using the waiting-time distribution that corresponds to each state and the threshold θ_{fc} is set beforehand by the user. We use a single-pass algorithm in order to scan the distribution of each state and find the *smallest* interval whose probability, i.e., the sum of probabilities of points included in the interval, exceeds θ_{fc} . In this paper, we assume stationarity, thus waiting-time distributions are computed once, after matrix estimation. An example of how forecasts are produced is shown in Figure 1. Figure 1a shows a simple DFA that can detect the pattern $R = a \cdot b \cdot b \cdot b$ on a stream composed of a and b events. For its four non-final states, Figure 1b shows their corresponding waiting-time distributions. For state 1, this figure also shows the forecast interval produced by scanning the corresponding (green and highlighted) distribution, when $\theta_{fc} = 50\%$. If the DFA moves to another state, then another distribution will be activated and a different forecast interval will be produced (these other intervals are not shown to avoid cluttering).

3 Forecasting with Symbolic Automata

One limitation of the method presented above is that it requires a finite alphabet, i.e., the input stream may be composed only of a finite set of symbols; on the other hand, events in a stream are in the form of tuples and look more like data words. Their attributes might be real-valued, taking values from an infinite set, and thus cannot be directly handled by DFA. In order to overcome this limitation, we have extended our method so that symbolic automata are employed [26, 10]. Symbolic automata, instead of having symbols from a finite set on their transitions, are equipped with predicates from a Boolean algebra, acting as transition guards. Such predicates can reference any event attribute, thus allowing for more expressive patterns.

We now present a formal definition of symbolic automata (for details, see [10]).

Definition 1 (Symbolic Automaton). *A symbolic finite automaton (SFA) is a tuple $M = (\mathcal{A}, Q, q^0, F, \Delta)$, where \mathcal{A} is an effective Boolean algebra, Q is a finite set of states, $q^0 \in Q$ is the initial state, $F \subseteq Q$ is the set of final states and $\Delta \subseteq Q \times \Psi_{\mathcal{A}} \times Q$ is a finite set of transitions, with $\Psi_{\mathcal{A}}$ being a set of predicates closed under the Boolean connectives.*

Most other definitions from classical automata carry over symbolic automata. Importantly, SFA are determinizable [10]. The definition for deterministic SFA is also similar to that for classical automata, with the important difference that it is not enough for all transitions from

a state to have different predicates. We require that predicates on transitions from the same state are mutually exclusive, i.e., at most one may evaluate to TRUE (see again [10]). The determinization process for *SFA* is similar to that for classical automata and is based on the construction of the power-set of the states of the non-deterministic automaton. Due to this issue of different predicates possibly both evaluating to TRUE for the same event/tuple, we first need to create the *minterms* of the predicates of a *SFA*, i.e., the set of maximal satisfiable Boolean combinations of such predicates. When constructing a deterministic *SFA*, these minterms are used as guards on the transitions, since they are mutually exclusive.

This result about *SFA* being determinizable allows us to use same technique of converting a deterministic automaton to an m -unambiguous automaton and then to a Markov chain. First, note that the set of minterms constructed from the predicates appearing in a *SFA* M , denoted by $Minterms(Predicates(M))$, induces a finite set of equivalence classes on the (possibly infinite) set of domain elements of M [10]. For example, if $Predicates(M) = \{\psi_1, \psi_2\}$, then $Minterms(Predicates(M)) = \{\psi_1 \wedge \psi_2, \psi_1 \wedge \neg\psi_2, \neg\psi_1 \wedge \psi_2, \neg\psi_1 \wedge \neg\psi_2\}$ and we can map each domain element, which, in our case, is an event/tuple, to exactly one of these 4 minterms: the one that evaluates to TRUE when applied to the element. Similarly, the set of minterms induces a set of equivalence classes on the set strings (event streams in our case). For example, if $S=t_1, \dots, t_n$ is an event stream, then it could be mapped to $S'=a, \dots, b$, with a corresponding to $\psi_1 \wedge \neg\psi_2$ if $\psi_1(t_1) \wedge \neg\psi_2(t_1) = \text{TRUE}$, b to $\psi_1 \wedge \psi_2$, etc. We say that S' is the stream induced by applying $Minterms(Predicates(M))$ on the original stream S . We first give a definition for an m -unambiguous deterministic *SFA*, by modifying the relevant definition for classical automata [21]:

