

Continual learning of a time series model using a mixture of HMMs with application to the IoT fuel sensor verification

Przemysław Głomb^{*✉}, Michał Cholewa^{*✉}, Paweł Foszner^{†✉} and Jakub Bularz[‡]

^{*}Institute of Theoretical and Applied Informatics, Polish Academy of Sciences,
 Bałtycka 5, 44-100 Gliwice, Poland, {pglomb,mcholewa}@iitis.pl

[†]Department of Computer Graphics, Vision and Digital Systems

Faculty of Automatic Control, Electronics and Computer Science, Silesian University of Technology

Akademicka 2A, 44-100 Gliwice, Poland, pawel.foszner@polsl.pl

[‡]AIUT Sp. z o.o. ul. Wyczółkowskiego 113, 44-109 Gliwice, Poland

Abstract—This paper presents an application of a mixture of Hidden Markov Models (HMMs) as a tool for verification of IoT fuel sensors. The IoT fuel sensors report the level of fuel in tanks of a petrol station, and are a key component for monitoring system reliability (billing), safety (fuel/oil leak detection) and security (theft prevention). We propose an algorithm for learning a mixture of HMMs based on a continual learning principle, i.e. it adapts the model while monitoring a sensor over time, signalling unexpected or anomalous sensor reports. We have tested the proposed approach on a real-life data of 15 fuel tanks being monitored with the FuelPrime system, where it has shown a very good performance (average area under ROC curve of 0.94) of detecting anomalies in the sensor data. Additionally we show that the proposed method can be used for trend monitoring and present qualitative analysis of the short and long term learning performance. The proposed method has promising performance score, the resulting model has a high degree of explainability, limited memory and computation requirements and can be easily generalized to other domains of sensor verification.

I. INTRODUCTION

A Key element of a recent change in the industry – dubbed ‘Industry 4.0’ or ‘The Fourth Industrial Revolution’ – is the proliferation of ‘smart’, connected Internet of Things (IoT) sensors, which have an ever increasing role in process monitoring, and as such require verification to achieve reliable, safe and secure systems [1]. In case of fuel tank sensors, which measure the state of large tanks e.g. within a petrol station, lack of sensor verification leads to problems being not detected, which results in reliability issues (billing errors) but can also have serious consequences for safety (not detecting a leak and subsequent environmental contamination) and security (facilitate theft of the fuel).

A sensor verification, or validation, is an internal, external or combined process to detect sensor faults and prevent control failures [2]. In case of IoT sensors, this consists of

This work was partially funded by National Center for Research and Development no POIR.01.01.01-00-0376/17-00 ‘A system for gathering and analysis of streaming data for fuel stations, for optimization of distribution costs and fuel sales as well as on-line monitoring for leakages-related issues’ (FuelPrime).

tasks like preparing models for denoising and missing data imputation, anomaly outlier detection, accuracy and semantic analysis [1]. The nature of the analysis can be complex, e.g. for subsequence outliers (set of consecutive data points that jointly behave unusually) [3]. On the other hand, sequential accumulation of data provides opportunity of continual learning of the properties of sensor behaviour over time; the challenge is to maintain learning ability without forgetting previously learned patterns [4]. Sensor verification method thus should be effective at its tasks [3], but at the same time provide adequate, explainable diagnostic data for the operator or maintenance engineer [2], support data processing within the technical constraints of the IoT sensor suite [1] while correctly dealing with new incoming data [4]. Recently, there has been an emphasis on proposing a approaches both explainable and comprehensive, which are able to deal with real world data, e.g. mobile networks scenarios [5] or oil well monitoring [6].

A large number of approaches have been proposed for application-oriented modelling of time series [7] and quality control [8], [9], including statistical tests, decomposition methods, autoregressive models, neural networks, and probabilistic models. Among them, the Hidden Markov Models (HMMs) [10], [11], [12] have proven to be versatile and effective across many fields, including fault diagnosis [13], [14]. HMM models are attractive, as they have three attractive properties: effective among many application fields; popular, thus well-studied; explainable, as their decisions can be easily traced to parameters of the underlying model. Original single-model HMM formulation have been extended with mixture or ensemble HMM models [15], [16], [17], [18], to further improve their modelling capability. However, while in continual learning setting incrementally learned HMMs can be almost as good as batch learned models [19], corruption of previously learned patterns is one of the main issue [20].

