

Using Self-Organizing Maps to Learn Hybrid Timed Automata in Absence of Discrete Events

Alexander von Birgelen
Institute Industrial IT
Ostwestfalen-Lippe of Applied Sciences
Lemgo, Germany
Email: alexander.birgelen@hs-owl.de

Oliver Niggemann,
Institute Industrial IT
Ostwestfalen-Lippe of Applied Sciences
Lemgo, Germany
Email: oliver.niggemann@hs-owl.de

Abstract—Modern industrial plants become more complex and consequently monitoring them often exceeds the capabilities of human operators. Model-based diagnosis is a commonly used approach to identify anomalies and root causes within a system through the use of models, which are often times manually created by experts. However, manual modelling takes a lot of effort and is not suitable for today's fast-changing systems. Today, the large amount of sensor data provided by modern plants enables data-driven solutions and models can be learned from data, significantly reducing the manual modelling efforts. These data-driven solutions enable tasks such as condition monitoring: anomalies can be detected automatically, giving operators the chance to restore the plant to a working state before production losses occur. The choice of the model depends on a couple of factors, one of which is the type of the available signals. Hybrid timed automata are one type of model which separate the systems behaviour into different modes, e.g. 'valve open' or 'motor is running' through discrete events which are for example created from binary signals of the plant or through real-valued signal thresholds, defined by experts. The real-valued signals are then separated into the corresponding modes to improve the anomaly detection process in comparison to unseparated data. The anomaly detection for hybrid timed automata combines the detection of timing errors and sequence errors in the mode changes and the detection of anomalies in the real-valued signals.

However, binary signals or expert knowledge to generate the much needed discrete events are not always available from the plant and automata can not be learned. The unsupervised, non-parametric approach presented and evaluated in this paper uses self-organizing maps and watershed transformations to allow the use of hybrid timed automata on data where learning of automata was not possible before.

I. INTRODUCTION

Increasing product variety, product complexity and pressure for efficiency in a distributed and globalized production chain cause production systems to evolve rapidly: they become modular, can be parameterized and contain a growing set of sensors [1].

In order to enable European SMEs to face these challenges and to utilize new technical possibilities, the Horizon2020 project IMPROVE is aimed at developing user support functions in terms of self-diagnosis (i.e. condition monitoring, predictive-maintenance) and self-optimization (e.g. energy optimization, output optimization).

Models can be constructed manually by an expert, but this is difficult, costly and time consuming in today's complex

and evolving production plant [2]. Instead of only relying on human expertise and additional engineering steps formalizing the necessary knowledge, the tasks stated above will be taken on in a data-driven way [3] where models are learned automatically from the data. For anomaly detection, the live data from the plant is compared to the predictions of the learned model and deviations from the normal behaviour are classified as anomalous.

The type of model depends on a variety of conditions such as available signals, the task of the model and the overall nature of the system. Modern cyber-physical production systems (CPPS) usually are hybrid systems, meaning they comprise both binary and real-valued signals. In general, these dynamic systems are state based, for example the system's state is defined by its current and previous binary control signals, and the actions taken out are time dependent [4].

One approach to perform anomaly detection in hybrid systems is to use hybrid timed automata [5] as a normal behaviour model which utilize discrete events to learn the system's normal behaviour. These events often cause so called mode or state changes in industrial plants, e.g. *conveyor is running* or *valve open*. Hybrid timed automata are well suited to learn the normal behaviour in terms of the modes/states, transitions and corresponding timings from data in an unsupervised manner. The real-valued signals are processed within the states using other types of models. Such models can for example be regression models or models which also reduce the dimensionality of the real-valued data. An example for this is the nearest-neighbor principal component analysis (NNPCA) [6] which is used for the experiments in section IV. The NNPCA was chosen because it is used frequently in a variety of research projects in our institute.

Changes in the binary control and sensor signal values of the system can be utilized as discrete events directly. Another way to get discrete events is to set thresholds on continuous signals but this requires additional expert knowledge. Unfortunately, expert knowledge is almost never available and binary sensor signals are also not always available or not meaningful. This occurs for example when data is recorded using internal trace functionalities from drive controllers, which offer very high sampling rates but lack knowledge of variables from the programmable logic controller (PLC).

