

Editorial Board Members

Joaquim Filipe 

Polytechnic Institute of Setúbal, Setúbal, Portugal

Ashish Ghosh

Indian Statistical Institute, Kolkata, India

Raquel Oliveira Prates 

Federal University of Minas Gerais (UFMG), Belo Horizonte, Brazil

Lizhu Zhou

Tsinghua University, Beijing, China


More information about this series at <http://www.springer.com/series/7899>

Joao Gama · Sepideh Pashami ·
Albert Bifet · Moamar Sayed-Mouchawe ·
Holger Fröning · Franz Pernkopf ·
Gregor Schiele · Michaela Blott (Eds.)


IoT Streams for Data-Driven Predictive Maintenance and IoT, Edge, and Mobile for Embedded Machine Learning

Second International Workshop, IoT Streams 2020
and First International Workshop, ITEM 2020
Co-located with ECML/PKDD 2020
Ghent, Belgium, September 14–18, 2020
Revised Selected Papers


Editors


Joao Gama 
University of Porto
Porto, Portugal

Albert Bifet 
Waikato University
Hamilton, New Zealand


Holger Fröning 
Heidelberg University
Heidelberg, Germany

Gregor Schiele 
University of Duisburg-Essen
Essen, Germany

Sepideh Pashami 
Halmstad University
Halmstad, Sweden

Moamar Sayed-Mouchawe 
University of Lille
Lille, France

Franz Pernkopf 
Graz University of Technology
Graz, Austria

Michaela Blott 
XILINX Research
Dublin, Ireland

ISSN 1865-0929 ISSN 1865-0937 (electronic)
Communications in Computer and Information Science
ISBN 978-3-030-66769-6 ISBN 978-3-030-66770-2 (eBook)
<https://doi.org/10.1007/978-3-030-66770-2>

© Springer Nature Switzerland AG 2020

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

IoT Streams 2020 Preface

Maintenance is a critical issue in the industrial context for preventing high costs and injuries. Various industries are moving more and more toward digitization and collecting “big data” to enable or improve the accuracy of their predictions. At the same time, the emerging technologies of Industry 4.0 empower data production and exchange, which leads to new concepts and methodologies for the exploitation of large datasets in maintenance. The intensive research effort in data-driven Predictive Maintenance (PdM) is producing encouraging results. Therefore, the main objective of this workshop is to raise awareness of research trends and promote interdisciplinary discussion in this field.

Data-driven predictive maintenance must deal with big streaming data and handle concept drift due to both changing external conditions and also normal wear of the equipment. It requires combining multiple data sources, and the resulting datasets are often highly imbalanced. The knowledge about the systems is detailed, but in many scenarios there is a large diversity in both model configurations as well as their usage, additionally complicated by low data quality and high uncertainty in the labels. Many recent advancements in supervised and unsupervised machine learning, representation learning, anomaly detection, visual analytics and similar areas can be showcased in this domain. Therefore, the overlap in research between machine learning and predictive maintenance has continued to increase in recent years.

This event was an opportunity to bring together researchers and engineers to discuss emerging topics and key trends. Both the previous edition of the workshop at ECML 2019 and the latest edition at ECML 2020 have been very popular.

Aims and Scope

This workshop welcomed research papers using Data Mining and Machine Learning (Artificial Intelligence in general) to address the challenges and answer questions related to the problem of predictive maintenance. For example, when to perform maintenance actions, how to estimate components’ current and future status, which data should be used, what decision support tools should be developed for prognostic use, how to improve the estimation accuracy of remaining useful life, and similar. It solicited original work, already completed or in progress. Position papers were also considered. The scope of the workshop covered, but was not limited to, the following:

- Predictive and Prescriptive Maintenance
- Fault Detection and Diagnosis (FDD)
- Fault Isolation and Identification
- Anomaly Detection (AD)
- Estimation of Remaining Useful Life of Components, Machines, etc.
- Forecasting of Product and Process Quality
- Early Failure and Anomaly Detection and Analysis
- Automatic Process Optimization