This paper proposes a mathematical model based on a mixture of Hidden Markov Models (HMMs), to be used as a tool for verification of IoT fuel sensors, together with

experiments documenting its performance. Our proposition leads to essentially nonparametric, lightweight (in terms of required computational resources, especially memory), continually learning modelling approach that is able to provide a verification of a sensor data series through the detection of structural changes, outliers and anomalies. Additionally, we present a case study, or experience report, of running the proposed approach through real-life historic data of 15 fuel tanks, with verification of detected patterns; the proposed method achieves a high average area under ROC curve of 0.94. While our study is, for the sake of clarity of presentation, limited to the case of IoT fuel sensor verification, the method can be easily generalized to other domains of sensor verification.

Our approach falls within the task-agnostic category described in [21], as we assume unknown both task boundaries – in our case changes in fuel tank usage characteristics over time – and task labels – in our cases the classes of anomalies to be detected. We note that this is the most general, and hence most desired in a practical application setting. Our learning setting from the point of view of anomaly detection is unsupervised, as all data is used without explicit consideration of the labels [22]. According to the taxonomy given in [23], our approach falls into the *Task-Free Continual Learning* (TFCL), with disjoint data label spaces and no task identities provided.

II. METHOD

A. Hidden Markov Models

A Hidden Markov Model is a model of a system that at any time is in one of n distinct states. At discrete time intervals, state switching occurs in time independent, first order Markovian dynamics (i.e. depends only on the current state). HMM states are not directly observable, however each state has an associated set of parameters describing the emission probability of observable symbols. For the fuel tank monitoring, the observed sequence is the volume of the fuel or its delta, while states correspond to the current situation (idle, refill from a tanker, distribution of fuel to clients, etc.).

A HMM λ of n states is described by initial state probability vector $\Pi = [\pi_i]_{n \times 1}$, state transition probability $A = [a_{ij}]_{n \times n}$, emission probability – typically Gaussian, with mean and standard deviation μ_i, σ_i defined for each state. Three main algorithms – Forward, Viterbi, Baum-Welch – provide tools for finding a probability of a sequence for a given model, state sequence for given data, and learning a model from data [10]. Other parameter identification schemes are possible [24]. In this work, all HMMs are ergodic, i.e. transition from a state to all other states is possible at the start of the training.

B. Mixture of HMMs

Mixed (hidden) Markov models were originally introduced for latent class models [15] in social sciences, and further adapted e.g. for accelerometer measurements [16] or data imputation [17]. Those models introduce hierarchical structure, with class membership dictating Markov model parameters. A

different approach has been proposed in [18], where an ensemble of HMMs is generated over time through incremental Boolean combination in the receiver operating characteristic (ROC) space.

In contrast to the above-mentioned propositions, we use a different approach. Our physical sensor model does not require latent class modelling, and absence of labels prevents from using ROC-based verification within the model operation. Our objective with using mixtures is to capture rare historical data patterns, and thus prevent them from being subject to catastrophic forgetting [4]. Our proposition is to model a IoT sensor time sequence with a set of m HMM models $\mathcal{H} = \{\lambda_1, \dots, \lambda_m\}$. We assume that the sensor data $\mathbf{x} \in \mathbb{R}^d$ is processed in windows or batches (e.g. a day's worth of data, d can vary between windows). We propose the following algorithm for mixed HMM sensor verification:

- 1) Initialize $\mathcal{H} = \emptyset$.
- 2) Read next window of sensor data \mathbf{x} . Use the Baum-Welch algorithm to identify parameters of a model for this data, $\lambda_{\mathbf{x}}$, for a predefined range of number of states (see Section II-C).
- 3) If \mathcal{H} is empty, then $\mathcal{H} = \{\lambda_{\mathbf{x}}\}$ and goto 2. Otherwise use the Forward algorithm to compute probabilities $P(\mathbf{x}|\lambda)$ and $N(\lambda)$ function to compute numbers of free parameters of the models

$$p_{\mathbf{x}} = P(\mathbf{x}|\lambda_{\mathbf{x}}) \quad n_{\mathbf{x}} = N(\lambda_{\mathbf{x}}) \quad (1)$$

$$p_{\mathcal{H}} = \max_{\lambda \in \mathcal{H}} P(\mathbf{x}|\lambda) \quad n_{\mathcal{H}} = \sum_{\lambda \in \mathcal{H}} N(\lambda) \quad (2)$$

