UNIVERSITY OF TARTU
Institute of Computer Science

Co-funded by the European Union

Investing in your future

Real-Time & Streaming AutoML for Industry

How AwareML Makes AI Transparent, Controllable & Sustainable for Your Business

Vikash Chander Maheshwari

Prof. Radwa El Shawi

Data Systems Group

Seminar: AI and Human-Centered Technologies in Industry

March 12, 2026 Kohtla-Järve, Estonia

Your Business Runs on Live Data.
Is Your AI Keeping Up?

Machines Fail Unexpectedly

Reactive maintenance costs **3-4x more** than predictive strategies. Downtime kills productivity.

Models Go Stale

Yesterday's AI model doesn't know today's reality. Data drifts, and static models degrade silently.

AI Decisions are a "Black Box"

If managers can't explain *why* an AI made a decision, they won't trust or adopt it.

AUTOML: YOUR AI EXPERT ON AUTOPILOT

Think of it this way:

Traditional ML is like hiring a specialist to manually tune an engine.

AutoML is an intelligent system that does the specialist's job automatically, 24/7.

Algorithm Selection

Automatically tests and finds the best algorithm for your specific data patterns.

Hyperparameter Tuning

Fine-tunes complex model settings that usually require hours of expert time.

Adaptive Rebuilding

Detects changes in data and retrains models instantly to maintain accuracy.

THE WORKFLOW

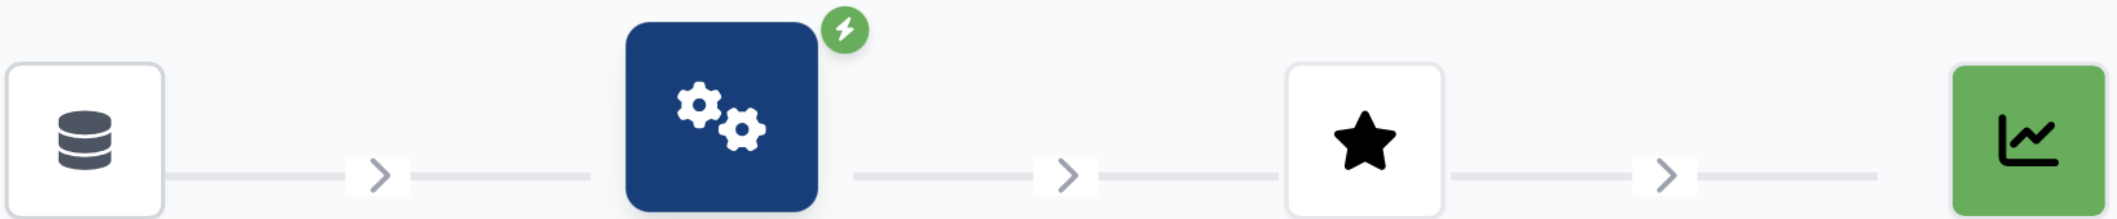

THE CHALLENGE: DATA NEVER STANDS STILL

What is Concept Drift?

In plain words: **customer behavior, production processes, and market environments change over time.**

⚠ **The Problem:** A model trained on January data may fail completely by June if it doesn't adapt.