In this paper we present an unsupervised, non parametric approach to learn hybrid timed automata in absence of discrete events. The approach uses self-organizing maps (SOM) and watershed transformations to extract modes and generate discrete events in an unsupervised manner using only real-valued signals. The generated events can then be used to capture the normal behaviour of the system by learning hybrid timed automata on data where they were not applicable before. The learned hybrid automaton is then used for anomaly detection in the real-valued signal values and in the time domain by analysing the transitions between the automaton's states.

The contents of this paper are structured as follows: Section II introduces the existing modelling formalisms which are then combined in section III to generate discrete events from real-valued signals in an unsupervised manner. Experimental results from three different data sets, one artificial and two real world ones, are given in section IV. Finally, the paper is concluded in section V.

II. METHODOLOGIES

A. Hybrid Timed Automata

Hybrid timed automata have proven to be a great tool to learn the normal behaviour of a system and detect deviations from it. Discrete events are required to learn an automaton. These events often cause mode changes in the system and the timing of these events is an important indicator for the health of the system. Hybrid timed automata are used to separate these modes, learn the transitions and timing between them and model the behaviour of the real-valued signals for each of the modes or states in the automaton.

An easy approach to obtain discrete events can be directly extracted from changes in the binary control and sensor signals of the system. It is also possible to obtain discrete events through thresholds for continuous signals such as *temperature* $< 19^\circ\text{C}$ [7]. However, setting the thresholds and combinations of conditions for the continuous signals requires expert knowledge which is usually not available for real world automation systems. For unsupervised learning of these automata only binary control signals are used to obtain the discrete events. Algorithms such as the online timed automaton learning algorithm (OTALA) [8] and its hybrid extension can work in an online, unsupervised manner, and do not require additional expert knowledge.

A hybrid automaton generated by the aforementioned algorithm can be defined as described in Definition 1.

Definition 1: a hybrid timed, probabilistic automaton is a tuple $A = (S, s_0, \Sigma, T, \delta, P, \theta)$, where

- S is a finite set of states where $s \in S$.
- s_0 is the initial state which can be given by the systems state at the start of the training.
- Σ is the set of discrete events. Events $a \in \Sigma$ is linked to the transitions of the automaton.
- T is the set of transitions with $t \in T$ and $t = (s, a, s')$, $s, s' \in S$ are source and destination state, $a \in \Sigma$ is the trigger event of the transition.
- The timing constraint $\delta : T \rightarrow I$ assigns a time interval to a transition $t \in T$, where I is a set of time intervals. The time here usually refers to the elapsed time since the last event occurred.
- P is a set of probabilities: for each transition $t \in T$ probability $p \in P$ is calculated.
- $\theta_{s \in S}$ describes a model for each state $s \in S$ which captures the behaviour of the real valued signals. Real valued signals are not captured by the discrete part of the automaton. These state models $\theta_{s \in S}$ are learned for each state of the automaton using other models such as linear regression, decision trees and others, such as the nearest neighbour principal component analysis (NNPCA) used in section IV of this paper.

The learned automaton can then be used to detect a variety of different classes of anomalies. This can for example be done using the anomaly detection algorithm (ANODA) [9] which can detect the following types of anomalies:

- **Unknown event / Wrong event sequence:** an event occurred which was not observed in the current state.
- **Timing error:** a transition occurred outside of the learned time bounds.
- **State remaining error:** when more time passed than for the latest event and the state is not a final state, then we have a state remaining error.
- **Probability error:** the probabilities of transitions for the new data are calculated and compared to the previously learned probabilities and an error is generated when deviations are too large.
- **Continuous error:** for each state and additional anomaly detection for the continuous signals can be performed using the internal state models.