- 4) Compute the information criteria values (e.g. AIC or BIC) for two cases: (C1) extending the current set of models with $\lambda_{\mathbf{x}}$ – possibly with better likelihood, but at the cost of expanding the total number of parameters – and (C2) staying with previous contents of the \mathcal{H} :

$$b_{\mathcal{H}+\lambda_{\mathbf{x}}} = IC(p_{\mathbf{x}}, n_{\mathcal{H}} + n_{\mathbf{x}}) \quad b_{\mathcal{H}} = IC(p_{\mathcal{H}}, n_{\mathcal{H}}) \quad (3)$$

- 5) If $b_{\mathcal{H}+\lambda_{\mathbf{x}}} < b_{\mathcal{H}}$ add the new model for the next iteration $\mathcal{H}' = \mathcal{H} \cup \{\lambda_{\mathbf{x}}\}$; if not, leave models as they were $\mathcal{H}' = \mathcal{H}$.
- 6) Regardless of decision in 5, calculate the anomaly score as $a_{\mathbf{x}} = p_{\mathbf{x}} - p_{\mathcal{H}}$. If there are remaining sequences, goto 2, otherwise stop.

The motivation for the algorithm presented above is as follows. While a HMM model has a very good performance in modelling time series, building a model of a long (monthly, yearly) series would require frequent, expensive re-training on a very long data history. For our case, initial observation of the data exposed dominant cycle or seasonality, in a similar way it is seen in energy consumption, water usage or weather patterns. Hence it's natural to treat the signal as a collection of cycle periods, in our case days, keeping in memory only an ensemble of HMM models, as the memory cost of a HMM model is much smaller than daily data sequence (see Section IV-A). Changes in the ensemble of HMMs occur between daily batches of data.

The proposed algorithm balances model complexity (number of parameters) and the ability to describe the signal (data fit). The collection of models retained on algorithm's progress over individual cycles serves additionally as a signal descriptor and source of diagnostic information.

C. Implementation of sensor verification

The verification system was incorporated into the fuel station tank monitoring system, which consists of the three main parts: the station part (implements software and hardware related to data acquisition, connects directly to devices and sends data to the central server); the server part (receives data from many stations, automatically processes this data and tries to draw and present preliminary conclusions [25]); the analytical part (responsible for analysing the results of the server part and for making decisions with human supervision).

For the actual verification, we focus on daily windows, as this corresponds to the rhythm of normal monitoring/verification applied in the system. A daily sequence of data \mathbf{x} is fed into the system, and its anomaly score $a_{\mathbf{x}}$ computed; if high enough, a 'require inspection' alert is generated. We note that there are additional possible ways to get information from the model, which are discussed in Section IV.

For each daily model identification (step 2), we use exhaustive search over a set of states $n \in \{1, \dots, 10\}$, with Bayesian Information Criterion (BIC) [26] for selection of the final model. As the HMM identification algorithm (the Baum-Welch procedure) can end in local optimum, $k = 10$ independent searches for given n are performed, and the best model is evaluated. The BIC is also used for mixture extension decision ($b_{\mathcal{H}+\lambda_{\mathbf{x}}}$ and $b_{\mathcal{H}}$ values).

III. RESULTS

A. Description of the sensor

Fuel and water level is measured by the Automatic Tank Gauging device (ATG). An ATG uses probes located in each tank or compartment to measure fuel and water levels. Each probe consists of a long rod with floats or sensors. The probe rod also has thermistors to measure the fuel temperature. The ATG sends an electrical impulse to both probes independently (product level float and water level float). The probe sends back the pulse and the ATG measures the time elapsed from sending to receiving. On this basis, the height is calculated. Measurements from all underground tanks are sent to a central unit located in the station building through wired or wireless connection. From here they are sent to the server (see Section II-C). The common risks with this type of device are: (1) suspension of the probe when it gets stuck at a certain height of the rod and (2) inertia of temperature measurements – especially important during the delivery of fuel with a significantly different temperature. Of the available data, we use fuel (product) level readout, which contains rich data about the tank situation; the remaining two (temperature and water) are used to diagnose specific, known problem conditions.