Energy Demand
Shifts seasonally & daily

Production Defects
Change with raw materials

Customer Behavior
Evolves with market trends

Model Accuracy Over Time

Visualizing the impact of drift and recovery

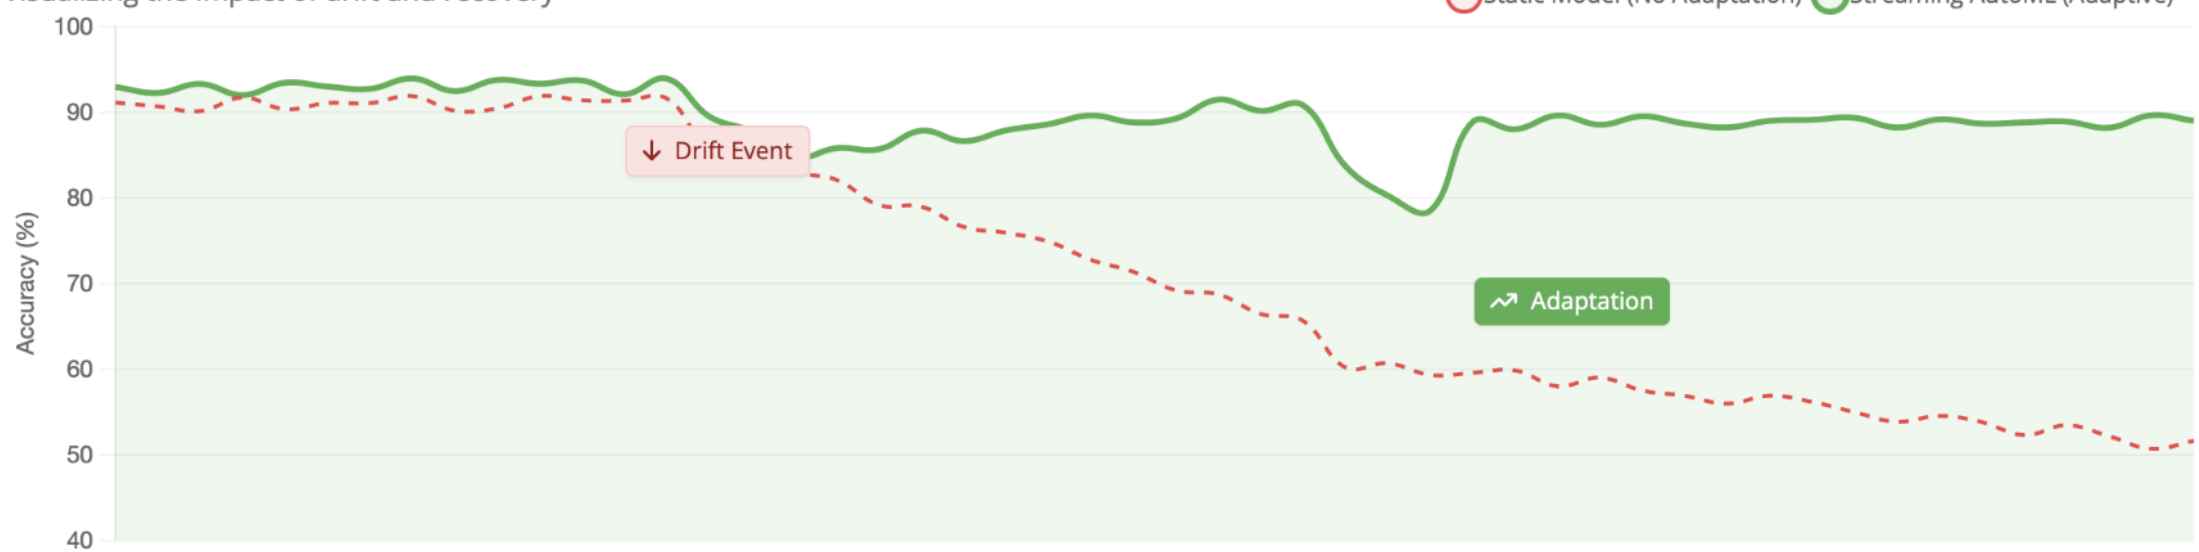

MEET AWAREML

Streaming AutoML You Can Trust and Control

University of Tartu

A unified platform integrating **5 state-of-the-art frameworks** to deliver full visibility and control over your streaming AI.

Controllable

You set the rules. Define algorithms, time budgets, and optimization goals (e.g., accuracy vs. energy) before execution.

Transparent

No more black boxes. See exactly **why** the AI made every decision with integrated SHAP explanations.

Fair

Monitor algorithmic fairness in real-time. Ensure your model treats all demographic groups equally over time.

Sustainable

Track the environmental cost. Measure energy consumption (kWh) and CO_2 emissions for every model run.

https://github.com/DataSystemsGroupUT/AwareML

FROM DATA TO PREDICTIONS

The 6-Stage AwareML Pipeline

A

CONTEXT SETUP

Upload data, set time budgets, and select algorithms. Define your constraints upfront.

B

SMART REC.

AI analyzes your goals (e.g., "low energy") and recommends the best framework.

C

LIVE TRACKING

Watch the model adapt to data drift in real-time. Visualize accuracy drops and recovery.

D

OUTCOME RANKING

Compare frameworks side-by-side. See trade-offs between accuracy, speed, and CO_2.

