

Silesian University of Technology

RESEARCH UNIVERSITY

Faculty of Automatic Control, Electronics and Computer Science
Department of Applied Informatics

BY: **Alexandre Niyomugaba**

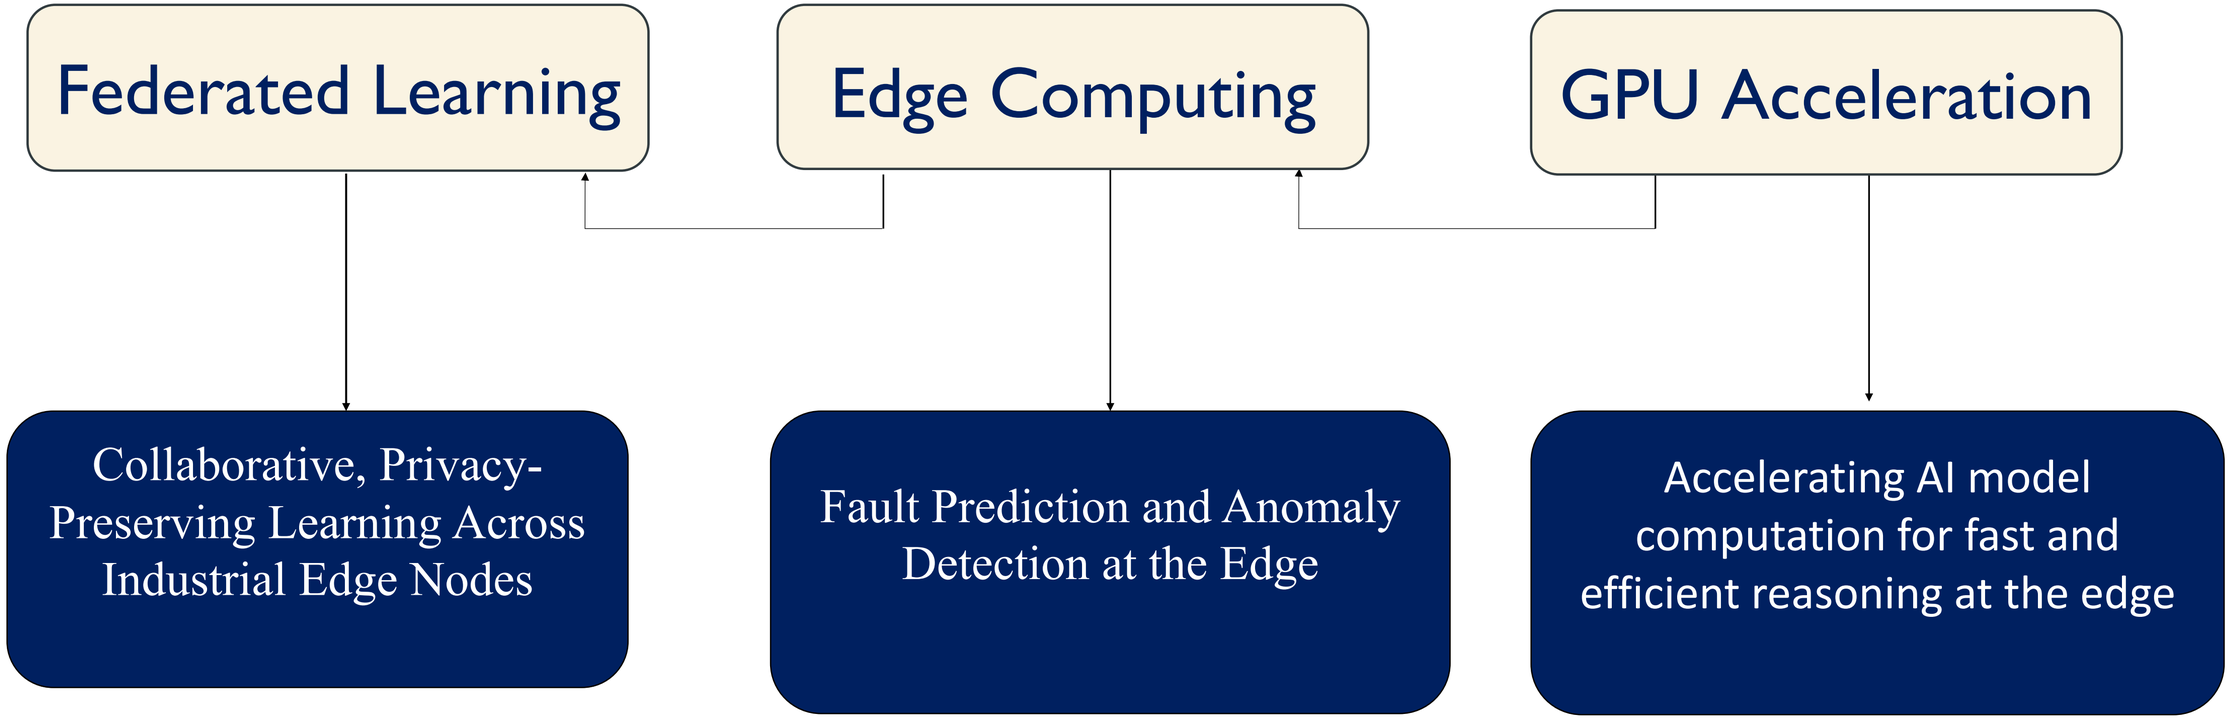
Main Supervisor: Prof. Dariusz Mrozek (SUT)

Co-supervisors: Prof. Francesco Piccialli (UNINA), Prof. David Camacho (UPM), Prof. Jia-Chun Lin (NTNU)

R&D cooperation: AIUT

MY ROLE IN TUAI PROJECT

Project Title: GPU-Accelerated Edge Computing for Federated Learning Reasoning in Industrial Environments



The project aims to develop Federated Learning (FL) framework with direct GPU acceleration at the edge devices, enabling efficiency training, updating and testing the prediction, and reasoning processes that are performed at the Edge of a network.