B. Selection of test cases

To test the method, data from $n_t = 15$ fuel tanks that have been previously known to have malfunctions and issues were selected for analysis. Both short and long term history sequences were selected, mean sequence length was 183 days (8 – 646 days) while mean sample count was $\approx 591\,189$ (30 151 – 2 531 218 samples). The sequences were annotated by experts (analytical team in charge of verification of the sensors), with a list of days with erroneous or anomalous readings. There are two types of outliers in the data: (1) related to real probe disturbances (e.g. when the float hangs on the rod and does not represent the height of the liquid accurately; or in the middle of a delivery and there are significant fuel fluctuations causing the float to sink temporarily) and (2) virtual errors (incorrect translation of the pulse length to the real height, occurs when raw current measurement values are mismeasured or misinterpreted).

C. Experimental procedure

As each data sequence consists of fuel volume measurements at irregular intervals, the sequences were differentiated, normalized by time delta, and standardized¹ prior to inclusion in the experiments. Each normalized and standardized sequence has been cut into day's windows \mathbf{x} and fed subsequently to the algorithm presented in Section II-B. For each sequence, the first half is treated as 'run-in' or training data, without using the labels (our case assumes they are unavailable during regular application). The resulting anomaly score $a_{\mathbf{x}}$ was recorded. The anomaly scores together with ground truth annotations were used to prepare ROC curves, with Area Under Curve (AUC) as the performance measure; for testing, positive labels were assigned to anomalies, while negative examples to normal levels.

Note that ground truth labels were generated especially for testing the proposed method, they are not required during the normal operation of the algorithm. We focus on evaluating the performance on the current set of data (current day), this is motivated by the performance measure of our underlying application setup, which is day to day monitoring of a fuel tank.

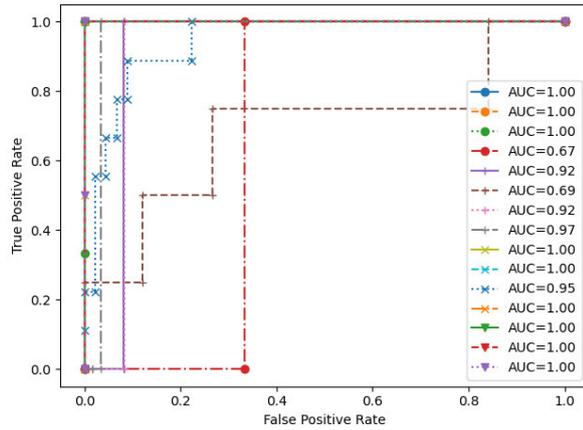
D. Algorithm performance

The average AUC score achieved by the method was 0.94 ± 0.11 . In nine cases (no. 1-3, 9, 10, 12-15), the anomaly score was precise enough to correctly single out all anomalies, achieving maximum possible AUC value of 1.0. In two cases (no. 4 and 6) the results were strongly affected by difficulty of the problem, as the anomalies show with small changes in the signal, resulting in AUC values of 0.67 and 0.69. Remaining four cases (no. 5, 7, 8, 11) achieve 0.92–0.97 (see Figure 1b). Almost all cases of mechanical failures (e.g. probe suspension) and faulty sensors were identified correctly. Anomalies (e.g. ripples or jumps within tank refill, especially in case 4)

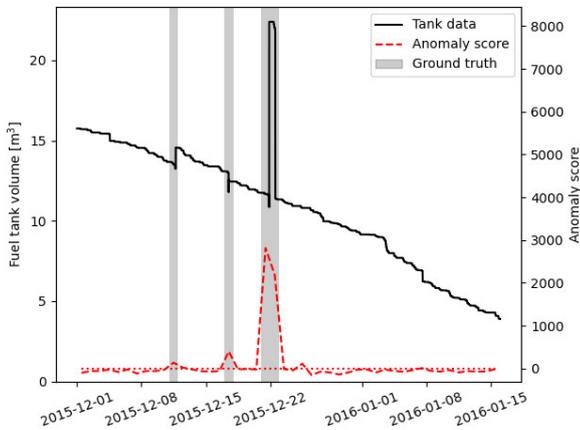
¹In actual application, the standardisation can be replaced by using mean and variance data from other tank with known history and similar characteristics, or estimated from physical tank properties, e.g. the total volume.



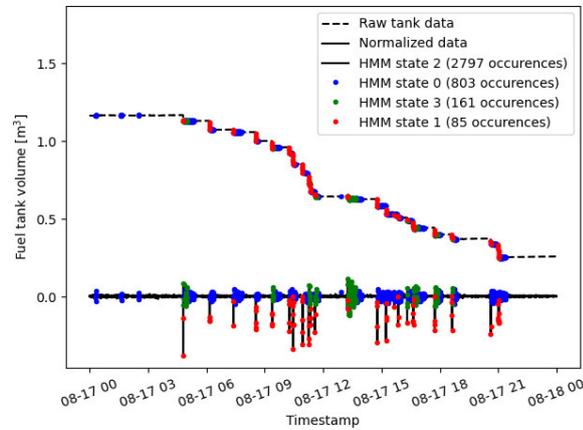
(a) An example of ATG sensor with two floating probes (for water and fuel).