E

FAIRNESS & XAI

Audit algorithmic fairness (parity) and generate SHAP explanations for decisions.

F

ASK THE AI

Chat with the integrated LLM assistant to get plain-language insights about your data.

This pipeline ensures a structured, compliant, and transparent path from raw streams to business value.

SMART RECOMMENDATION

Let AI Choose the Best Setup for You

Tell AwareML your goal, and it selects the optimal framework.

Mode 1: Natural Language

Type your requirement:

"I need a setup for a low-energy edge deployment with strong performance."

Analyzing...

Mode 2: Metric Selection

Select optimization priorities:

- Energy
- Speed
- CO_2
- Accuracy

RECOMMENDED FRAMEWORK

AutoStreamML

Best match for Low Energy and Sustainability goals.

73.28%	84%
CONFIDENCE SCORE	META-MODEL ACCURACY

CASE STUDY

Predicting Workforce Income in Real-Time

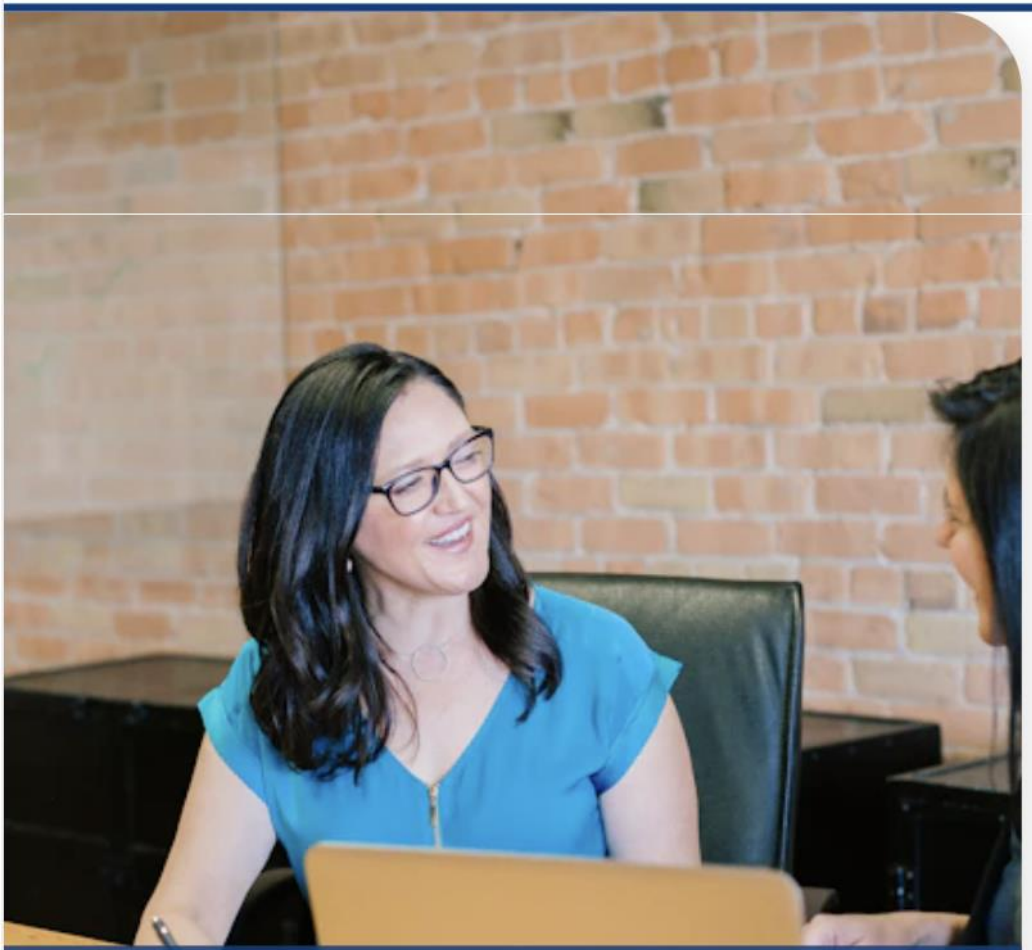

Business Context

An HR analytics platform continuously scores incoming employee profiles. The model must adapt as demographics and job market conditions shift in real-time.