RESEARCH CHALLENGE

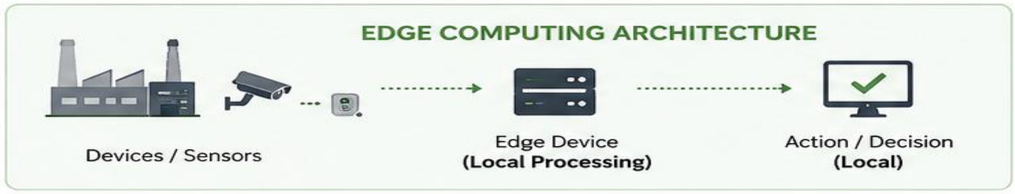
GPU-Accelerated
Federated
Edge AI

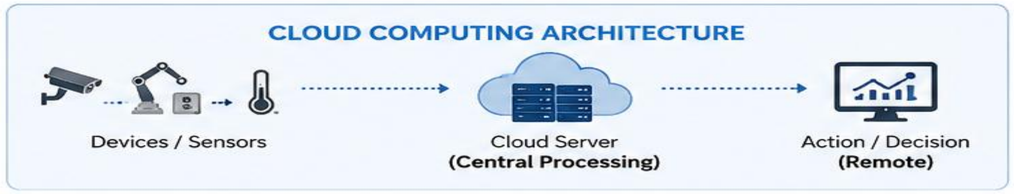
Achieving efficient distributed anomaly detection and fault prediction on resource-constrained federated industrial edge devices

Why Edge Computing for Industrial AI matters

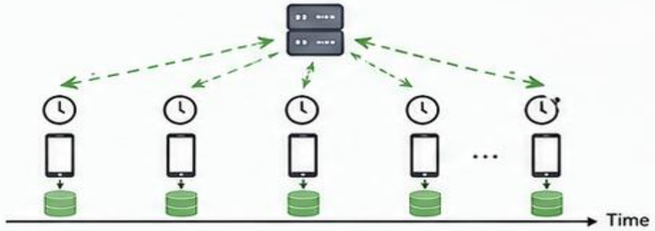
Edge Computing vs Cloud Computing

Features	Edge Computing	Cloud Computing
Data Location	On-site	Remote Data
Connection	Works Offline (Autonomous)	Online (Requires constant connectivity)
Memory requirements	Limited	High
Data Handling	Processes data locally	Sends all data to remote server
Privacy / Security	High Privacy	Lower Privacy
Application	Decentralized Intelligent systems	Centralized Intelligent systems





Federated Learning (FL) vs Federated Continual Learning (FCL)

Features	Federated Learning (FL)	Federated Continual Learning (FCL)
Clients training	Clients usually train on their entire available local dataset. 	Non-stationary data that changes over time. 
Data Nature	Fixed data distributions across clients. 	Evolving or sequential Data. 
Memory requirements	Clients usually train on their entire available local dataset. 	Clients often cannot store all past data to manage the memory. 
Aggregation	Fixed rounds. 	Task-based. 
Knowledge Transfer	Focuses on spatial transfer (from client to server and back). 	Focuses on both spatial (client-to-server) and temporal (past-to-present) transfer. 





PROGRESS

GPU-Enabled Edge-based Federated Continual Learning with Resource Adaptation for Efficient Non-Stationary Visual Anomaly Detection

Publisher: IEEE | Cite This | PDF

Alexandre Niyomugaba ; Dariusz Mrozek All Authors

Problem definition:

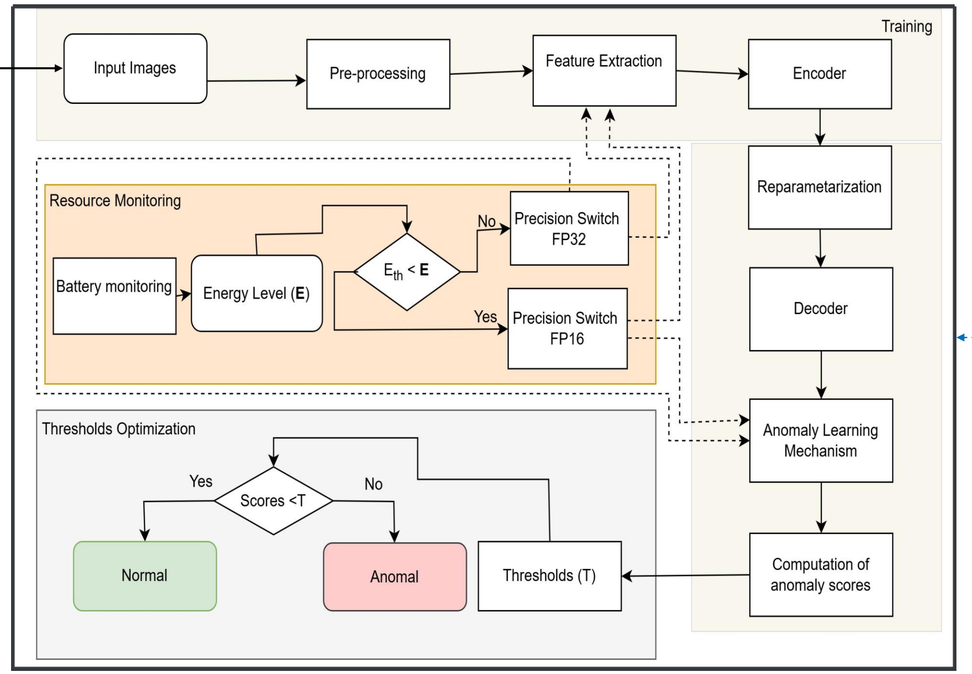
- Continual detection on non-stationary visual inputs remains challenging due to limited local data, statistical heterogeneity.
- Single-precision deep models lead to performance degradation when operating under resource-scarce conditions, such as low battery, due to the absence of adaptive resource management.