Definition 2 (m -unambiguous deterministic *SFA*). *A deterministic SFA (A, Q, q^0, F, Δ) is m -unambiguous if there exist $q \in Q$ and $a, b \in T^m$ such that $a \neq b$ and $\delta(q, a) = \delta(q, b)$. A deterministic SFA which is not m -unambiguous is m -unambiguous.*

In other words, upon reaching a state q of an m -unambiguous *SFA*, we know which last m minterms evaluated to TRUE, i.e., we know the last m symbols of the induced stream S' . The following proposition then follows:

Proposition 1. *Let M be a deterministic m -unambiguous SFA, $S = t_1, \dots, t_n$ an event stream and $S' = X_1, \dots, X_n$ the stream induced by applying $T = Minterms(Predicates(M))$ on S . Assume that S' is a m -order Markov process, i.e., $P(X_i = a \mid X_1, \dots, X_{i-1}) = P(X_i = a \mid X_{i-m}, \dots, X_{i-1})$. Then the sequence $Y = Y_m, \dots, Y_n$ defined by $Y_0 = q^0$ and $Y_i = \delta(Y_{i-1}, t_i)$ (i.e., the sequence of states visited by M) is a 1-order Markov chain whose transition matrix is given by:*

$$\Pi(p, q) = \begin{cases} P(X_{m+1} = a \mid X_1, \dots, X_m = \phi(\delta^{-m}(p))) & \text{if } \delta(p, a) = q \\ 0 & \text{if } q \notin \delta(p, T) \end{cases}$$

where $\delta^{-m}(q) = \{a \in T^m \mid \exists p \in Q : \delta(p, a) = q\}$ is the set of concatenated labels of length m that can lead to q and, by definition, is a singleton for m -unambiguous *SFA*.

Proof. This result holds for classical m -unambiguous deterministic automata [21]. For the symbolic case, note that, from an algebraic point of view, the set $T = Minterms(Predicates(M))$ may be treated as a generator of the monoid T^* , with concatenation as the operation. If the cardinality of T is k , then we can always find a set $\Sigma = \{a_1, \dots, a_k\}$ of k distinct symbols and then a morphism (in fact, an isomorphism) $\phi : T^* \rightarrow \Sigma^*$ that maps each minterm to exactly one, unique a_i . A classical deterministic automaton N can then be constructed by relabelling

the *SFA* M under ϕ , i.e., by copying/renaming the states and transitions of the original *SFA* M and by replacing the label of each transition of M by the image of this label under ϕ . Then, the behavior of N (the language it accepts) is the image under ϕ of the behavior of M [24]. Based on these observations, we can see that the sequence of states visited by an m -unambiguous deterministic *SFA* is indeed a 1-order Markov chain, as a direct consequence of the fact that a deterministic *SFA* has the same behavior (up to isomorphism) to that of a classical deterministic automaton, constructed through relabelling. \square

Therefore, after constructing a deterministic *SFA* we can use it to construct a PMC and learn its transition matrix. The conditional probabilities in this case are essentially probabilities of seeing an event that will satisfy a predicate, given that the previous event(s) have satisfied the same or other predicates. For example, if $m = 1$ and $Predicates(M) = \{\psi_1, \psi_2\}$, one such conditional probability would be $P(\psi_1(t_{i+1}) \wedge \psi_2(t_{i+1}) = \text{TRUE} \mid \neg\psi_1(t_i) \wedge \neg\psi_2(t_i) = \text{TRUE})$, i.e., the probability of seeing an event that will satisfy both ψ_1 and ψ_2 given that the current, last seen event satisfies neither of these predicates. As with classical automata, we can then provide again forecasts based on the waiting-time distributions of the PMC derived from a *SFA*.