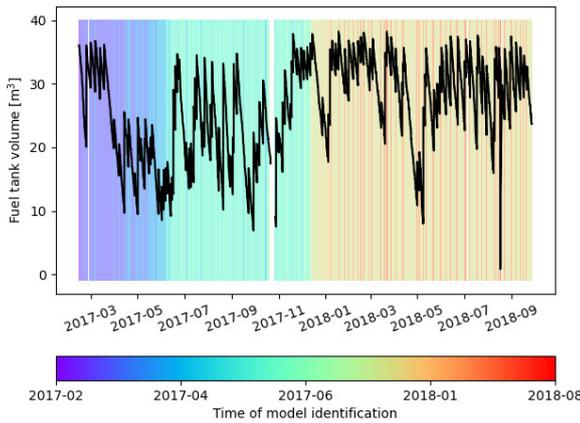
(b) Receiver operating characteristics (ROC) and area under ROC (AUC) values achieved in the experiments.



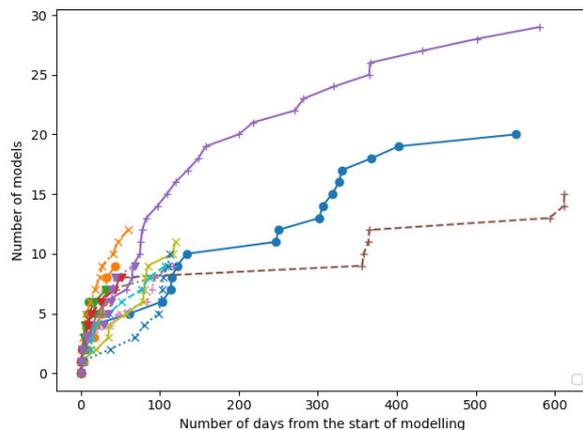
(c) Example of outlier detection – tank data and anomaly score computed by the proposed method. Note high values where outliers were identified.



(d) Example of typical HMM model – tank data unnormalized (raw) and normalized (see Section III-C), with state labels superimposed. Note the easy physical interpretation of identified states (see Section IV).



(e) Illustration of model assignments over time. Colour denotes time of identification of a model assigned to given sample. Note two visible trend changes at 2017-04 and 2018-01; e.g. through the second half of the 2017 most assigned models were identified in the period right after 2017-04.



(f) Number of models as a function of days from modelling start, for given set of tanks. Consistent patterns are visible (see Section IV).

Fig. 1: Illustration of the proposed method and experiments' results.

were more difficult to spot, but the performance remained acceptable. In rare cases of tanks with a longer history, model adaptation (i.e. the process of conditional adding of a new HMM to an ensemble, see step 5 in Section II-B) had been seen interfering with a detection (outlier present within a new pattern was learned and not detected subsequently); however those cases could be isolated at the cost of increasing the sensitivity and potentially the number of false positives. A small portion of the errors were traced back to imprecise labelling of bad cases. Example detections are presented in Figure 1c. For our current results, a correlation study could not be carried out (correlation coefficients not statistically significant); an analysis of the results suggests that better AUC scores are achieved by bigger models, but not necessarily with longer training or test length, both when measured in number of days or number of samples.

IV. DISCUSSION

A. Performance summary

Overall performance of the method was evaluated as very good, both in terms of quantitative score, and qualitative evaluation. Detailed inspection of the models revealed additional usage patterns, beyond the use of anomaly score. Adding a new model is usually connected with trend change of the series. If a high $p_{\mathcal{H}}$ value comes from a rare of old model (not matched lately, and previously matched to only a few cases), it may be additional signal of an anomaly. Finally some of the learned models could be tagged as interesting by the human observer and alert could be generated when they appear. The method has excellent aggregation properties; average number of parameters at the end of modelling was $n_p = 334$, which is less than 0.1% of the number of original samples.

