

SMARTHEP Meets Industry @CERN


TECHNOLOGY OFFERS



Funded by the European Union's Horizon 2020 research and innovation programme, call H2020-MSCA-ITN-2020, under Grant Agreement n. 956086



CONFINDUSTRIA
PIEMONTE



FONDAZIONE
PIEMONTE
INNOVA




DIGITAL
INNOVATION
HUB
PIEMONTE

- **MOBILITY/AEROSPACE.....1 - 5**
- **SMART MANUFACTURING/INDUSTRY 4.0.....6 - 9**
- **HEALTHCARE.....10 - 11**
- **FINANCE/INSURANCE.....10 - 15**
- **ENERGY MANAGEMENT.....16**


*N.B. For each technology offer, the **main (non-exhaustive) industrial application areas** are indicated, along with the **industrial partner of the SMARTHEP project** where the technologies have been developed or tested, if available.*



Funded by the European Union's Horizon 2020 research and innovation programme, call H2020-MSCA-ITN-2020, under Grant Agreement n. 956086



**CONFINDUSTRIA
PIEMONTE**



**FONDAZIONE
PIEMONTE
INNOVA**



**DIGITAL
INNOVATION
HUB
PIEMONTE**

LEARNING TRAFFIC ANOMALIES FROM GENERATIVE MODELS ON REAL- TIME OBSERVATIONS

ESR: Fotis
Giasemis

Sorbonne
University

 x i m a n t i s
CERN

Accurate **detection of traffic anomalies** is crucial for effective urban **traffic management** and **congestion mitigation**. The Spatiotemporal Generative Adversarial Network (STGAN) framework, combining Graph Neural Networks and Long Short-Term Memory networks, is employed to capture complex spatial and temporal dependencies in traffic data. The STGAN is applied to real-time, minute-by-minute observations from 42 traffic cameras across Gothenburg, Sweden, collected over several months. The images are processed to compute a flow metric representing **vehicle density**, which serves as input for the model. Results demonstrate that the model effectively detects traffic anomalies with high precision and low false positive rates. Detected anomalies include camera signal interruptions, visual artifacts, and extreme weather conditions affecting traffic flow.

GRAPH NEURAL NETWORK TRAFFIC PREDICTIONS ON REAL-TIME OBSERVATIONS

ESR: Joachim
Carlo Kristian
Hansen

Lund
University

 x i m a n t i s

The technology utilises real-time traffic data from cameras in Gothenburg and applies advanced Graph Neural Networks (GNNs) to predict road **traffic density** on an hourly basis. Each traffic camera is modeled as a point in a network, where connections represent the road layout and distance between locations. To detect vehicles, YOLOv5, a **state-of-the-art object detection algorithm**, is employed, accurately measuring traffic density while filtering out background noise, such as shadows and reflections. Testing different GNN models on datasets of varying sizes revealed that increasing the training data from 3 days to 14 days improved prediction accuracy, reducing the Mean Absolute Percentage Error (MAPE) from 1.5% to 1.0%. This approach demonstrates the potential of GNNs in **traffic forecasting**, showing that more extensive datasets enhance **predictive performance**. The method captures not only short-term rush-hour variations but also provides insights into longer-term traffic patterns, making it valuable for urban mobility planning and smart traffic management.

SIMULATION OF STREET VIDEO SCENES FOR OBJECT DETECTION MODELS IN COMPUTER VISION

**ESR: Carlos
Cocha**

**University of
Heidelberg**

**verizon[✓]
connect**

At Verizon Connect, **street video scenes** were simulated for different shifts/rotations of the object detector (camera) to be used as input for testing the performance of a **novel object detection model** based on Convolutional Neural Networks, widely applied in **Computer Vision**. The results showed that **simulated street videos** could potentially perform as well as real street videos that are currently used for computer vision.

ROAD SCENE UNDERSTANDING FOR RISK ANTICIPATION FROM EGO-VEHICLE DATA

ESR: Henrique
Piñeiro
Monteagudo

University
of Bologna

verizon
connect

This technology represents the environment surrounding a road vehicle and **anticipates potential risks** using cameras installed on the vehicle dashboards. The main innovation compared to existing technology is the ability to train machine learning models with less annotated data and in settings with more affordable sensors. This technology can be applied in the **intelligent vehicles and transportation sector**.