4 Experimental Results

Wayeb is a Complex Event Forecasting engine based on symbolic automata, written in the Scala programming language. It was tested against two real-world datasets coming from the field of maritime monitoring. When sailing at sea, (most) vessels emit messages relaying information about their position, heading, speed, etc.: the so-called AIS (automatic identification system) messages. Such a stream of AIS messages can then be used in order to detect interesting patterns in the behavior of vessels [22]. Two AIS datasets were used, made available in the dataAcron project¹: the first contains AIS kinematic messages from vessels sailing in the Atlantic Ocean around the port of Brest, France, and span a period from 1 October 2015 to 31 March 2016 [23]; the second was provided by IMISG² and contains AIS kinematic messages from most of Europe (the entire Mediterranean Sea and parts of the Atlantic Ocean and the Baltic Sea), spanning an one-month period from 1 January 2016 to 31 January 2016. AIS messages can be noisy, redundant and typically arrive at unspecified time intervals. We first processed our datasets in order to produce clean and compressed trajectories, consisting of critical points, i.e., important points that are a summary of the initial trajectory, but allow for an accurate reconstruction [22]. Subsequently, we sampled the compressed trajectories by interpolating between critical points in order to get trajectories where each point has a temporal distance of one minute from its previous point. After excluding points with null speed (in order to remove stopped vessels), the final streams consist of ≈ 1.3 million points for the Brest dataset and ≈ 2.4 million points for the IMISG dataset. The experiments were run on a machine with Intel Core i7-4770 CPU @ 3.40GHz processors and 16 GB of memory.

We have chosen to demonstrate our forecasting engine on two important patterns. The first concerns a movement pattern in which vessels approach a port and the goal is to forecast when a vessel will enter the port. This is a key target of several vessel tracking software platforms, as indicated by the 2018 challenge of the International Conference on Distributed and Event-Based Systems [15]. The second pattern concerns a fishing maneuver inside a fishing area. Forecasting when vessels are about to start fishing could be important in order to manage the pressure exerted on fishing areas [17].

¹<http://datacron-project.eu/>

²<https://imisglobal.com/>

The *approaching* pattern may be defined as follows:

$$\begin{aligned}
 R_1 := & x \cdot y^+ \cdot z \text{ WHERE} \\
 & \textit{Distance}(x, \textit{PortCoords}, 7.0, 10.0) \text{ AND } \textit{Distance}(y, \textit{PortCoords}, 5.0, 7.0) \text{ AND} \\
 & \textit{WithinCircle}(z, \textit{PortCoords}, 5.0) \text{ PARTITION BY } \textit{vesselId}
 \end{aligned} \tag{1}$$

Concatenation is denoted by \cdot and $+$ stands for Kleene+. We want to start detecting a vessel’s movement whenever a vessel is between 7 and 10 km away from a specific port, then it approaches the port, with its distance from it falling to the range of 5 to 7 km, stays in that range for 1 or more messages, and finally enters the port, defined as being inside a circle with a radius of 5 km around the port. We have chosen the above syntax in order to make clear what the regular part of a pattern is (before the WHERE keyword) and what its logical part is (after WHERE), but predicates can be placed directly in the regular expression. Note also that the first argument of a predicate is the event upon which it is to be applied, but we also allow for other arguments to be passed as constants. The *partition contiguity* selection strategy is employed [28], i.e., for each new vessel appearing in the stream, a new automaton run is created, being responsible for this vessel. This is just a special case of parametric trace slicing typically used in runtime verification tools [7]. If we use this pattern to build a PMC, the transition probabilities will involve the three predicates that appear in it. However, it is reasonable to assume that other features of a vessel’s kinematic behavior could affect the accuracy of forecasts, e.g., its speed. For this reason, we have also added a mechanism to our module that can incorporate in the PMC some extra features, declared by the user, but not present in the pattern itself. The extra features we decided to add concern the vessel’s speed and its heading: *SpeedBetween*($x, 0, 10.0$), *SpeedBetween*($x, 10.0, 20.0$), *SpeedBetween*($x, 20.0, 30.0$) and *HeadingTowards*($x, \textit{PortCoords}$). The first three try to use the speed level of a vessel, whereas the last one uses the vessel’s heading and checks whether it is headed towards the port. Our experiments were conducted using the main Brest port as the port of reference.