B. Observed model behaviour

As expected, individual HMM within the ensemble represent daily behaviours of the tank being monitored. Example is presented in Figure 1d, note how the identified states (four in this example) correspond to known physical phenomena: stable level (s. 2), fuel unloading (s. 1), level oscillations after unloading (s. 3) and general or temperature induced oscillation (s. 0). Typically, the number of states is found in the range 3 – 7, and most of the time they can be assigned some interpretation based on what happens inside the tank. We consider this correspondence a qualitative validation of our approach. Individual HMM inspection revealed that they contain features common mainly for the fuel type (e.g. diesel, gasoline, premium); this may make it possible to produce a tank-independent dictionary of HMMs that could be used as an initialization of the algorithm.

The step 5 of the algorithms prevents adding new models if previous adequately explain current data – until there's a trend change, when a set of new models must be added to keep the model accurate. In the example presented in Figure 1e two trend changes can be easily observed. Those trend changes are usually explainable by process or physical change (e.g.

reassignment of fuel type for the tank, seasonal changes, general station usage type change resulting from roadworks). This trend change could be identified through analysis of model addition times, and provide additional monitoring information. Another view of this phenomenon can be seen in Figure 1f, where model count at given number of days from the start is presented. Often addition of models is seen in batches, on the beginning or when some trend change occurs.

C. Conclusions and future work

The proposed method has promising performance score, the resulting model has a high degree of explainability, limited memory and computation requirements and can be easily generalized to other domains of sensor verification.

As the objective of this work was a case study of the proposed algorithm, based on a mixture of HMMs, to the application of IoT sensor verification, further work will focus on extended analysis of the proposed approach, including: comparison with other approaches; detailed investigation of parameters, e.g. the effect of training sequence length; pruning the set of models \mathcal{H} to remove rarely used, old models.

ACKNOWLEDGEMENTS

Authors wish to thank Bartosz Grabowski and Wojciech Masarczyk for insightful cooperation on research issues with neural networks sensor models, which led to the development of the method presented here. Authors also wish to thank anonymous reviewers for their insightful comments and remarks, which were basis for improvement of this paper and plans for future work.

REFERENCES

- [1] R. Krishnamurthi, A. Kumar, D. Gopinathan, A. Nayyar, and B. Qureshi, "An overview of IoT sensor data processing, fusion, and analysis techniques," *Sensors*, vol. 20, no. 21, 2020. doi: 10.3390/s20216076
- [2] M. Henry and G. Wood, "Sensor validation: principles and standards," *atp International*, vol. 3, no. 2, 2005.
- [3] A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, "A review on outlier/anomaly detection in time series data," *ACM Computing Surveys*, vol. 54, no. 3, 2021. doi: 10.1145/3444690
- [4] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars, "A continual learning survey: Defying forgetting in classification tasks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3366–3385, 2022. doi: 10.1109/TPAMI.2021.3057446
- [5] J. M. Ramírez, F. Díez, P. Rojo, V. Mancuso, and A. Fernández-Anta, "Explainable machine learning for performance anomaly detection and classification in mobile networks," *Computer Communications*, vol. 200, pp. 113–131, 2023. doi: 10.1016/j.comcom.2023.01.003
- [6] N. Aslam, I. U. Khan, A. Alansari, M. Alrammah, A. Alghwairy, R. Alqahtani, R. Alqahtani, M. Almushikes, and M. A. Hashim, "Anomaly detection using explainable random forest for the prediction of undesirable events in oil wells," *Applied Computational Intelligence and Soft Computing*, vol. 2022, p. 1558381, 2022. doi: 10.1155/2022/1558381
- [7] F. Dama and C. Sinoquet, *Time Series Analysis and Modeling to Forecast: a Survey*. arXiv:2104.00164 [cs.LG], 2021.
- [8] M. Staniszewski, A. Skorupa, Ł. Boguszewicz, M. Sokół, and A. Polański, "Quality control procedure based on partitioning of NMR time series," *Sensors*, vol. 18, no. 3, p. 792, 2018. doi: 10.3390/s18030792
- [9] M. Staniszewski, A. Skorupa, Ł. Boguszewicz, A. Michalczyk, K. Wereszczyński, M. Wicher, M. Konopka, M. Sokół, and A. Polański, "Application of reiteration of hankel singular value decomposition in quality control," *AIP Conference Proceedings*, vol. 1863, no. 1, p. 400006, 2017. doi: 10.1063/1.4992575