EFFECTS OF VIEWPOINT SHIFTS IN BIRD-EYE VIEW SEGMENTATION

**ESR: Patin
Inkaew**

**University
of Helsinki**

**verizon
connect**

Using **cameras** and **sensors** around the **vehicle** to reconstruct a **Bird-Eye View (BEV)** image is an important piece in many driving tasks because BEV provides a compact overview description of the road scene. **BEV segmentation** aims to identify different objects in the scene, for example, cars, streets, and sidewalks in BEV perspectives. Models are often trained with images collected with a single fixed camera pose. However, in practice, it is mostly impossible to install cameras with the same camera pose the model is trained on, in particular, when the cameras are installed in different vehicles, for example, personal cars versus transport trucks. This leads to viewpoint shifts and a degradation in model performance. This project aims to analyze the effects of **viewpoint shifts in BEV segmentation tasks** and bring more attention to this problem within the community.

STUDY OF INDUSTRIAL TIME SERIES DATA FOR SEGMENTATION AND PATTERN IDENTIFICATION

**ESR: Daniel
Magdalinski**

**National Institute
for Subatomic
Physics / VU
Amsterdam**

point8
data matters.

Accurate **industrial process analysis** can be used to optimize processes, detect faults, or extract meaningful insights. **Industrial time-series data** were analyzed, focusing on the detection and segmentation of peaks, an approach useful for identifying signals, anomalies, and extracting signal regions in various industrial processes. By leveraging deep learning techniques, this approach aimed to enhance the reliability and automation of peak identification, reducing the need for manual inspection and calibration, and **improving decision-making in industrial environments.**

DEVELOPMENT OF SIMULATION-BASED INFERENCE INTERFACES IN THE ROOT SOFTWARE FRAMEWORK FOR INDUSTRIAL APPLICATIONS

ESR: James Gooding

TU Dortmund

CERN

Simulation-based inference (SBI) is a statistical approach used when traditional likelihood functions are difficult or expensive to compute. Instead of relying on explicit probability distributions, SBI uses machine learning classifiers trained on simulations to infer parameters from data, making it ideal for complex systems. This method is widely applicable in industries such as **predictive maintenance**, where it models equipment behavior to forecast failures, and **financial forecasting**, where it infers market parameters. The ROOT software framework, commonly used in high-energy physics, has integrated SBI and machine learning tools within the RooFit toolkit, extending its functionality for industrial decision-making.

REAL-TIME DATA ACQUISITION AND ANALYSIS FOR HIGH-FREQUENCY MONITORING IN INDUSTRY 4.0 AND BEYOND

ESR: Max Amerl

University of Manchester

CERN

Searches for low-mass particles using the ATLAS detector are constrained by data acquisition limits, as only a finite amount of data can be stored per second. High event rates from low-momentum signatures require stringent selection criteria, which may lead to data loss. To address this, **real-time analysis strategies** optimize data acquisition by reducing the amount of information stored per event, allowing for more flexible selection criteria. These strategies have **industrial applications** where **real-time monitoring** and **fast data analysis** are essential. Industry 4.0 applies them to **predictive maintenance** and **quality control**. Energy management benefits from optimized grid control. In healthcare, they improve real-time diagnostics and patient monitoring. In aerospace and mobility, they enhance sensor fusion for autonomous systems. In finance, they support high-frequency trading and fraud detection.

OPTIMIZED MACHINE LEARNING INFERENCE THROUGH HETEROGENEOUS ARCHITECTURES

ESR: Sanjiban
Sengupta

University of
Manchester

CERN

Machine Learning supports a wide range of tasks in LHC experiments, from simulation and reconstruction to anomaly detection and real-time data analysis. These applications require fast, low-latency inference integrated into complex physics workflows—an increasing challenge with growing data rates.

To address this, the ML4EP team at CERN is developing SOFIE, a ROOT/TMVA tool that converts trained ML models (from Keras, PyTorch, ONNX) into optimized C++ code with minimal dependencies. With a Python interface and support for ONNX operations—including Graph Neural Networks—SOFIE enables **real-time inference within event-based workflows**. Recent improvements include better memory management, kernel-level optimizations, and support for heterogeneous architectures (via SYCL and ALPAKA), enabling GPU inference on both NVIDIA and AMD platforms. While designed for high-energy physics, SOFIE is also well suited to other high-throughput, resource-constrained environments such as **industrial monitoring, real-time decision-making, traffic prediction, fraud detection, and autonomous systems**.