B. Self-Organizing Map

The self-organizing map (SOM), also referred to as self-organizing feature map or kohonen network, is a neural network that can be associated with vector quantization, visualization and clustering but it can be used as an approach for non-linear, implicit dimensionality reduction [10]. The reduction is performed in a qualitative, implicit way. A SOM consists of a collection of neurons which are connected in a topological arrangement which is usually a two dimensional rectangular or hexagonal grid. The input data is mapped to the neurons forming the SOM. Each neuron is essentially a weight vector of original dimensionality.

Definition 2: the self-organizing map $SOM = (M, G, d)$ forms a topological mapping of an input space $O \subset \mathbb{R}^m$, $m \in \mathbb{N}$ and consist of

- a set of neurons M .
- each neuron $n \in M$ has a weight vector $\mathbf{w}_n \in \mathbb{R}^m, m \in \mathbb{N}$.
- G is usually a two-dimensional rectangular or hexagonal lattice in which the neurons $n \in M$ are arranged. Toroidal versions of these topologies are also common.
- $d(\mathbf{x}, \mathbf{y})$ is the distance measure to calculate the distance between two vectors \mathbf{x} and \mathbf{y} which can for example be




Fig. 1: Different u-matrix visualizations of a 120x120 SOM.

weight vectors and/or vectors in the input space. Usually, the euclidean distance is used but other measures, such as the mahalanobis distance, can be used.

- an input sample $\mathbf{o}_i \in \mathbb{R}^m, i \in \mathbb{N}$ is mapped to the SOM through its best matching unit (BMU). The BMU is given by $bm_u(\mathbf{o}_i) = \operatorname{argmin}_{k \in M} d(\mathbf{o}_i, \mathbf{w}_k)$

One way to learn a SOM from data is a random batch training approach: The initial values of the neuron's weight vectors for the training can be randomly initialized or sampled from the training data. All samples from the training data are presented to the algorithm within one epoch. A best matching unit (BMU) is calculated for each input sample from the training data by finding the neuron which has the smallest distance to the sample. The BMU and all of its neighbouring neurons, assigned through the topology and neighbourhood radius, are shifted towards the input sample. Both the neighbourhood radius and strength of the shift decrease over time. The training stops after a chosen amount of epochs.

Each neuron of the SOM represents a part of the training data. Areas in the training input space with few examples are represented by few neurons of the SOM while dense areas in the input space are represented by a larger number of neurons.

The unified distance matrix (u-matrix) [11] allows a three-dimensional visual identification of clusters contained in the self-organizing map. It calculates the average distance to neighbouring neurons according to the SOM's topology and visualizes clusters contained in the, usually high dimensional, training data. The X and Y coordinates of the neurons represent the first two dimensions. The third dimension is given by the average distance to neighbouring neurons as in definition 3.

It can be visualized directly in 3D or in 2D using a color gradient as shown in Figure 1.

Definition 3: for each neuron $n \in M$ and its associated weight vector \mathbf{w}_n , the u-matrix height is given by $U(n) = \sum_{k \in NN(n, G)} d(\mathbf{w}_n, \mathbf{w}_k)$, where $NN(n, G)$ is the set of neighbouring neurons of n defined by grid G and $d(x, y)$ is the distance used in the SOM algorithm.

The u-matrix representation illustrates why SOM's were chosen: SOM's tend to keep neurons with similar signal weights closely together, which results in a topographic landscape with valleys, where weights of neighbours are similar, and ridges, where weights of neighbours are not similar. Valleys represent regions where the contained neurons weight




Fig. 2: Watershed transformations of a u-matrix.

vectors are very similar. These valleys are separated by ridges which mark transitions between the different feature spaces.

C. Watershed Transformation

Clustering algorithms from the image processing domain, such as the watershed transformation [12], can be used on the u-matrix representation of a SOM to identify the clusters in a mathematical way.

This works analogous to rain falling on top of the u-matrix. The water runs from higher regions to the lower regions and consequently flooding the basins. When the water level gets high enough so two basins merge, a ridge forms which separates them.