The *fishing* pattern may be defined as follows:

$$\begin{aligned}
 R_2 := & x \cdot y^* \cdot z \text{ WHERE} \\
 & (\textit{IsFishingVessel}(x) \wedge \neg \textit{InArea}(x, \textit{FishingArea})) \text{ AND} \\
 & (\textit{InArea}(y, \textit{FishingArea}) \wedge \textit{SpeedBetween}(y, 9.0, 20.0)) \text{ AND} \\
 & (\textit{InArea}(z, \textit{FishingArea}) \wedge \textit{SpeedBetween}(z, 1.0, 9.0)) \\
 & \text{PARTITION BY } \textit{vesselId}
 \end{aligned} \tag{2}$$

This definition attempts to capture a movement of a fishing vessel, in which it is initially outside a specific fishing area, then enters the area with a traveling speed (between 9 to 20 knots), remains there for zero or more messages, and finally starts moving with a fishing speed (between 1 to 9 knots) while still in the same area. We also used the same extra features as in Pattern 1, with the difference that the *Heading* feature now concerns the area. Note that by partitioning by *vesselId* we don’t have to add the *IsFishingVessel* predicate for y and z .

Figure 2 shows *precision* and *spread* results on the Brest dataset for different values of the forecasting confidence threshold θ_{fc} and of the assumed order m . Precision is defined as the percentage of forecasts that were correct, i.e., whose interval includes the timestamp of the complex event’s actual occurrence. Spread is defined as the length of a forecast interval. The smaller the spread, the more focused a forecast is and thus more valuable. Good results are therefore considered those with high precision scores (ideally 1.0) and low spread scores (ideally 0, i.e., single point forecasts). For each confidence threshold, each bar represents a different