- Self-healing and Self-correction
- Incremental and evolving (data-driven and hybrid) models for FDD and AD
- Self-adaptive time-series-based models for prognostics and forecasting
- Adaptive signal processing techniques for FDD and forecasting
- Concept Drift issues in dynamic predictive maintenance systems
- Active learning and Design of Experiment (DoE) in dynamic predictive maintenance
- Industrial process monitoring and modelling
- Maintenance scheduling and on-demand maintenance planning
- Visual analytics and interactive Machine Learning
- Analysis of usage patterns
- Explainable AI for predictive maintenance

It covered real-world applications such as:

- Manufacturing systems
- Transport systems (including roads, railways, aerospace and more)
- Energy and power systems and networks (wind turbines, solar plants and more)
- Smart management of energy demand/response
- Production Processes and Factories of the Future (FoF)
- Power generation and distribution systems
- Intrusion detection and cybersecurity
- Internet of Things
- Smart cities

We received a total of 19 papers and 13 of those were accepted. Each paper was reviewed by three PC members and camera ready papers were prepared based on the reviewers' comments.

Many people contributed to making this workshop a successful event. We would like to thank the Program Committee members and additional reviewers for their detailed and constructive reviews, the authors for their well-prepared presentations, and all workshop attendees for their engagement and participation.

IoT Streams 2020 Organization

Program Committee Members

Bruno Veloso	LIAAD - INESC TEC, Portugal
Maria Pedroto	University of Porto, Portugal
Sofia Fernandes	University of Aveiro, Portugal
Vinícius Souza	University of New Mexico, USA
Paula Branco	University of Ottawa, Canada
Hadi Fanaee-T	Halmstad University, Sweden
Jonathan De Andrade Silva	University of São Paulo, Brazil
Rita P. Ribeiro	University of Porto, Portugal
Shazia Tabassum	University of Porto, Portugal
João Vinagre	INESC TEC, Portugal
Mario Cordeiro	ProDEI, Portugal
Nuno Moniz	INESC TEC/University of Porto, Portugal
Chetak Kandaswamy	FEUP, Portugal
Jacob Montiel	University of Waikato, New Zealand
Gustavo Batista	UNSW Sydney, Australia
Douglas Castilho	Federal University of Minas Gerais, Brazil
Raquel Sebastião	IEETA, Portugal
Heitor Murilo Gomes	University of Waikato, New Zealand
Carlos Ferreira	LIAAD INESC Porto LA, Portugal
Paulo Paraíso	INESCTEC, Portugal
Florent Masseglia	Inria, France
Indre Zliobaite	University of Helsinki, Finland
Brais Cancela	Universidade da Coruña, Spain
Thiago Andrade	INESC TEC, Portugal
Richard Hugh Moulton	Queen's University, Canada
Andre de Carvalho	University of São Paulo, Brazil
Anders Holst	RISE, Sweden
Elaine Faria	Federal University of Uberlândia, Brazil
Ana Nogueira	LIAAD - INESC TEC, Portugal
Herna Viktor	University of Ottawa, Canada

Workshop Organizers

Joao Gama	University of Porto, Portugal
Albert Bifet	Télécom Paris, France
Moamar Sayed Mouchaweh	IMT Lille Douai, France
Grzegorz J. Nalepa	Jagiellonian University, Poland
Sepideh Pashami	Halmstad University, Sweden

ITEM 2020 Preface

Background

There is an increasing need for real-time intelligent data analytics, driven by a world of Big Data and society's need for pervasive intelligent devices. Application examples include wearables for health and recreational purposes, infrastructure such as smart cities, transportation and smart power grids, e-commerce and Industry 4.0, and autonomous robots including self-driving cars. Most applications share facts like large data volumes, real-time requirements and limited resources including processor, memory and network. Often, battery life is a concern, data might be large but possibly incomplete, and probably most important, data can be uncertain. Notably, often powerful cloud services are unavailable, or not an option due to latency or privacy constraints.