UCI Adult Dataset (Census Income)

48,842	14	Binary
RECORDS	FEATURES	TARGET (>$50K)

Key Features:

Age | Education | Occupation | Hours/Week | Marital Status | Relationship

Experimental Setup

Streaming Configuration

- Time Budget: 60s
- Window Size: 500
- Fairness Attribute: Sex (M/F)

Frameworks Benchmarked

Compared **AutoStreamML** against industry baselines: AutoClass, EvoAutoML, OAML, and ChaCha.

CASE STUDY RESULTS

Live Performance Under Concept Drift

ACCURACY

88.14%

↑ Top-1 vs. Benchmark

TOTAL RUNTIME

770s

Processing ~35k Instances

ENERGY CONS.

0.013 kWh

Ultra Low Edge Ready

CO_2 EMISSIONS

5,352 µg

Minimal Footprint

Accuracy Evolution Over Time (with Drift Events)

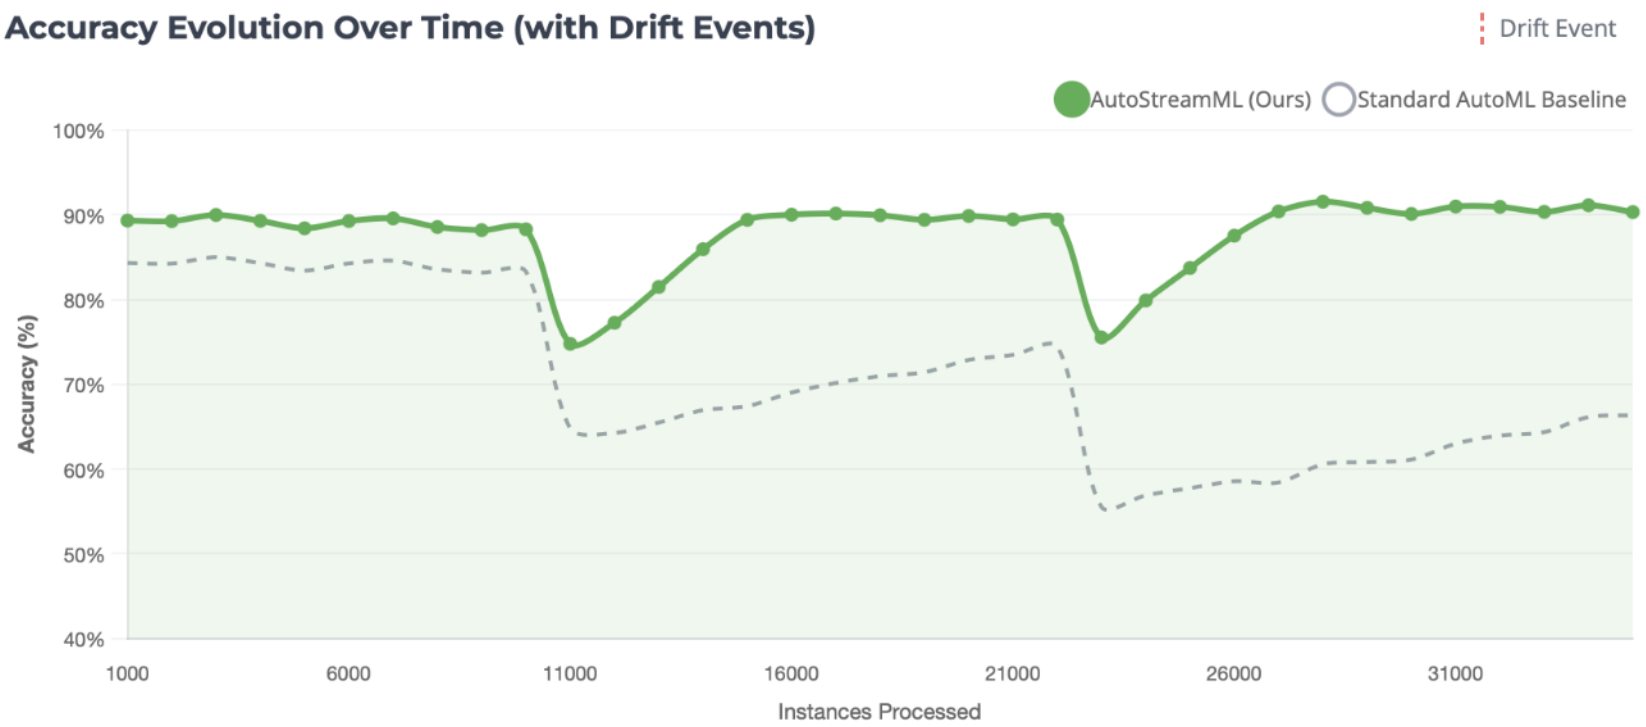

Key Insight

AutoStreamML (Green Line) recovers sharpest after drift events compared to static baselines.