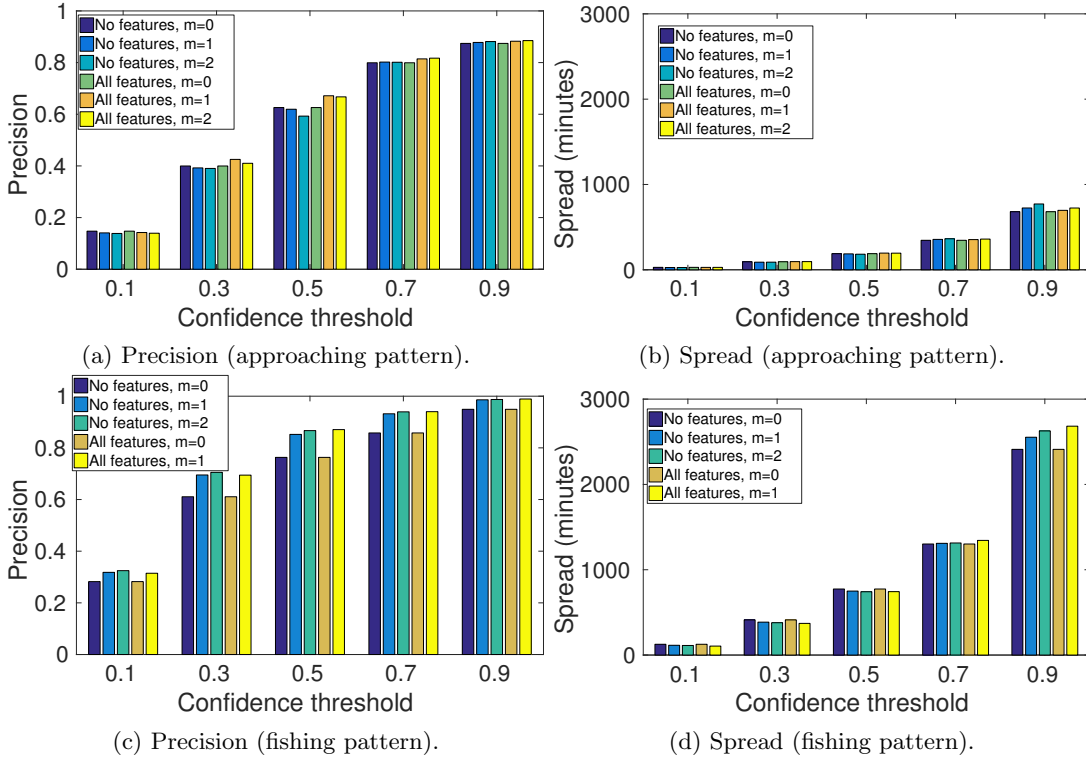


Figure 2: Results for approaching the Brest port and for fishing.

variation of the pattern. We vary the order m of the PMC in order to investigate the possible gains of looking deeper into the past. There are also variations where no extra features are used (i.e., only the predicates of the original pattern are included in the PMC) and where all the extra features are included. Our engine can always achieve precision scores that are above the confidence threshold. However, note that, as the confidence threshold increases, the spread also tends to increase. This happens because the engine tries to satisfy the θ_{fc} constraint by expanding the forecast intervals. Looking back at the example of Figure 1b, if, instead of $\theta_{fc} = 50\%$, we set $\theta_{fc} = 70\%$, then the green interval shown at the top would have to be expanded to the right so that its probability exceeds the new increased threshold (expanding to the left would be pointless since only points with zero probability exist in this region). It is interesting to also note that, when we include the extra features, the spread is lower for $\theta_{fc} = 0.9$, indicating that these features help in producing more focused forecasts. Although a similar pattern can be observed for the *fishing* pattern as well, a more careful examination reveals that this pattern is more challenging. The precision scores are higher, but the spread is also significantly higher for all values of θ_{fc} . This means that accurately pinpointing when a fishing pattern will be detected is more difficult. This result is also expected, since fishing maneuvers (often occurring in the open sea) exhibit a greater degree of variability, whereas movement patterns while approaching a port are more or less straightforward. The trade-off between precision and spread is thus always present, but its exact nature also depends on both the pattern itself and the included features. We also see that increasing m does not seem to have a significant impact, at least, with the predicates chosen here. Note that this is not always

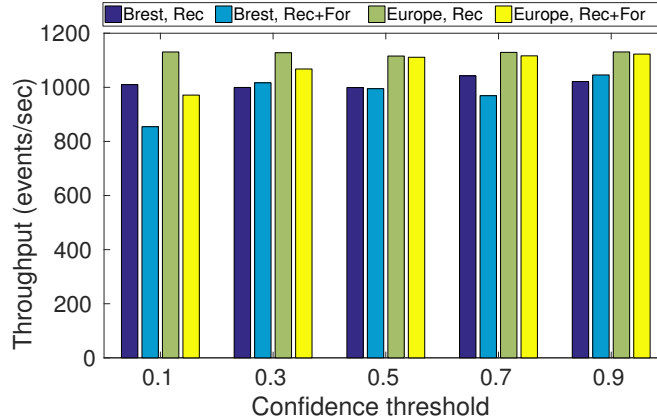


Figure 3: Throughput results. “Rec” denotes recognition only, with forecasting disabled. “Rec+For” denotes that both recognition and forecasting are enabled.

the case, since m can indeed play a significant role in other domains and/or patterns [1]. As a general comment, Figure 2 can help a user determine a satisfactory set of parameter values. For example, for the approaching pattern, a user could choose to set $\theta_{fc} = 50\%$ and $m = 0$, which gives a high enough precision with relatively low spread and avoids the cost of disambiguation that accompanies any higher values of m .

In order to assess Wayeb’s throughput, we run another series of experiments. We tested the *approaching* pattern on both the Brest and the Europe datasets. For the Brest dataset, all of the 222 ports of Brittany were included, each with its own *SFA* and *PMC*. For the Europe dataset, 222 European ports were randomly selected. We tested for throughput both when forecasting is performed and when it is disabled, in which case only recognition is performed. The results are shown in Figure 3. As expected, throughput is lower when forecasting is enabled. However, the overhead is not significant. We also see that throughput is higher for the Europe dataset. It is possible to achieve higher throughput because the event rate of the input stream is also higher. The throughput is higher in this case, despite the fact that it contains almost ten times as many vessels and has a higher incoming event rate, indicating that Wayeb can scale well as the number of monitored objects increases.