For these tasks, Machine Learning (ML) is among the most promising approaches to address learning and reasoning under uncertainty. In particular deep learning methods in general are well-established supervised or unsupervised ML methods, and well understood with regard to compute/data requirements, accuracy and (partly) generalization. Today's deep learning algorithms dramatically advance state-of-the-art performance in terms of accuracy of the vast majority of AI tasks. Examples include image and speech processing such as image recognition, segmentation, object localization, multi-channel speech enhancement and speech recognition, and signal processing such as radar signal denoising, with applications as broad as robotics, medicine, autonomous navigation, recommender systems, etc.

As a result, ML is embedded in various compute devices, ranging from power cloud systems over fog and edge computing to smart devices. Due to the demanding nature of this workload, which is heavily compute- and memory-intensive, virtually all deployments are limited by resources, this being particularly true for edge, mobile and IoT. Among the results of these constraints are various specialized processor architectures, which are tailored for particular ML tasks. While this is helpful for a particular task, ML is advancing fast and new methods are introduced frequently. Notably, one can observe that very often the requirements of such tasks advance faster than the performance of new compute hardware, increasing the gap between application and compute hardware. This observation is emphasized by the slowing down of Moore's law, which used to deliver constant performance scaling over decades.

Furthermore, to address uncertainty and limited data, and to improve in general the robustness of ML, new methods are required, with examples including Bayesian approaches, sum-product networks, capsule networks, graph-based neural networks and many more. One can observe that, compared with deep convolutional neural networks, computations can be fundamentally different, compute requirements can substantially increase, and underlying properties like structure in computation are often lost.

As a result, we observe a strong need for new ML methods to address the requirements of emerging workloads deployed in the real world, such as uncertainty, robustness and limited data. In order to not hinder the deployment of such methods on various computing devices, and to address the gap between application and compute hardware, we furthermore need a variety of tools. As such, this workshop aimed to gather new ideas and concepts on ML methods for real-world deployment, methods for compression and related complexity reduction tools, dedicated hardware for emerging ML tasks, and associated tooling like compilers and mappers. Similarly, the workshop also aimed to serve as a platform to gather experts from ML and systems to jointly tackle these problems, creating an atmosphere of open discussions and other interactions.

Workshop Summary

In September 2020, the first edition of ITEM took place, collocated with ECML-PKDD as the premier European machine learning and data mining conference. Even though the workshop had to take place virtually, there was a lively discussion and interaction, also due to inspiring keynote presentations by Luca Benini from ETH Zürich (“From Near-Sensor to In-Sensor AI”) and Song Han from MIT (“MCUNet: TinyNAS and TinyEngine on Microcontrollers”). Those keynotes created the right environment for a couple of contributed talks in the areas of hardware, methods and quantization, coming from institutions including Universidade da Coruña, Heidelberg University, Bosch Research, University of Duisburg-Essen, Technical University of Munich, KU Leuven, Università di Bologna and Graz University of Technology, among others.

Early take-aways include on the hardware side observations on open-source digital hardware (PULP) as well as analog hardware (BrainScaleS-2) as promising emerging alternatives to established architectures, that code generation for specialized hardware can be challenging, and that designing processor arrays is more difficult than one might think. From a methodological point of view, 8bit seems to be a natural constant when it comes to quantization, and time-multiplexing as well as on-device learning can be viable options. On the compression side, observations include that predictive confidence can help dynamic approaches to switching among models, heterogeneous uniform quantization as well as application-specific (radar) quantization.