The watershed transformation dissects the u-matrix into different clusters, separated by the so-called watershed lines. Watershed lines separate the different basins and do not belong to any of the clusters. The basins can be interpreted as stationary process phases while the watershed lines represent transient process phases [13].

The implementation used here is the Vincent-Soille watershed algorithm which performs the watershed transformation in a non-recursive manner [14]. The sensitivity of the algorithm can be adjusted by setting a number of levels which in turn influences the number of final clusters found. Figure 4 shows examples using different levels on the same u-matrix.

In the end, we receive a mapping for each neuron of the map to its corresponding cluster:

Definition 4: the watershed transformation maps each neuron $n \in M$ to a cluster c , with C being a set of clusters and $c = [0, |C| - 1] \in \mathbb{N}$.

III. LEARNING HYBRID TIMED AUTOMATA WITHOUT DISCRETE EVENTS

In order to learn hybrid timed automata, without binary control signals and without expert knowledge about thresholds for continuous signals, it is necessary to derive the discrete events using an alternative way. For complex systems, the events can also be related to combinations of different continuous signal values. Here, SOMs are used as a preprocessing step to extract the different modes of the system in an unsupervised manner.

This section describes the generation of discrete events using SOMs. Figure 3 shows the general steps of the presented approach. A SOM is learned from the input data and the

transitions between emerging clusters are used to generate discrete events. This is essentially a preprocessing procedure which allows learning hybrid timed automata in case the input data does not contain discrete events. After events are generated an automaton can be learned and then used to detect anomalies within the real-valued signal values and also in the timing and probabilities of the events.

First, real-valued signals are recorded during normal operation of the plant to generate dataset O which consists of many observations $o \in O$ and represents the input space of the model. Then, the SOM-Discretization algorithm (Algorithm 1) generates discrete events for the data to allow learning of hybrid timed automata:

A SOM is trained on the recorded normal behaviour data (step 3). The size of the SOM is automatically calculated according to [15], where the number of neurons $|M| \approx 5\sqrt{N}$ with N being the number of observations. The ratio of the side lengths is the ratio of the two largest eigenvalues of the data's covariance matrix. A normalization should generally be done before the SOM training but this depends on the input space and is therefore optional (step 2).

The SOM's u-matrix is calculated (step 4) and then clustered using the watershed transformation (step 5). The watershed levels are adjusted so the final cluster count is close to the shorter side length of the SOM. Each basin represents a stationary process phase and gets a unique number $i = [0, |C| - 1]$, $i \in \mathbb{N}$ for identification. The watershed lines are transient process phase all receive a negative identification number to distinguish them from the stationary process phases. Observations mapped to transient process phases, the borders left after the watershed transformation, are assigned to the same cluster as the previous observation and therefore receive the same event vector as the previous observation (step 12). The event vector for cluster $c \in C$ in step 10 is created such that:

$$\mathbf{v}(c) = \left(v_1, v_2, \dots, v_i, \dots, v_{|C|-1} \mid v_i = \begin{cases} 1 & \text{if } i=c \\ 0 & \text{else} \end{cases} \right)$$

The event vector is then concatenated with the original signal vector in step 14. The data now contain binary signals which are then interpreted as discrete events by hybrid automaton learning algorithms such as the hybrid OTALA used in this paper.

The anomaly detection can be performed either offline or online. Algorithm 2 shows the procedure for an online anomaly detection, where each observation o is tested as soon as it is received. When normalization was used for the training each new observation must also be normalized based on the normalization parameters calculated before (step 2). The best matching unit is calculated and linked to its corresponding cluster (steps 3, 4). When the observation falls onto a transient process phase it is assumed to be in the same cluster as the previous observation (step 8). When mapped to a stationary process phase, the discrete event signal vector is generated and appended to the observation in steps 6 and 10. The observation