JET ENERGY CALIBRATION WITH REAL-TIME MACHINE LEARNING TRAINING AND ANOMALY DETECTION IN HIGH-ENERGY PHYSICS AND FINANCE

ESR: Laura Boggia

Sorbonne University

IBM

Anomaly detection in multivariate time series is an important problem **across various fields**, including healthcare, financial services, manufacturing, and physics detector monitoring. Accurately **identifying the moments when unexpected errors or defects occur** is essential for these applications. However, this can be challenging as the types of anomalies are unknown beforehand. In healthcare, for example, time series data, such as heart rates, can be analyzed to identify irregularities. This research evaluates, compares, and develops algorithms for fraud detection.

LEVERAGING MACHINE LEARNING FOR REAL-TIME DECISION MAKING: APPLICATIONS IN FINANCE AND HEALTHCARE

ESR: Danielle
Joan Wilson

University of
Manchester

CERN

At high energy physics experiments, like ATLAS, the trigger system is responsible for deciding, in a fraction of a second, whether to keep or discard collision data. With high luminosity, many collisions happen at once (pileup), making it harder to extract meaningful physics signals. Machine learning methods can be applied at the trigger, in order to **reject pileup in real time in large datasets**, so that only "interesting" events are saved. Such **real-time decision-making** in noisy environments has **potential applications in various sectors**. For example, in fraud detection, high-frequency trading, and medical diagnostic tools, such as Magnetic Resonance Imaging and Computed Tomography.

NEXT-GEN BANKING FRAUD SIMULATION WITH AGENT-BASED MODELING AND LLMs

**ESR: Micol
Olocco**

**TU
Dortmund**

IBM

The research focuses on developing an **advanced fraud simulation framework for banking**, integrating agent-based modeling and probabilistic reasoning. The key innovation lies in **using Large Language Models (LLMs)** to generate adaptive agent behaviors, improving realism. Unlike static rule-based fraud detection, this approach dynamically generates **realistic transaction sequences based on behavioral patterns**. It simulates both legitimate and fraudulent activities, allowing for testing detection algorithms against evolving fraud tactics. Current public datasets for banking transactions are limited in realism, often lacking the complexity and variability of real-world behavior, making it difficult to develop effective anomaly detection algorithms. This simulation aims to address these limitations by providing high-fidelity, diverse fraud scenarios, **enhancing AI-driven fraud detection systems**.

REAL-TIME MARKET ANALYSIS WITH RECURRENT NEURAL NETWORKS FOR PREDICTIVE INSIGHTS

**ESR: Leon
Bozianu**

**University of
Geneva**

LIGHTBOX

Recurrent Neural Networks (RNNs) have become a powerful tool in modern machine learning, particularly for analyzing complex sequential data. Unlike traditional networks, RNNs excel at capturing temporal dependencies, making them highly effective in financial market analysis. By maintaining an internal memory, these models can identify **long-term patterns and trends in time-series data**, such as **stock prices or market behavior**. RNNs are also widely used in forecasting, helping financial institutions predict market fluctuations based on historical data. This capability is crucial for **real-time market insights** and **decision-making**. Additionally, RNNs have potential **business applications** like **sales forecasting and market positioning** across various sectors.

STUDY OF HISTORICAL TIMESERIES OF FINANCIAL DATA TO DEVELOP MACHINE LEARNING APPLICATIONS FOR ANALYSIS AND PREDICTION OF FINANCIAL ASSETS

**ESR: Sofia
Cella**

**CERN /
University of
Geneva**

LIGHTBOX

Advanced machine learning techniques, including boosted decision trees, Bayesian models, and recurrent neural networks, are applied and adapted to volatile financial markets. **AI systems for predictive analysis** are developed using **financial time-series data**, providing insights into financial concepts and highlighting differences in data processing when compared to high-energy physics. Transfer learning is employed to enhance model training, particularly when working with limited historical data.

ACCELERATED ANOMALY DETECTION FOR FRAUD PREVENTION

**ESR: Pratik
Jawahar**

**University of
Manchester**

CERN

Anomaly detection is essential for **fraud detection in financial institutions**. The work aims to select anomalies out of a wide range of "normal" samples with as few assumptions about the anomaly as possible. This ensures that the algorithm is robust to any type of anomaly and not just the most commonly occurring ones.

ACCELERATED ANOMALY DETECTION FOR ENERGY MANAGEMENT

**ESR: Pratik
Jawahar**

**University of
Manchester**

CERN

The work on **CPU-GPU (Central Processing Unit - Graphics Processing Unit)** hybrid architecture aims to optimize parallelization in code, **improving both runtime performance and energy efficiency**. It builds on the philosophy that fast algorithms are inherently low-power consumers, addressing both aspects equally.