- **Best Accuracy**
- **Lowest Environmental Cost**

DRIFT RECOVERY SPEED

2x Faster

than standard retraining

FAIRNESS MONITORING

Does the AI Treat Everyone Equally?

Metric: Demographic Statistical Parity (DDSP)

We measure whether sensitive groups (e.g., Male vs. Female) receive positive predictions at similar rates. In streaming, we track this **continuously over time** to detect if drift introduces bias.

EU AI Act Ready

Continuous monitoring is essential for compliance with high-risk AI regulations and GDPR non-discrimination requirements.

FRAMEWORK	FAIRNESS STABILITY	BEHAVIOR UNDER CONCEPT DRIFT	RESULT
AutoStreamML	Most Stable	**Maintains consistent parity even after sudden data shifts. Minimal bias fluctuation.**	🏆
AutoClass	Moderate	Shows minor fluctuations in fairness metrics when adapting to new concepts.	–
EvoAutoML	Fluctuates	Evolutionary adaptation sometimes sacrifices fairness for short-term accuracy gains.	–
ChaCha	High Fluctuation	Largest parity swings observed. High risk of temporary bias during model updates.	✕

** Analysis based on 'Sex' attribute (Male/Female) in Adult dataset stream.*

EXPLAINABILITY (SHAP)

Why Did the AI Make That Decision?

What is SHAP?

SHAP assigns a **"blame score"** to each data point. It shows exactly which factors (features) pushed the AI's prediction up or down, making the "Black Box" transparent.