Algorithm 1 Generation of discrete events

```

1: procedure SOM-DISCRETIZATION( $O$ )
2:    $O \leftarrow \text{NORMALISE}(O)$ 
3:    $SOM \leftarrow \text{TRAIN SOM}(O)$ 
4:    $U \leftarrow \text{U-MATRIX}(SOM)$ 
5:    $C \leftarrow \text{WATERSHED}(U)$ 
6:   for all  $o \in O$  do
7:      $bm u_o \leftarrow \text{BMU}(SOM, o)$ 
8:      $c_o \leftarrow C(bm u_o)$ 
9:     if  $c_o \geq 0$  then
10:       $v_o \leftarrow \mathbf{v}(c_o)$ 
11:     else
12:       $v_o \leftarrow v_{o-1}$ 
13:     end if
14:      $p \leftarrow \{v_o, o\}$ 
15:   end for
16: end procedure

```

is then given to the ANODA algorithm to perform the anomaly detection using the learned automaton.

Algorithm 2 Preprocessing for online anomaly detection

```

1: procedure SOM-DISCRETIZATION-ONLINEAD( $o$ )
2:    $o \leftarrow \text{NORMALISE}(o)$ 
3:    $bm u_o \leftarrow \text{BMU}(SOM, o)$ 
4:    $c_o \leftarrow C(bm u_o)$ 
5:   if  $c_o \geq 0$  then
6:      $v_o \leftarrow \mathbf{v}(c_o)$ 
7:   else
8:      $v_o \leftarrow v_{o-1}$ 
9:   end if
10:   $o \leftarrow \{v_o, o\}$ 
11:  return  $o$ 
12: end procedure

```

IV. EXPERIMENTS

This section presents experimental results of the presented approach using one artificial and two real world data sets.

A. Artificial test data

We created an artificial dataset through a simple PLC application. The PLC moves a virtual, linear SoftMotion [16] axis back and force between two target positions. The drive uses trapezoid ramps with limited speed and accelerations. The drive related data from the software PLC are acquired through OPC-UA subscriptions in a 100ms publishing interval.

The data contains three real-valued signals: target position, actual velocity and actual speed. The training set contains 15 cycles, while the evaluation set contains five normal and five anomalous cycles. Maximum acceleration and deceleration are decreased during the anomalous cycles. An excerpt of the data is shown in Figure 4. The anomalies are labelled, to later evaluate the score for the anomaly detection. All of the

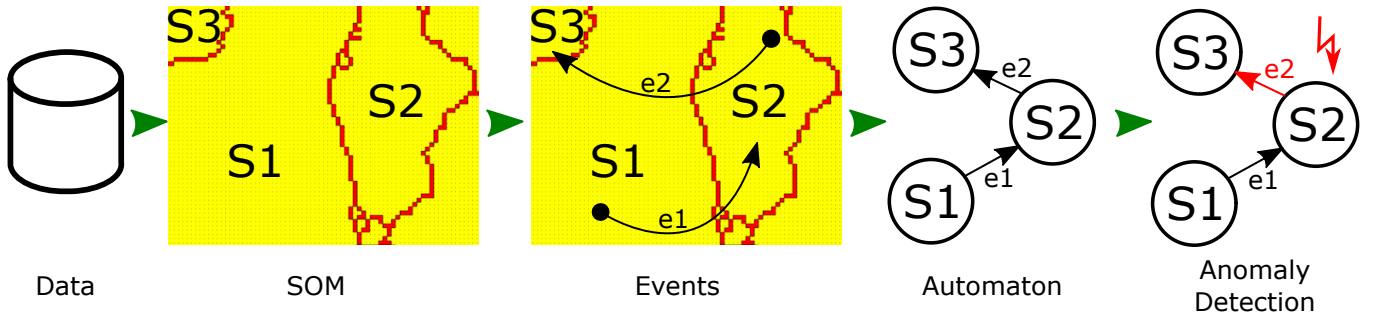


Fig. 3: General steps of the presented algorithm.




Fig. 4: One normal (827-991) and one anomalous (992-1251) cycle, scaled to range [0,1].