Outlook

It is planned to continue ITEM for the next couple of years, so any interested researcher or scientist is invited to contribute to future editions. Also, while ITEM’s main focus is to be an academic platform with peer-reviewed contributions, there is also a more informal counterpart called the Workshop on Embedded Machine Learning (WEML), which is held annually at Heidelberg University. WEML is distinguished from ITEM by being a platform that only includes invited presentations from the community for mutual updates on recent insights and trends, but without the rigorous demands of scientific peer review. For more information about these two workshops, please refer to:

ITEM: <https://www.item-workshop.org>

WEML: <https://www.deepchip.org>

Finally, the co-organizers of ITEM would like to acknowledge the comprehensive commitment of the Workshop Co-Chairs at ECML-PKDD (Myra Spiliopoulou and Willem Waegeman), who had to face a mandatory shift to online events due to the pandemic, which they handled very swiftly and with excellent communication and organization. Similar acknowledgements go to the time and effort spent by our program committee, and last but not least the strong commitment of our program chair (Benjamin Klenk). Ultimate acknowledgement goes to Springer for publishing the workshop's proceedings.

Holger Fröning
Franz Pernkopf
Gregor Schiele
Michaela Blott

ITEM 2020 Organization

Co-organizers

Holger Fröning	Heidelberg University, Germany (holger.froening@ziti.uni-heidelberg.de)
Franz Pernkopf	Graz University of Technology, Austria (pernkopf@tugraz.at)
Gregor Schiele	University of Duisburg-Essen (gregor.schiele@uni-due.de)
Michaela Blott	XILINX Research, Ireland (michaela.blott@xilinx.com)

Program Chair

Benjamin Klenk	NVIDIA Research, USA
----------------	----------------------

Technical Program Committee

Costas Bekas	Citadel Securities, Switzerland
Herman Engelbrecht	Stellenbosch University, South Africa
Giulio Gambardella	XILINX, Ireland
Tobias Golling	University of Geneva, Switzerland
Domenik Helms	OFFIS e.V. - Institut für Informatik, Germany
Eduardo Rocha Rodrigues	IBM, Brazil
David King	Air Force Institute of Technology, USA
Manfred Mücke	Materials Center Leoben Forschung GmbH, Austria
Dimitrios S. Nikolopoulos	Virginia Tech, USA
Robert Peharz	Eindhoven University of Technology, The Netherlands
Marco Platzner	Paderborn University, Germany
Thomas B. Preußner	ETH Zürich, Switzerland
Johannes Schemmel	Heidelberg University, Germany
Wei Shao	Royal Melbourne Institute of Technology (RMIT), Australia
David Sidler	Microsoft, USA
Jürgen Teich	Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany
Sebastian Tschitschek	University of Vienna, Austria
Nicolas Weber	NEC Labs Europe, Germany
Matthias Zöhrer	Evolve.tech, Austria

Contents

IoT Streams 2020: Stream Learning

Self Hyper-parameter Tuning for Stream Classification Algorithms	3
<i>Bruno Veloso and João Gama</i>	
Challenges of Stream Learning for Predictive Maintenance in the Railway Sector	14
<i>Minh Huong Le Nguyen, Fabien Turgis, Pierre-Emmanuel Fayemi, and Albert Bifet</i>	
CycleFootprint: A Fully Automated Method for Extracting Operation Cycles from Historical Raw Data of Multiple Sensors	30
<i>Hadi Fanaee-T, Mohamed-Rafik Bouguelia, Mahmoud Rahat, Jonathan Blixt, and Harpal Singh</i>	
Valve Health Identification Using Sensors and Machine Learning Methods	45
<i>M. Atif Qureshi, Luis Miralles-Pechuán, Jason Payne, Ronan O'Malley, and Brian Mac Namee</i>	
Failure Detection of an Air Production Unit in Operational Context	61
<i>Mariana Barros, Bruno Veloso, Pedro M. Pereira, Rita P. Ribeiro, and João Gama</i>	

IoT Streams 2020: Feature Learning

Enhancing Siamese Neural Networks Through Expert Knowledge for Predictive Maintenance	77
<i>Patrick Klein, Niklas Weingarz, and Ralph Bergmann</i>	
Explainable Process Monitoring Based on Class Activation Map: Garbage In, Garbage Out	93
<i>Cheolhwan Oh, Junhyung Moon, and Jongpil Jeong</i>	
AutoML for Predictive Maintenance: One Tool to RUL Them All	106
<i>Tanja Tornede, Alexander Tornede, Marcel Wever, Felix Mohr, and Eyke Hüllermeier</i>	
Forklift Truck Activity Recognition from CAN Data	119
<i>Kunru Chen, Sepideh Pashami, Sławomir Nowaczyk, Emilia Johansson, Gustav Stenelöv, and Thorsteinn Rögnvaldsson</i>	