PROGRESS……

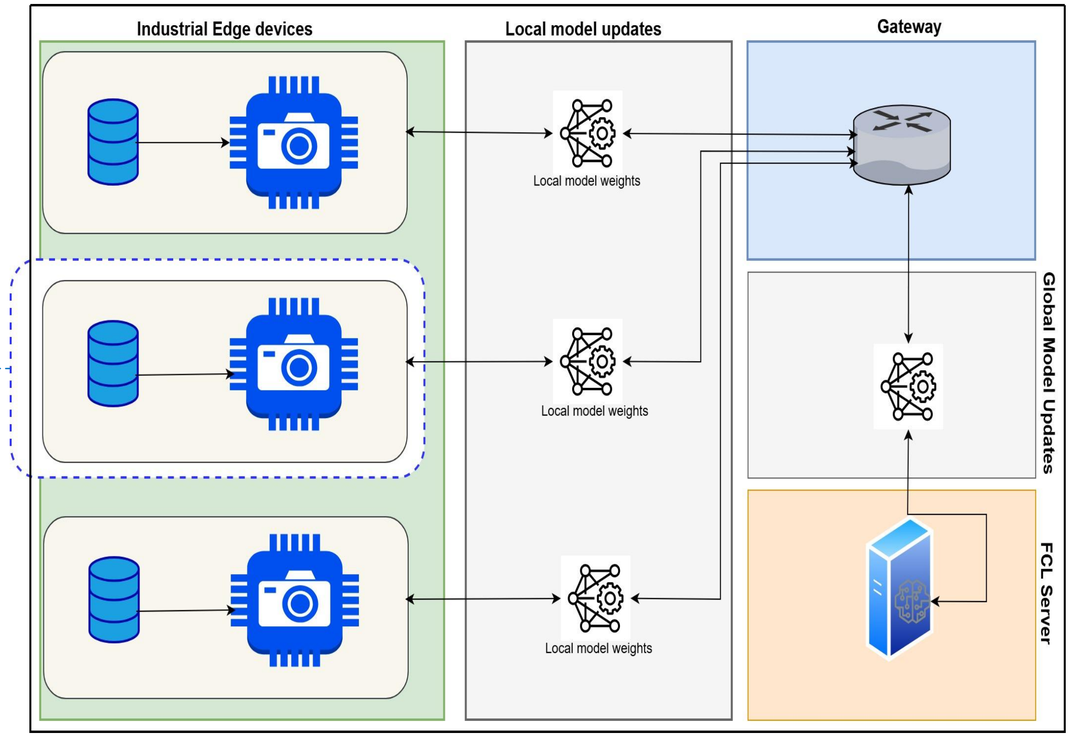
METHODOLOGY

Dataset	MVTec AD Dataset
Data type	Visual Data
Categories	15 Categories
Category types	10 Objects and 5 textures
Dataset size	5,354 Images
Training set	3629 Images
Test set	1725 Images
Application	Anomaly detection

Anomaly detection architecture at the edge



Industrial Federated learning Architecture



Solutions:

Edge-FCL, a Federated Continual Learning framework for dynamic, energy-efficient anomaly detection at the industrial edge.

PROGRESS……

EXPERIMENT



Server



Gateway



Edge device 1

Local



Edge device 2

Local



Edge device 3

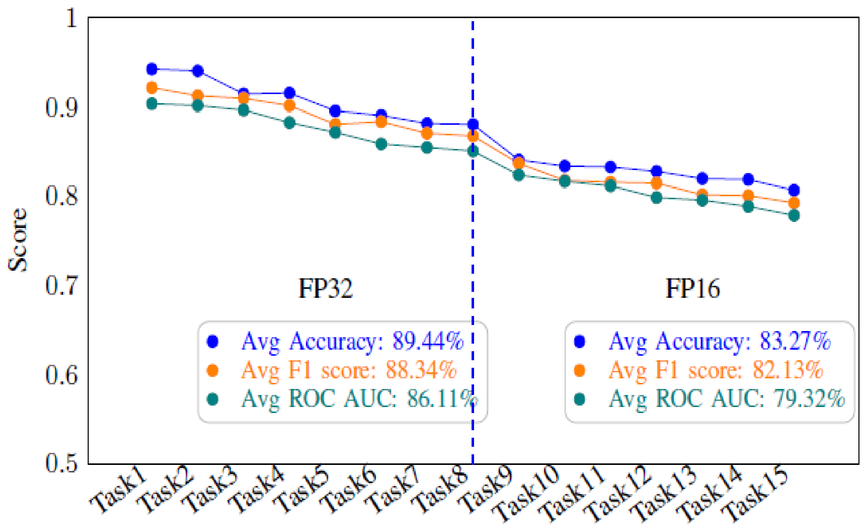
Local

Architecture	Details
AI performance	**40 TOPS** (Tera Operations Per Second), providing enough compututational power
GPU Architecture	**NVIDIA Ampere architecture** with **1024 CUDA cores** and **32 Tensor Cores**
Memory Configulation	8GB 128-bit LPDDR5memory (68 GB/s bandwidth)
CPU Power	**6-core Arm® Cortex®-A78AE v8.2** 64-bit
Power Efficiency:	Operates power ranging from **7W to 15W**