Top 5 Features Driving Predictions

AutoStreamML Model

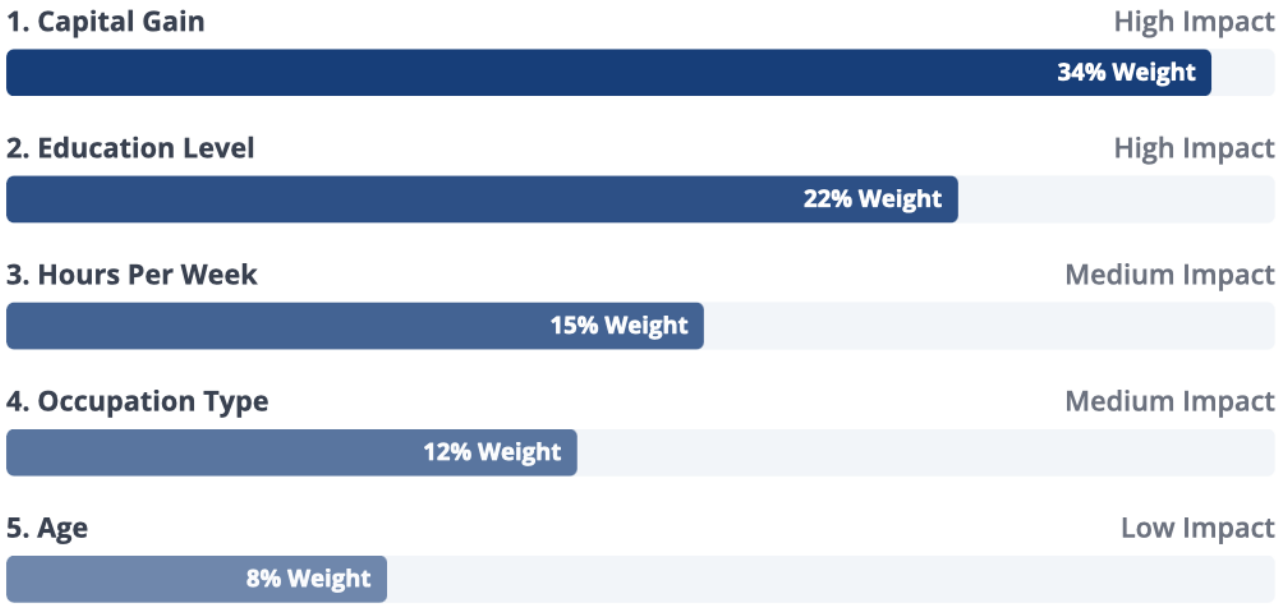

EXPLANATION QUALITY METRICS

STABILITY
✓ High

CONSISTENCY
✓ High

SPARSITY
✓ Good

LLM Insight

"AutoStreamML provides reliable, repeatable explanations driven by a small set of informative features (like Capital Gain and Education) — making it ideal for real-time deployment where trust is critical."

Business Value

You can now EXPLAIN any AI decision to your manager, your auditor, and your customer with confidence.

LLM ASSISTANT

Have a Question About Your AI? Just Ask.

AwareML Assistant
Online

Show me the most important features for AutoStreamML.

Here is the SHAP feature importance analysis for the current AutoStreamML model:

Insight: The top features driving income predictions are **Capital Gain**, **Education Level**, and **Hours Per Week**. These factors account for **67%** of the model's decision weight.

Ask a follow-up question...

Example Questions

"Why was this customer flagged as high risk?"
Get instant explanation for individual predictions

"Has performance dropped in the last 1,000 records?"
Check real-time model health checks

"Which features changed most after the last drift?"
Understand concept drift impact

No PhD Required

AwareML Speaks Your Language

You don't need a data science degree to understand your AI. The built-in LLM translates complex model metrics into clear business insights.