5 Discussion

Wayeb is one of the few CEP systems capable of the form of forecasting presented in this paper. The only other system with similar capabilities is presented in [19], using automata and Markov chains as well. The advantage of our method is that it allows for a deeper investigation of the past, if this is needed, by being able to handle higher order Markov processes. By using symbolic automata, our computational model also has nice compositional properties (a feature generally lacking in CEP) and can accommodate any regular expression, without being restricted to sequential patterns. Having predicates on the transitions also allows for an easy incorporation of Boolean expressions and of background knowledge. For example, note that the *IsFishingVessel(x)* predicate in Pattern 2 is evaluated not by using information from the streaming AIS messages (which may not always contain correct information about a vessel’s type) but by using the relevant background knowledge of the list of fishing vessels. A significant

number of forecasting methods comes from the field of temporal pattern mining, e.g., [27, 18, 29]. These are typically unsupervised methods with a focus on predicting what the next *input* event(s) in a stream might be. The same is true for the various sequence prediction methods [5]. However, predicting the next input event(s) is not of the highest priority in CEP. In fact, CEP engines typically employ *selection strategies* that allow for ignoring input events that are not relevant for a given pattern. Most input events might thus be irrelevant for a pattern and making predictions about them is not useful. Instead, it is more important to forecast when a complex event will be detected, as can be done with our method. This is the reason why we cannot compare Wayeb to these methods, since they have a different formulation of the forecasting task. Moreover, sequence prediction methods, as well as predictive CEP methods inspired by sequence prediction approaches [14], are based on the assumption of both a finite alphabet and a language of finite cardinality, basically excluding iteration. Both of these assumptions do not hold in real-world applications and our method does not require them.

Although sequence prediction methods are not suitable for complex event forecasting as they are, they employ techniques that could be useful in our case as well. For example, due to the high cost of increasing the assumed order m of the Markov process, they employ variable-order Markov models [6]. We also intend to explore this direction, in order to avoid the combinatorial explosion on the number of states of a PMC. Another research direction is that of finding ways to compactly represent the past [25] without having to enumerate every possible combination. Finally, the predicates of symbolic automata are unary and are applied only to the last event read from the stream. This is a serious limitation for CEP where patterns are required having constraints between the last event and events seen in the past. We intend to investigate if and how our method can be extended to other automata models that provide this functionality, like extended symbolic automata [9] and quantified event automata [4].

6 Acknowledgments

This work was supported by the EU H2020 datAcron project (grant agreement No 687591).

References

- [1] E. Alevizos, A. Artikis, and G. Paliouras. Event forecasting with pattern markov chains. In *Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems*. ACM, 2017.
- [2] E. Alevizos, A. Skarlatidis, A. Artikis, and G. Paliouras. Probabilistic complex event recognition: a survey. *ACM Computing Surveys (CSUR)*, 50(5):71, 2017.
- [3] T. Anderson and L. Goodman. Statistical Inference about Markov Chains. *The Annals of Mathematical Statistics*, 1957.
- [4] H. Barringer, Y. Falcone, K. Havelund, G. Reger, and D. Rydeheard. Quantified event automata: Towards expressive and efficient runtime monitors. In *International Symposium on Formal Methods*, pages 68–84. Springer, 2012.
- [5] R. Begleiter, R. El-Yaniv, and G. Yona. On prediction using variable order markov models. *Journal of Artificial Intelligence Research*, 22:385–421, 2004.
- [6] P. Bühlmann, A. Wyner, et al. Variable length markov chains. *The Annals of Statistics*, 1999.
- [7] F. Chen and G. Roşu. Parametric trace slicing and monitoring. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2009.
- [8] G. Cugola and A. Margara. Processing flows of information: From data stream to complex event processing. *ACM Computing Surveys (CSUR)*, 44(3):15, 2012.