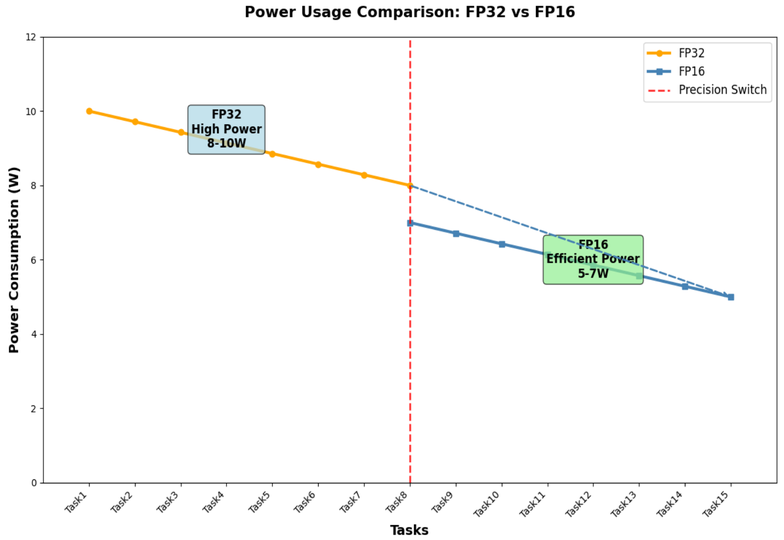
PROGRESS

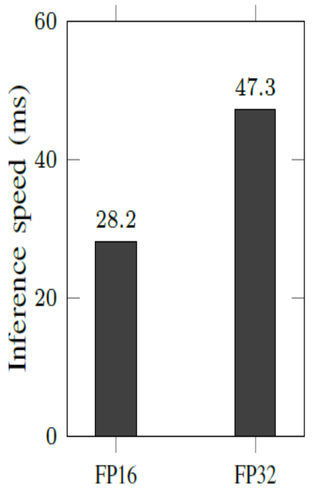
EXPERIMENTAL EVALUATION

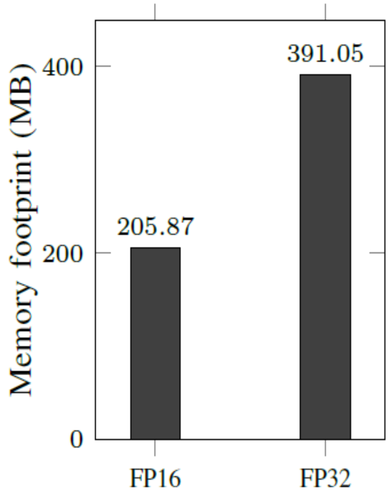
Performance



Precision Switch







By constraining the model to run solely in FP32 precision, we observed a critical reduction in energy efficiency, increasing the average inference time from 28.2 ms to 47.3 ms with an increasing of 1.8 of the memory footprint to 391.05MB, which, while retaining the highest performance 89.44 % accuracy, significantly, limited the operational longevity of the resource-constrained edge device.

PROGRESS ….