SUSTAINABILITY METRICS

Green AI for Responsible Business

ENERGY CONSUMPTION

0.013 kWh

✓ Extremely Low

CO_2 EMISSIONS

5,352 µg

Minimal

RUNTIME BUDGET

60 sec

Strict Limit

EQUIVALENT TO

~5 mins

Resource Footprint Comparison

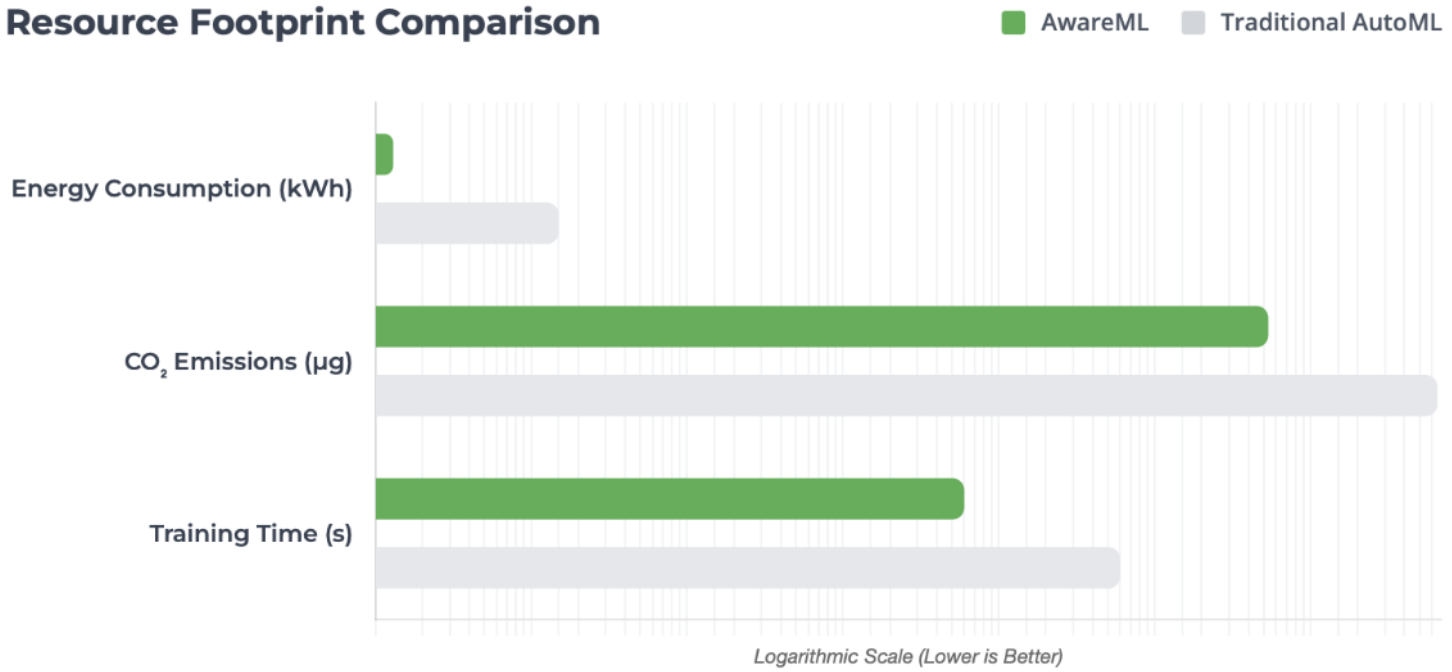

Regulatory Ready

EU Green Deal & CSRD

New directives push companies to measure their digital carbon footprint. AwareML tracks and reports this automatically.

- Audit-Ready Logs
- Transparent Reporting

USABILITY RESULTS

Tested and Validated by Industry Users

Key Findings

n=15 Participants

Ease of Use	Dashboard Navigation	80%
Configuration	System Setup	80%
Visualization	Understanding Drift	80%
Decision Support	Fairness & Sustainability	73%
Confidence	LLM Explanations	73%

STUDY METHODOLOGY

- 15 Industry Participants
- Likert-Scale Evaluation

You don't need to be an AI expert to use AwareML — that's the whole point.

Study Participant
Industry Engineer

BUSINESS VALUE

From Research to Your Factory Floor

INDUSTRY CHALLENGE	AWAREML SOLUTION	BUSINESS BENEFIT
Model Degradation Accuracy drops over time	**Streaming AutoML** Adapts continuously	**Reduced Maintenance** Fewer manual retrains
Black-Box AI Decisions are untrusted	**SHAP + LLM** Plain language insights	**Trust & Adoption** Stakeholder confidence
Fairness Risk Potential legal liability	**DDSP Monitoring** Continuous auditing	**Compliance Ready** EU AI Act & GDPR
Hidden Energy Cost Unknown footprint	**CO_2 & Energy Tracking** Built-in metrics	**Sustainability Reporting** CSRD / ESG Data
Complexity Choosing AI is hard	**Meta-Recommender** AI selects the best setup	**Faster Deployment** Goal-aligned startup

ROI Takeaway

AwareML turns technical complexity into measurable business value.

✓ Risk Reduced

GET STARTED

Try AwareML Today — 3 Simple Steps

1

Upload Data

Upload your streaming dataset in **CSV format**. AwareML works with data from any industry domain (energy, finance, manufacturing).

2

Describe Goal

Use **natural language** ("optimize for low energy") or manually select metrics like Accuracy, Speed, or CO_2 efficiency.

3

Explore Results

View the live dashboard. Monitor performance, audit **fairness**, analyze **explainability**, and track sustainability.

Access AwareML

Open source and ready to demo.

CODE
AwareML

DEMO
AwareML Demo

Let's Pilot Together

We are actively looking for industry partners. We can pilot AwareML on **your company's data** to demonstrate value.

Contact us to start →

LET'S COLLABORATE

Connect with the Research Team

Vikash Chander Maheshwari

PhD Researcher

vikash.chander.maheshwari@ut.ee

Institute of Computer Science

Prof. Radwa El Shawi

Assoc. Professor, Head of Group

radwa.elshawi@ut.ee

Data Systems Group

Data Systems Group

University of Tartu • Institute of Computer Science

bigdata.cs.ut.ee

WAYS WE CAN WORK TOGETHER

Joint R&D Projects

Collaborative research grants and innovation partnerships

Workshops & Training

Upskill your engineering team on AutoML & AI Ethics

Pilot Project

Test AwareML on your company's data today

Thank You

Questions?

We'd love to hear from you.

PRESENTED BY

University of Tartu
Institute of Computer Science

Seminar
AI and Human-Centered Technologies in Industry

AwareML

Transparent, Controllable,
and **Sustainable** AI for Industry

Explore the Code
https://github.com/DataSystemsGroupUT/AwareML

Watch the Demo
https://streamable.com/nm5mkv