- [10] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989. doi: 10.1109/5.18626
- [11] E. Lach, D. Grzechca, A. Polański, J. Rutkowski, and M. Staniszewski, "Analysis of semestral progress in higher technical education with HMM models," in *Computational Science – ICCS 2021*. Springer, 2021. doi: 10.1007/978-3-030-77967-2_18 pp. 214–228.
- [12] M. Cholewa and P. Glomb, "Natural human gestures classification using multisensor data," in *Proceedings 3rd IAPR Asian Conference on Pattern Recognition ACPR 2015*. IEEE, 2015. doi: 10.1109/ACPR.2015.7486553 pp. 499–503.
- [13] J. Ying, T. Kirubarajan, K. Pattipati, and A. Patterson-Hine, "A hidden Markov model-based algorithm for fault diagnosis with partial and imperfect tests," *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 30, no. 4, pp. 463–473, 2000. doi: 10.1109/5326.897073
- [14] W. Zhao, T. Shi, and L. Wang, "Fault diagnosis and prognosis of bearing based on Hidden Markov Model with multi-features," *Applied Mathematics and Nonlinear Sciences*, vol. 5, no. 1, pp. 71–84, 2020. doi: 10.2478/amns.2020.1.00008
- [15] F. van de Pol and R. Langeheine, "Mixed Markov latent class models," *Sociological Methodology*, vol. 20, pp. 213–247, 1990. doi: 10.2307/271087
- [16] M. D. R. de Chaumaray, M. Marbac, and F. Navarro, "Mixture of hidden Markov models for accelerometer data," *Annals of Applied Statistics*, vol. 14, no. 4, pp. 1834 – 1855, 2020. doi: 10.48550/arXiv.1906.01547
- [17] D. Vidotto, J. K. Vermunt, and K. V. Deun, "Multiple imputation of longitudinal categorical data through bayesian mixture latent Markov models," *Journal of Applied Statistics*, vol. 47, no. 10, pp. 1720–1738, 2020. doi: 10.1080/02664763.2019.1692794
- [18] W. Khreich, E. Granger, A. Miri, and R. Sabourin, "Adaptive ROC-based ensembles of HMMs applied to anomaly detection," *Pattern Recognition*, vol. 45, no. 1, pp. 208–230, 2012. doi: 10.1016/j.patcog.2011.06.014
- [19] P. R. Cavalin, R. Sabourin, C. Y. Suen, and A. S. Britto Jr., "Evaluation of incremental learning algorithms for HMM in the recognition of alphanumeric characters," *Pattern Recognition*, vol. 42, no. 12, pp. 3241–3253, 2009. doi: 10.1016/j.patcog.2008.10.012
- [20] W. Khreich, E. Granger, A. Miri, and R. Sabourin, "A survey of techniques for incremental learning of HMM parameters," *Information Sciences*, vol. 197, pp. 105–130, 2012. doi: 10.1016/j.ins.2012.02.017
- [21] K. Faber, R. Corizzo, B. Sniezynski, and N. Japkowicz, "Vlad: Task-agnostic vae-based lifelong anomaly detection," *Neural Networks*, vol. 165, pp. 248–273, 2023. doi: 10.1016/j.neunet.2023.05.032
- [22] D. Samariya and A. Thakkar, "A comprehensive survey of anomaly detection algorithms," *Annals of Data Science*, vol. 10, no. 3, pp. 829–850, 2023. doi: 10.1007/s40745-021-00362-9
- [23] L. Wang, X. Zhang, H. Su, and J. Zhu, *A Comprehensive Survey of Continual Learning: Theory, Method and Application*. arXiv:2104.00164 [cs.LG], 2023.
- [24] P. Glomb, M. Romaszewski, A. Sochan, and S. Opozda, "Unsupervised parameter selection for gesture recognition with vector quantization and hidden markov models," in *Human-Computer Interaction – INTERACT 2011*, P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler, Eds. Springer Berlin Heidelberg, 2011. doi: 10.1007/978-3-642-23768-3_14 pp. 170–177.
- [25] P. Foszner, A. Gruca, and J. Bularz, "Fuel pipeline thermal conductivity in automatic wet stock reconciliation systems," in *Advances in Data Mining. Applications and Theoretical Aspects: 16th Industrial Conference*. Springer, 2016. doi: 10.1007/978-3-319-41561-1_22 pp. 297–310.
- [26] G. Schwarz, "Estimating the dimension of a model," *Annals of Statistics*, vol. 6, pp. 461–464, 1978. doi: 10.1214/aos/1176344136