TITLE: EDGE-BASED SAFETY OPTIMIZATION: LIGHTWEIGHT FUZZY INTEGRATED MODEL FOR TRUSTWORTHY AGV BRAKING

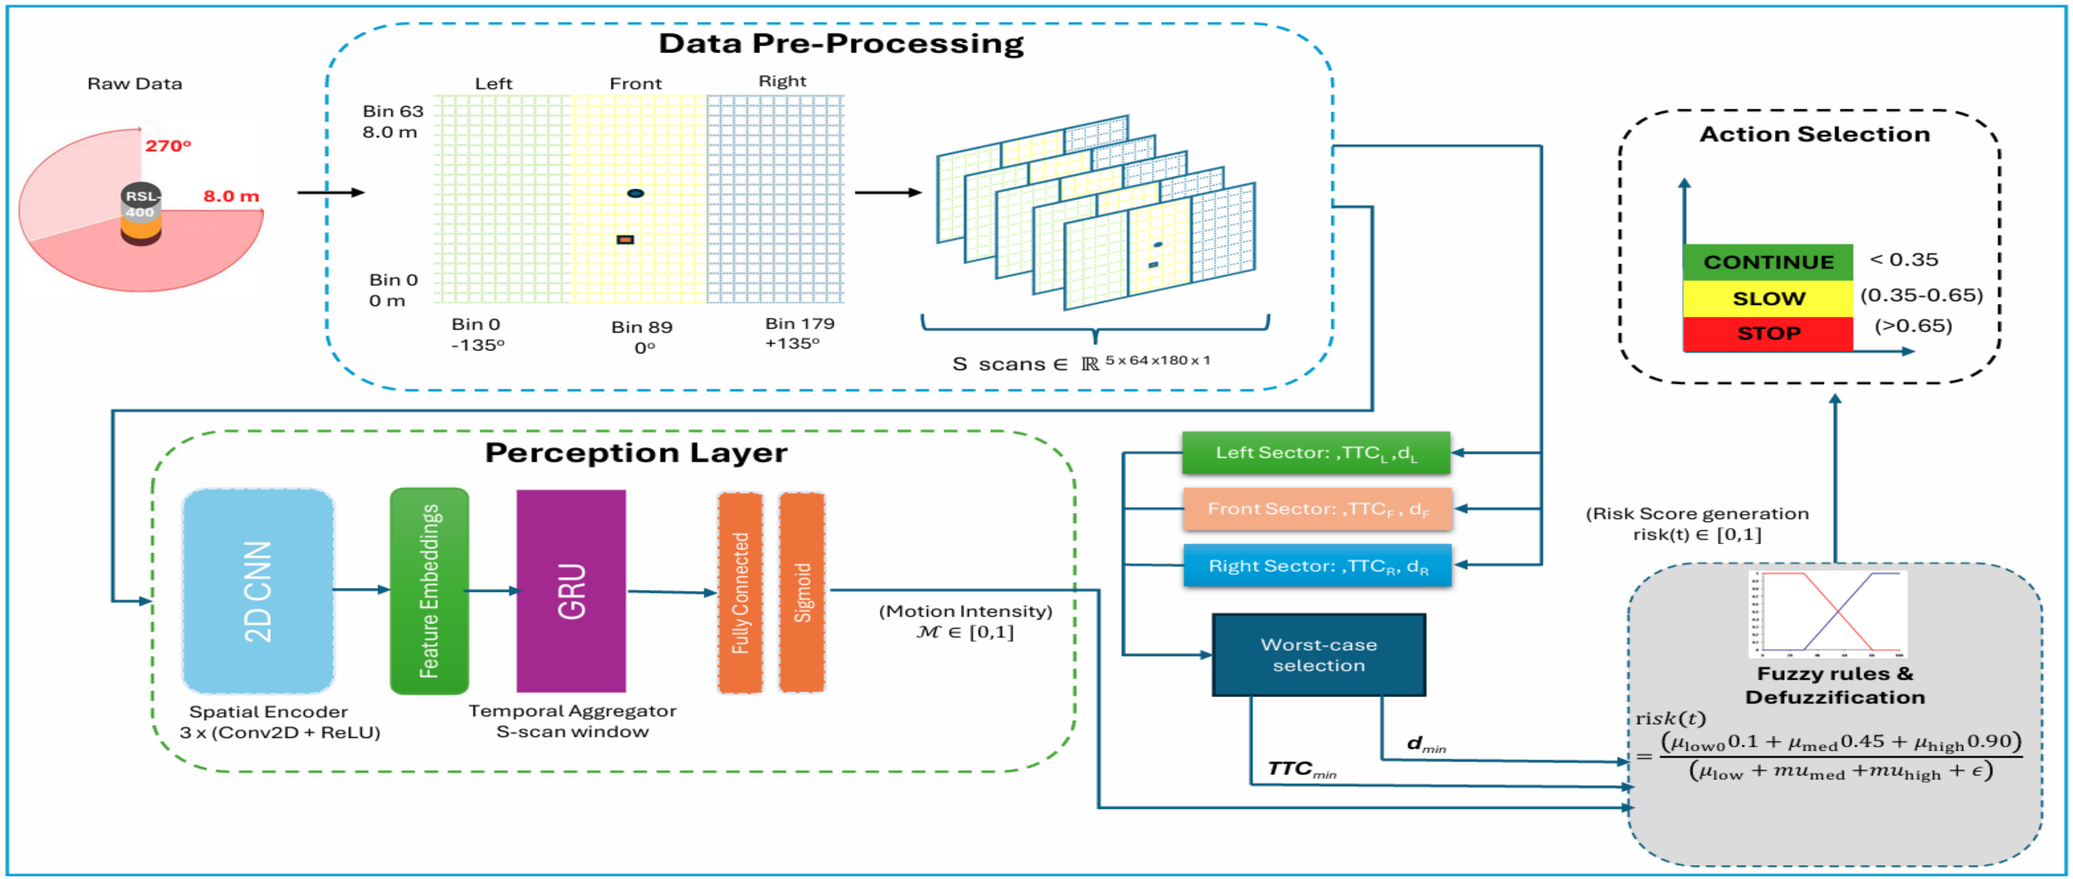


Problem Definition

Limitations of Traditional Rule-Based AGV Braking:

- **Rigid Rule-Based Safety:** Lack of predictive nuance required in dense traffic, leading to abrupt stops.
- **Mechanical Impact**: Frequent reactive braking, causes heavy wear and shortening component lifespan.
- **Efficiency Loss**: Hard stops disrupt the reliable logistics flow in dynamic environments

PROPOSAL SOLUTIONS



A lightweight Edge-AI framework employing a fuzzy logic controller to achieve smooth, proactive deceleration rather than reactive braking

METHODOLOGY

Perception Layer: CNN-GRU Model

2D-CNN Encoder

Extracts essential spatial features and obstacle patterns from the individual polar occupancy grids, avoiding heavy point-based processing.

GRU Temporal

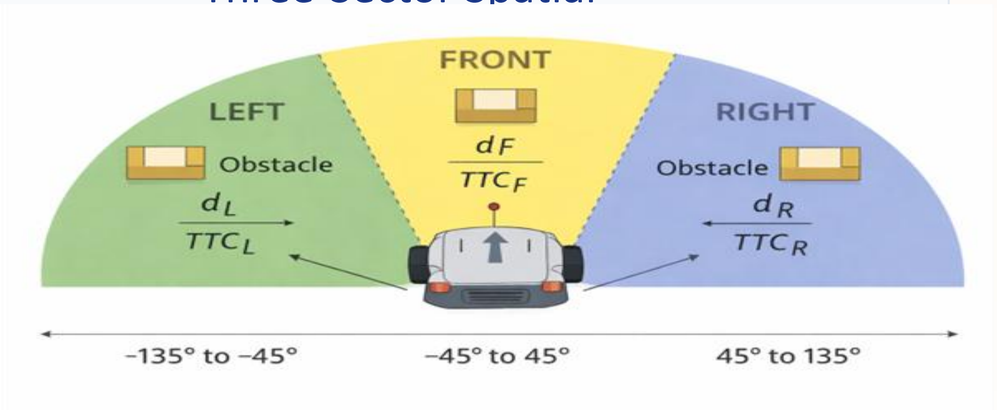
Sequentially processes the temporal vector frames over the S-scan window to effectively capture obstacle motion trajectories.