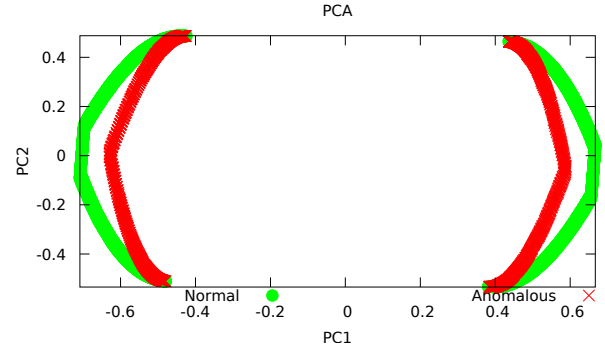


Fig. 5: Plot of the full dataset in the two-dimensional PCA space. Only normal behaviour is known to the learned model. Anomalies plotted to show the separation between normal and anomalous.

machine learning methods here work unsupervised and have no knowledge about the labels.

The nearest-neighbor PCA (NNPCA) [6] was first used on the full training data and then second as state-model for the hybrid automaton.

The data was reduced to two dimensions and then used for the anomaly detection. The training observations provide the reference to calculate the distance for each evaluation observation. When the distance exceeds a certain threshold, the observation is considered anomalous. The threshold is calculated using a mexican-hat wavelet, so the distance is given as an error probability.

Setting the threshold is not trivial and highly depends on the input data. A 100% threshold is good against false positives but also might be not sensitive enough to find the true positives. Lowering the threshold usually increases the true positive rate but at the risk of more false positives.

Figure 5 shows a plot of two-dimensional NNPCA learned from the normal behaviour used as training data. The anomalies from the evaluation data are unknown to the model and are shown here to visualize the separation between normal and anomalous behaviour. A threshold of 25% was selected, so every observations which gets a probability greater or equal to 25% of being anomalous is marked anomalous. Even with this low threshold, not a single anomaly is found by the




Fig. 6: Clusters on the SOM. Grey areas mark transient process phases.

NNPCA on the full dataset as the separation between normal and anomalous behaviour is not large enough to be found with the given threshold. Furthermore, the NNPCA model does not include time in an explicit way and detection of anomalies in the time domain difficult or often not possible.

Now, algorithm 1 is used to generate discrete events so the data can be separated by a hybrid automaton. The automatic selection of parameters results in a SOM with 24x10 neurons and 11 clusters as shown in Figure 6.

With this, a hybrid automaton is trained using the hybrid OTALA algorithm and the NNPCA as a state model with the exact same parameters as before. The resulting automaton is




Fig. 7: The learned automaton with 12 states. State 1 is an initial state. Each cluster is mapped to one state.




Fig. 8: Two state models from the automaton in Figure 7 showing the normal and mapped anomalous behaviour.

shown in Figure 7. Behind each state stands the state model from which two are shown in Figures 8a and 8b showing the separation of the normal and anomalous behaviour.

The presented approach for the event generation in preparation for automaton learning calculates the necessary parameters from the training data as described in section III. The settings for the NNPCA within the states were kept identical and the hybrid automaton achieves F1 measure of 92.68%, compared to the 0% F1 achieved by the NNPCA on the full dataset.

Table I shows some more details about the scores. The automaton offers another advantage: deviations in the timing and sequence of the system's behaviour can now be detected, which is not possible with the NNPCA alone.

B. High Rack Storage System

The High Rack Storage System or HRSS is a demonstrator from the SmartFactoryOWL which transports wares between its different shelves. The data from the system's drives includes position, power and voltage signals to a total of 17 real-valued signals, after removal of signals with zero-variance. The training data contains 107 cycles of the same transport operation. Evaluation data contains 113 cycles and was generated by modifying the training cycles in different ways such as increasing or decreasing one or multiple signals by different amounts (e.g. 20%), shortening of cycles and inserting




Fig. 9: The learned automaton for the storage system.

constant sequences. This artificially generated anomalies are labelled and are used to calculate the scores of the anomaly detection.

The NNPCA reduces the data to 15 dimensions, keeping a variance coverage of 99.93%. A 60% threshold for the anomaly detection, successfully identifies 38 of the 7365 labelled anomalies resulting in an atrocious F1 score of 1%. The event generation creates a 34x24 SOM resulting in 24 clusters which are then captured by the automaton (Figure 9).