Embeddings Based Parallel Stacked Autoencoder Approach for Dimensionality Reduction and Predictive Maintenance of Vehicles	127
<i>Vandan Revanur, Ayodeji Ayibiowu, Mahmoud Rahat, and Reza Khoshkangini</i>	

IoT Streams 2020: Unsupervised Machine Learning

Unsupervised Machine Learning Methods to Estimate a Health Indicator for Condition Monitoring Using Acoustic and Vibration Signals: A Comparison Based on a Toy Data Set from a Coffee Vending Machine	145
<i>Yonas Tefera, Maarten Meire, Stijn Luca, and Peter Karsmakers</i>	

Unsupervised Anomaly Detection for Communication Networks: An Autoencoder Approach	160
<i>Pieter Bonte, Sander Vanden Haute, Annelies Lejon, Veerle Ledoux, Filip De Turck, Sofie Van Hoecke, and Femke Ongenae</i>	

Interactive Anomaly Detection Based on Clustering and Online Mirror Descent.	173
<i>Lingyun Cheng, Sadhana Sundaresh, Mohamed-Rafik Bouguelia, and Onur Dikmen</i>	

ITEM 2020: Hardware

hxtorch: PyTorch for BrainScaleS-2: Perceptrons on Analog Neuromorphic Hardware	189
<i>Philipp Spilger, Eric Müller, Arne Emmel, Aron Leibfried, Christian Mauch, Christian Pehle, Johannes Weis, Oliver Breitwieser, Sebastian Billaudelle, Sebastian Schmitt, Timo C. Wunderlich, Yannik Stradmann, and Johannes Schemmel</i>	

Inference with Artificial Neural Networks on Analog Neuromorphic Hardware	201
<i>Johannes Weis, Philipp Spilger, Sebastian Billaudelle, Yannik Stradmann, Arne Emmel, Eric Müller, Oliver Breitwieser, Andreas Grübl, Joscha Ilmberger, Vitali Karasenko, Mitja Kleider, Christian Mauch, Korbinian Schreiber, and Johannes Schemmel</i>	

Search Space Complexity of Iteration Domain Based Instruction Embedding for Deep Learning Accelerators	213
<i>Dennis Rieber and Holger Fröning</i>	

On the Difficulty of Designing Processor Arrays for Deep Neural Networks.	229
<i>Kevin Stehle, Günther Schindler, and Holger Fröning</i>	

ITEM 2020: Methods

When Size Matters: Markov Blanket with Limited Bit Depth Conditional Mutual Information	243
<i>Laura Morán-Fernández, Eva Blanco-Mallo, Konstantinos Sechidis, Amparo Alonso-Betanzos, and Verónica Bolón-Canedo</i>	
Time to Learn: Temporal Accelerators as an Embedded Deep Neural Network Platform	256
<i>Christopher Cichiwskyj, Chao Qian, and Gregor Schiele</i>	
ML Training on a Tiny Microcontroller for a Self-adaptive Neural Network-Based DC Motor Speed Controller	268
<i>Frederik Funk, Thorsten Bucksch, and Daniel Mueller-Gritschneider</i>	

ITEM 2020: Quantization

Dynamic Complexity Tuning for Hardware-Aware Probabilistic Circuits	283
<i>Laura I. Galindez Olascoaga, Wannes Meert, Nimish Shah, and Marian Verhelst</i>	
Leveraging Automated Mixed-Low-Precision Quantization for Tiny Edge Microcontrollers	296
<i>Manuele Rusci, Marco Fariselli, Alessandro Capotondi, and Luca Benini</i>	
Author Index	309