Motion Intensity

Outputs a continuous scalar $M \in [0, 1]$ for mapping how aggressively obstacles close in.

Geometric safety cues

Three-Sector Spatial



Percentile-based filtering isolates the robust sector distance while mitigating outlier returns.

Worst-Case Kinematic Safeguards

To prioritize safety-critical constraints, the lowest bounds across all sectors are forwarded to the controller:

Minimum Distance: $d_{min} = \min d_s$

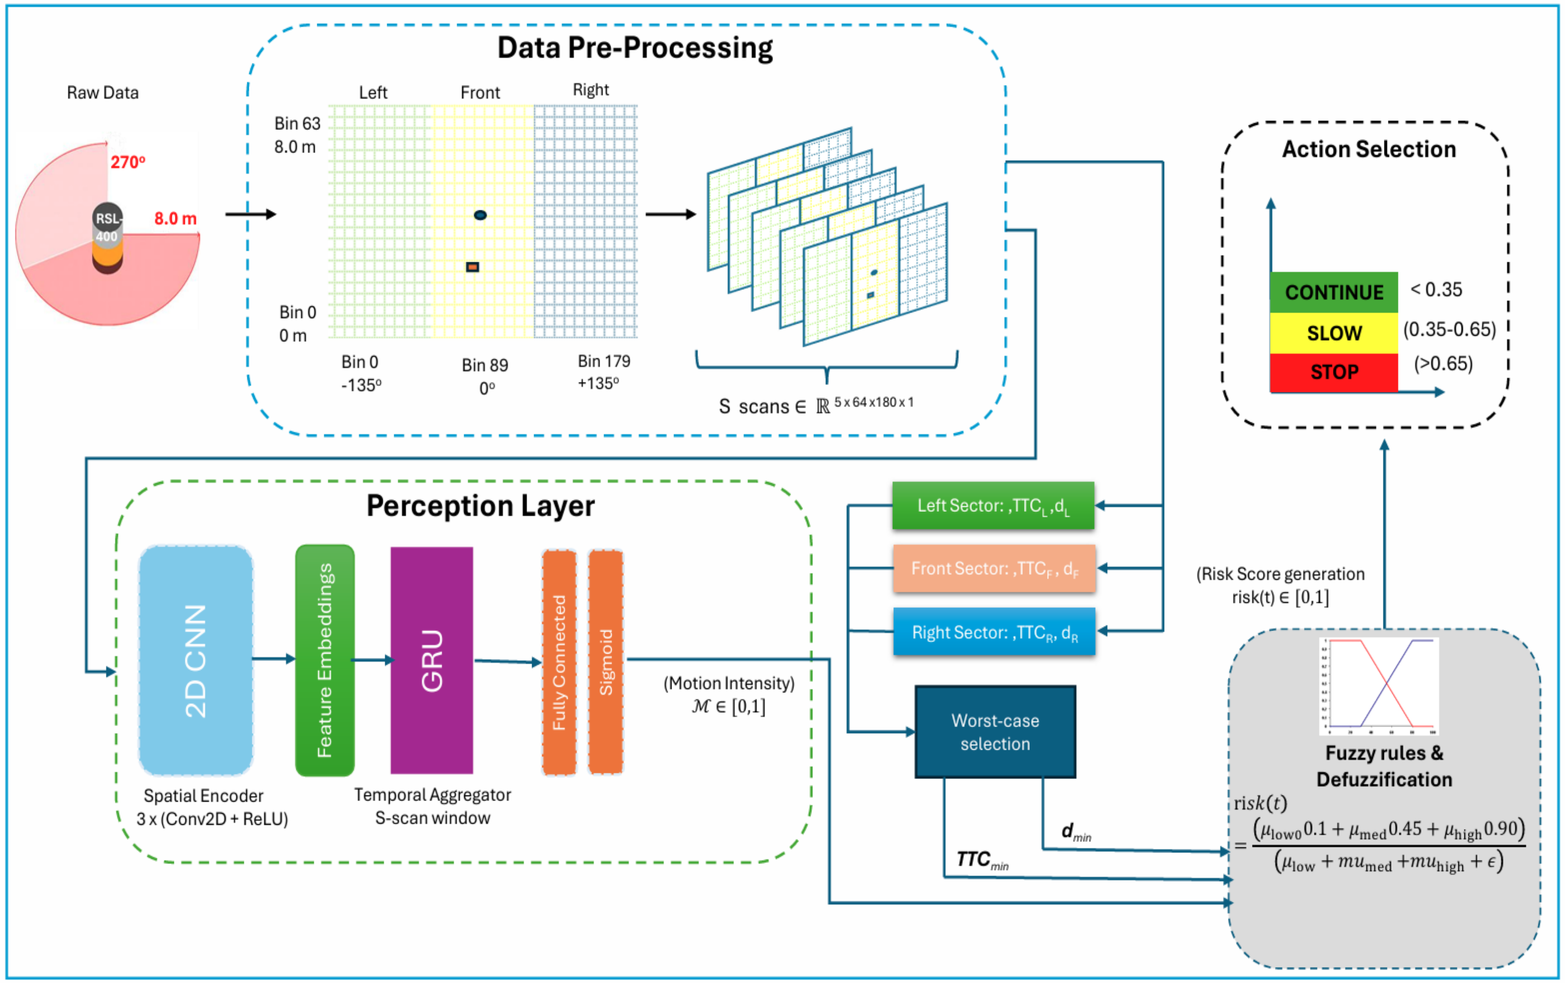
Minimum Time to Collision: $TTC_{min} = \min TTC_s$