The automaton with the 15-dimensional NNPCA inside the states identifies 1516 true positives and reaches an F1 measure of 30.76% which is a drastic increase to the full-dataset NNPCA. It is to be noted that the anomalies were generated without respect to the overall dataset, so a 20% increase of a signal might not be a significantly different from the overall training data and therefore the absolute scores on this dataset might not be representative.

C. Film-Spool Unwinder

The third dataset presented in this paper was recorded from the cutting unit of a Vega shrink-wrap packer by OCME [17]. The machine groups loose bottles or cans into set package sizes, wraps them in plastic film and then heat-shrinks the plastic film to combine them into a package.

The drive controllers within the film cutting unit recorded chunks of data which each contain 2048 observations using their built-in scope functionality at a 4ms resolution. Here, we only consider two signals from the film unwinder drive: actual speed and lag error. 40 chunks were used for training and another 32 chunks were used as evaluation data. The film spool depleted during the last cycles, so every observation in the last 12 cycles was labelled anomalous. A dimensionality reduction is not necessary for two dimensions, but the NNPCA was still used to keep the model consistent for the experiments presented here. Again, the NNPCA failed to detect any anoma-




Fig. 10: Unwinder Automaton

lies on the full dataset. The automatically calculated size of the SOM is 64x22 neurons resulting in the automaton shown in Figure 10. The NNPCA on the full data and a selection of the 24 state models is shown in Figure 11.

The absolute scores for this dataset have to be taken with a grain of salt because the anomalies are not labelled observation-perfect but a generous range of observations was marked as anomalous. The F1 score for the anomaly detection with the automaton reaches 40.38%, resulting in an increase of over 21% (Table I) compared to the NNPCA on the full dataset.

V. CONCLUSION

This paper presented an unsupervised, non-parametric approach which allows application of hybrid timed automata on data which do not allow their use due to a lack of knowledge about discrete events which are needed to learn and use hybrid timed automata. Hybrid automata are used for anomaly detection in the real-valued signal values as well as in the time domain by analysing the transitions between the states of the automaton.

The presented approach separates the data into different stationary and transient process phases using self-organizing maps and watershed transformations. Also, all necessary parameters for the SOM and watershed transformation are automatically estimated from data. The discrete events generated from the transitions between the extracted process phases are used to learn a hybrid timed automaton which in turn is used for anomaly detection. The presented algorithms work offline during the learning phase and can later be used online with live data from the plant. Further, the presented algorithms are not linked to a specific automaton learning algorithm and can be used as a preprocessing step prior to automaton learning.

The anomaly detection for real-valued signals through a model within the states of the learned hybrid automaton increases the performance of the anomaly detection in comparison to the same model working with the same parameters on the full dataset. The experiments in section IV show results from three different datasets, two of which come from real-world machines. For all tested datasets, the separation of modes through the automaton improves the scores for the anomaly detection of the real-valued signals, without any additional tuning of parameters for the anomaly detection.

ACKNOWLEDGEMENT

This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No. 678867.