- [9] L. D’Antoni and M. Veanes. Extended symbolic finite automata and transducers. *Formal Methods in System Design*, 2015.
- [10] L. D’Antoni and M. Veanes. The power of symbolic automata and transducers. In *International Conference on Computer Aided Verification*. Springer, 2017.
- [11] Y. Engel and O. Etzion. Towards proactive event-driven computing. In *Proceedings of the 5th ACM international conference on Distributed event-based system*, pages 125–136. ACM, 2011.
- [12] J. C. Fu and W. Y. W. Lou. *Distribution theory of runs and patterns and its applications: a finite Markov chain imbedding approach*. World Scientific, 2003.
- [13] L. Fülöp, A. Beszédes, G. Tóth, H. Demeter, L. Vidács, and L. Farkas. Predictive Complex Event Processing: A Conceptual Framework for Combining Complex Event Processing and Predictive Analytics. In *Fifth Balkan Conference in Informatics*, 2012.
- [14] S. Gillani, A. Kammoun, K. Singh, J. Subercaze, C. Gravier, J. Fayolle, and F. Laforest. Pi-cep: Predictive complex event processing using range queries over historical pattern space. In *Data Mining Workshops (ICDMW), 2017 IEEE International Conference on*. IEEE, 2017.
- [15] V. Gulisano, Z. Jerzak, P. Smirnov, M. Strohbach, H. Ziekow, and D. Zissis. The DEBS 2018 grand challenge. In *Proceedings of the 12th ACM International Conference on Distributed and Event-based Systems*. ACM, 2018.
- [16] J. Hopcroft, R. Motwani, and J. Ullman. *Introduction to automata theory, languages, and computation*. Pearson/Addison Wesley, 2007.
- [17] A.-L. Jousselme, C. Ray, E. Camossi, M. Hadzagic, C. Claramunt, K. Bryan, E. Reardon, and M. Iteris. Deliverable D5.1: Maritime use case description. Project datAcron.
- [18] S. Laxman, V. Tankasali, and R. W. White. Stream prediction using a generative model based on frequent episodes in event sequences. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2008.
- [19] V. Muthusamy, H. Liu, and H.-A. Jacobsen. Predictive publish/subscribe matching. In *Proceedings of the Fourth ACM International Conference on Distributed Event-Based Systems*. ACM, 2010.
- [20] P. Nicodème, B. Salvy, and P. Flajolet. Motif statistics. *Theoretical Computer Science*, 2002.
- [21] G. Nuel. Pattern Markov Chains: Optimal Markov Chain Embedding through Deterministic Finite Automata. *Journal of Applied Probability*, 2008.
- [22] K. Patroumpas, E. Alevizos, A. Artikis, M. Vodas, N. Pelekis, and Y. Theodoridis. Online event recognition from moving vessel trajectories. *GeoInformatica*, 2017.
- [23] C. Ray, R. Dreo, E. Camossi, and A. Jousselme. Heterogeneous Integrated Dataset for Maritime Intelligence, Surveillance, and Reconnaissance, 10.5281/zenodo.1167595, 2018.
- [24] J. Sakarovitch. *Elements of automata theory*. Cambridge University Press, 2009.
- [25] P. Tino and G. Dorffner. Predicting the future of discrete sequences from fractal representations of the past. *Machine Learning*, 2001.
- [26] M. Veanes, P. De Halleux, and N. Tillmann. Rex: Symbolic regular expression explorer. In *2010 Third International Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 2010.
- [27] R. Vilalta and S. Ma. Predicting rare events in temporal domains. In *Proceedings of the 2002 IEEE International Conference on Data Mining*, page 474. IEEE Computer Society, 2002.
- [28] H. Zhang, Y. Diao, and N. Immerman. On complexity and optimization of expensive queries in complex event processing. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 2014.
- [29] C. Zhou, B. Cule, and B. Goethals. A pattern based predictor for event streams. *Expert Systems with Applications*, 2015.