FUZZY MEMBERSHIP PARAMETERS

Mamdani trapezoidal membership function's structure the inputs: $\mu(x;a,b,c,d)$

Fuzzy Input	Linguistic Term Set	a	b	c	d
Sector Distance (Normalized)	NEAR	0.00	0.00	0.06	0.09
	MEDIUM	0.07	0.10	0.14	0.18
	FAR	0.16	0.20	0.30	0.30
Time-to-Collision (Seconds)	LOW	0.0	0.0	1.0	2.0
	MEDIUM	1.5	3.0	10.0	20.0
	HIGH	15.0	30.0	10^9	10^9
Motion-Intensity ($M \in [0, 1]$)	LOW	0.00	0.00	0.20	0.40
	MEDIUM	0.30	0.45	0.60	0.75
	HIGH	0.60	0.80	1.00	1.00

METHODOLOGY



Algorithm 1 CNN-GRU with fuzzy risk to action mapping for trustworthy AGV braking

Require: Raw LiDAR $\mathbf{x}(t)$, previous geometry $d_s(t-1)$, sliding window buffer $\mathcal{B}$

Ensure: Motion Action $A \in \{\text{CONTINUE, SLOW, STOP}\}$

$\mathbf{G}_t \leftarrow$ Generate polar occupancy grid (64×180)

$\mathcal{B} \leftarrow$ Update $T = 5$ sliding window with $\mathbf{G}_t$

$\mathcal{M} \leftarrow$ CNN-GRU$(\mathcal{B})$ {Output: Motion Intensity $[0, 1]$}

for all sectors $s \in \{Left, Front, Right\}$ **do**

  $d_s(t) \leftarrow Q_q(\{\tilde{x}_i \mid \tilde{x}_i \in (d_{\min}, d_{\max}]\})$ {Percentile distance}

  $\dot{d}_s(t) \leftarrow \frac{d_s(t) - d_s(t-1)}{\Delta t_{\text{eff}}}$ {Range rate calculation}

  $\text{TTC}_s(t) \leftarrow \begin{cases} \frac{d_s(t)}{-\dot{d}_s(t)} & \text{if } \dot{d}_s(t) < 0 \\ \infty & \text{otherwise} \end{cases}$

end for

$d_{\min} \leftarrow \min_{s \in \mathcal{S}} d_s(t)$ and $\text{TTC}_{\min} \leftarrow \min_{s \in \mathcal{S}} \text{TTC}_s(t)$

$\boldsymbol{\mu} \leftarrow$ Fuzzification of $\{d_{\min}, \text{TTC}_{\min}, \mathcal{M}\}$ {Using Trapezoidal MFs}

$\mathcal{R}_{agg} \leftarrow$ Evaluate Mamdani Rule Base (H1-L2) {Max-Min Aggregation}

$risk \leftarrow \frac{\sum(\mu_i \cdot w_i)}{\sum \mu_i + \epsilon}$ {Rule aggregation and defuzzification}