REFERENCES

- [1] *Factories of the Future: MultiAnnual Roadmap for the contractual PPP under HORIZON 2020*. Luxembourg: European Union, 2013.
- [2] B. Vogel-Heuser, C. Diedrich, A. Fay, S. Jeschke, M. Kowalewski S. and Wollschlaeger, and P. Goehner, “Challenges for software engineering in automation,” *Journal of Software Engineering and Applications*, vol. 7, no. 5, 2014.
- [3] O. Niggemann, A. Maier, A. Vodencarevic, and B. Jantscher, “Fighting the modeling bottleneck - learning models for production plants,” *Workshop ”Modell-basierte Entwicklung Eingebetteter Systeme” (MBEES)*, Jul. 2011.
- [4] O. Niggemann, A. Maier, R. Just, and M. Jäger, “Anomaly detection in production plants using timed automata : Automated learning of models from observations,” in *Proceedings of the 8th International Conference on Informatics in Control, Automation and Robotics*, 2013, pp. 363 –369.
- [5] A. Maier and O. Niggemann, “On the learning of timing behavior for anomaly detection in cyber-physical production systems,” in *International Workshop on the Principles of Diagnosis (DX)*, Paris, France, 2015.
- [6] J. Eickmeyer, T. Krueger, A. Frischkorn, T. Hoppe, P. Li, F. Pethig, S. Schriegel, and O. Niggemann, “Intelligente zustandsberwachung von windenergieanlagen als cloud-service,” in *Automation 2015*, Baden-Baden, 2015.
- [7] T. A. Henzinger, “The theory of hybrid automata,” in *Proceedings 11th Annual IEEE Symposium on Logic in Computer Science*, 1996, pp. 278–292.
- [8] A. Maier, “Online passive learning of timed automata for cyber-physical production systems,” in *2014 12th IEEE International Conference on Industrial Informatics (INDIN)*, 2014, pp. 60–66.
- [9] —, “Identification of timed behavior models for diagnosis in production systems,” PhD thesis, Paderborn, Univ., 2015.




Fig. 11: PCA of the unwinder data (11a) and examples from the automaton states (11b, 11c, 11d).

TABLE I: Scores for different datasets and methods. Explanations can be found in the corresponding sections.

Dataset	Method	TP	TN	FP	FN	ACC	F1
IV-A: Artificial test data	NNPCA	0	991	0	1692	36.94%	0%
IV-A: Artificial test data	Automaton, NNPCA only	1443	983	8	249	90.42%	91.82%
IV-A: Artificial test data	Automaton, all errors	1481	968	23	211	90.42%	92.68%
IV-B: High Rack Storage System	NNPCA	38	17990	0	7327	71.10%	1.03%
IV-B: High Rack Storage System	Automaton, NNPCA only	1430	17055	935	5935	72.90%	29.39%
IV-B: High Rack Storage System	Automaton, all errors	1516	17013	977	5849	73.08%	30.76%
IV-C: Film-Spool Unwinder	NNPCA	2528	40926	34	22048	66.31%	18.63%
IV-C: Film-Spool Unwinder	Automaton, NNPCA only	6688	37637	3323	17888	67.63%	38.67%
IV-C: Film-Spool Unwinder	Automaton, all errors	7082	37538	3422	17494	68.08%	40.38%

- [10] H. Yin, "The self-organizing maps: Background, theories, extensions and applications," in *Computational Intelligence: A Compendium*, J. Fulcher and L. C. Jain, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 715–762, ISBN: 978-3-540-78293-3.
- [11] A. Ultsch and H. P. Siemon, "Kohonen's self-organizing feature maps for exploratory data analysis," in *Proceedings of the International Neural Network Conference (INNC'90)*, 1990.
- [12] F. Meyer, "Topographic distance and watershed lines," *Signal Processing*, vol. 38, no. 1, pp. 113–125, 1994.
- [13] C. Frey, "Monitoring of complex industrial processes based on self-organizing maps and watershed transformations," in *Industrial Technology (ICIT), 2012 IEEE International Conference on*, 2012.
- [14] L. Vincent and P. Soille, "Watersheds in digital spaces: An efficient algorithm based on immersion simulations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 6, pp. 583–598, 1991.
- [15] J. Tian, M. H. Azarian, and M. Pecht, "Anomaly detection using self-organizing maps-based k-nearest neighbor algorithm," *Second European Conference of the Prognostics and Health Management Society 2014*, 2014.
- [16] 3S-Smart Software Solutions GmbH. (2017). Codesys softmotion: Integrierte bewegungssteuerung in einem iec 61131-3 programmiersystem, [Online]. Available: <https://de.codesys.com/produkte/codesys-motion-cnc/softmotion.html> (visited on 03/28/2017).
- [17] OCME. (2017). Shrink-wrap packers vega, [Online]. Available: <http://www.ocme.com/en/our-solutions/secondary-packaging/vega> (visited on 03/27/2017).