$A \leftarrow$ Map $risk$ **based on** $\{0.35, 0.65\}$ thresholds

return A

EXPERIMENTAL EVALUATION

PERCEPTION PERFORMANCE

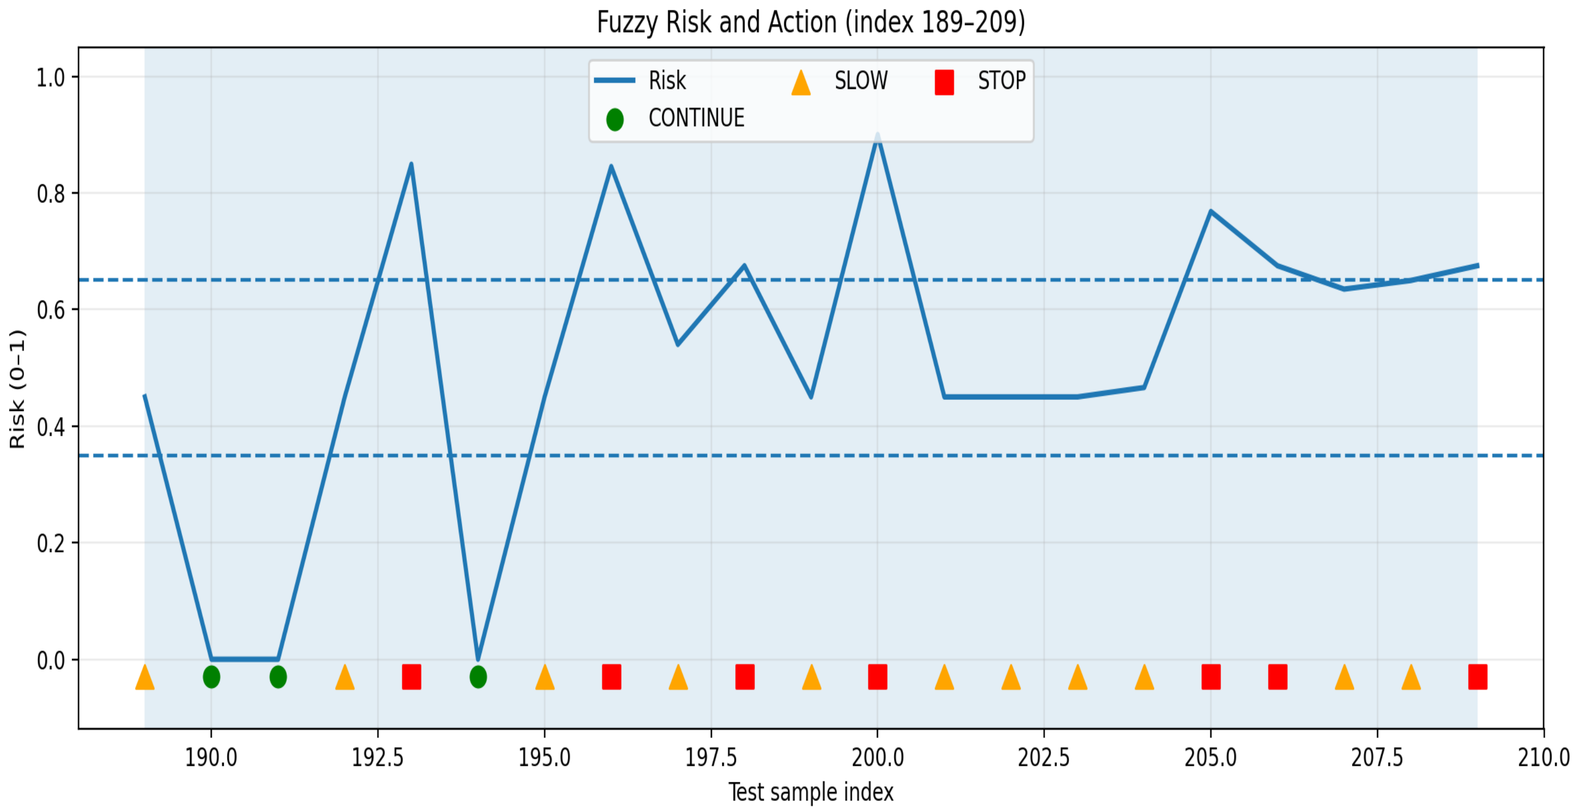
Key Statistics

- **MSE:** Precise regression with final error of 0.00245.
- **Complexity:** 814,595 parameters; optimized for edge deployment.
- **Reliability:** Robust capture of non-linear short-term dynamics.

R^2 Prediction Fidelity

89.15%

CONTINUOUS ACTION SELECTION



Three-Tier Speed Control

CONTINUE: Risk $r < 0.35$

SLOW: $0.35 \leq r < 0.65$

STOP: Risk $r \geq 0.65$

CONCLUSION

Transparency

The Framework provided a human-interpretable rationale for AGV braking decisions.

Tunability

Site-specific thresholds can be adjusted within the FIS to align with warehouse safety policies without retraining neural weights.

Sustainability

Finaly the framework provide a sustainable solution for Proactive deceleration that reduce mechanical jerk, leading to longer AGV operational in Industrial Environment

Future work

The collaborative perception of trajectories of AGV fleets is planned for future work

FLICS

International Conference on Federated Learning and Intelligent Computing Systems

Valencia - Spain

PROGRESS

SCIENTIFIC IMPACT

GPU-Enabled Edge-based Federated Continual Learning with Resource Adaptation for Efficient Non-Stationary Visual Anomaly Detection

Publisher: IEEE | Cite This | PDF

Alexandre Niyomugaba ; Dariusz Mrozek All Authors



RESEARCH-ARTICLE | OPEN ACCESS

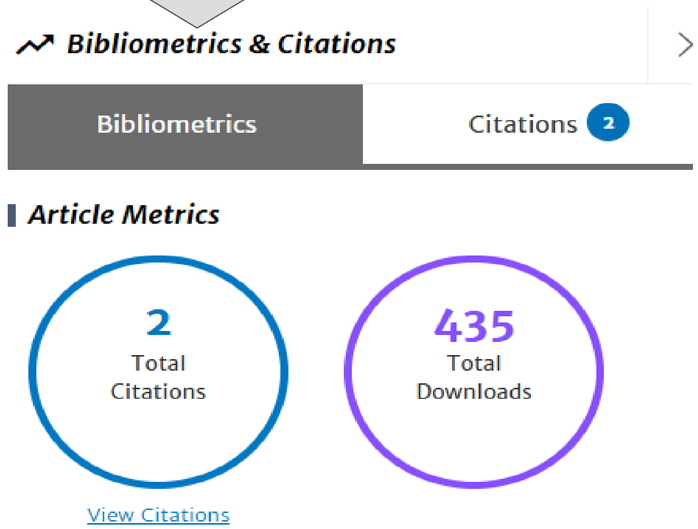
Edge-Optimized Hybrid DNNs for Efficient High-Dimensional Visual Anomaly Detection in Industrial Federated Learning Environments

Authors: Alexandre Niyomugaba, Dariusz Mrozek | Authors Info & Claims

MobiHoc '25: Proceedings of the Twenty-sixth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing
Pages 436 - 443 • https://doi.org/10.1145/3704413.3765514

Published: 23 October 2025 Publication History | Check for updates

2 435 | PDF/eReader



Thank you for your kind attention


+48507222392


Niyomugaba.alexandre@polsl.pl


X
@NIYO_ALEXA

ORCID '0000-0002-8133-5253