



Case Study:

Revolutionizing Industrial Laundry Operations with Edge Impulse and Edge AI



EDGE IMPULSE

A Qualcomm Company

Executive Summary

In the high-volume, low-margin commercial laundry industry, operational efficiency and uncompromising quality control are the primary drivers of profitability. Facilities processing hundreds of thousands of pounds of textiles daily for premium hospitality and healthcare clients face strict Service Level Agreements (SLAs). A single stained bedsheet reaching a luxury hotel can result in severe financial penalties and damaged client relationships.

This case study examines how a major industrial laundry provider successfully transformed its sorting and quality assurance processes by pivoting from a failed traditional, rule-based machine vision initiative to a highly adaptable Edge AI solution. By leveraging the Edge Impulse platform and deploying the proprietary FOMO (Faster Objects, More Objects) computer vision algorithm on production optimized hardware, the facility achieved unprecedented results.

The business impact was immediate and profound: the company was able to completely reallocate **100% of their Full-Time Equivalent (FTE) manual-sorting workforce to higher-value roles**. Furthermore, they minimized client returns due to textile defects and stains. Most remarkably, after a previous 18-month project using fixed, proprietary machine vision tools stalled completely, the engineering team utilized Edge Impulse to deploy a production-ready, highly accurate system in **less than three months — a 6x acceleration in time-to-value**.



With Edge Impulse it just works. We were able to accomplish what we were NEVER able to do with the previous vendor in 18 months.



— Operations and Customer Service Manager, Laundry Industry

The Business Context:

The Industrial Laundry Bottleneck

Commercial laundry is a relentless logistical challenge. The modern facility is a symphony of massive tunnel washers, automated dryers, and high-speed conveyor belts. However, the critical junction of sorting and quality control has historically remained a manual, human-driven bottleneck.

As damp textiles exit the pressing machines, they travel down conveyor belts at speeds exceeding 1.5 meters per second. At this stage, two critical decisions must be made in fractions of a second:

- 1. Routing:** What type of textile is this? (e.g., king sheet, queen sheet, bath towel, pillowcase).
- 2. Quality Assurance (QA):** Does this item contain a defect? (e.g., makeup stains, grease, blood, or physical tears).



High-throughput sorting conveyors present a massive bottleneck for manual inspection.

Historically, a team of four manual operators per shift stood beside these fast-moving belts. Their job was to visually scan the passing white sheets under harsh industrial lighting and manually pull stained or defective items off the line, while physically routing the clean items into their respective finishing bins.

This manual approach presented severe operational vulnerabilities:

- 1. Cognitive Fatigue:** Staring at a continuous stream of white textiles under bright lights causes rapid eye fatigue, leading to a steep drop in accuracy after just 45 minutes of a shift.
- 2. Labor Costs and Turnover:** The work is physically demanding and repetitive, resulting in high employee turnover and immense training overhead.
- 3. The Cost of Escapes:** Human error is inevitable. When a stained sheet "escapes" the facility and is unfolded in a luxury hotel room, the laundry provider incurs immediate financial penalties and risks losing multimillion-dollar contracts.

A Previous Approach: The Limits of Traditional Fixed Machine Vision

Prior to working with Edge Impulse, the facility's executive team initially approved a highly ambitious, high-cost system to automate the inspection using industry-standard, fixed machine vision systems. Unfortunately, that project's rigid architecture, slow pace of deployment, and overall uncertain viability doomed it from the start.

The initial automation strategy relied on traditional, deterministic machine vision — the kind of rule-based systems (similar to standard solutions commonly offered commercially) that have dominated manufacturing automation for decades. The facility installed expensive, proprietary smart cameras over the belts and attempted to program explicit logic to find defects using standard tools like **Blob Analysis** and **Contrast Thresholding**.

After 18 months of development and hundreds of thousands of dollars in sunk costs, the project was completely scrapped. The failure was driven by three insurmountable technical barriers inherent to rule-based vision when applied to organic, natural variability.

1. The Shadow Dilemma: Folds vs. Stains

Traditional machine vision relies heavily on pixel intensity and contrast. These systems excel at measuring the rigid, machined edges of a metal automotive part or reading a perfectly flat barcode.

Textiles, however, are non-rigid, amorphous, and constantly deforming. When a king-sized sheet is ejected onto a high-speed belt, it bunches, folds, and overlaps. Under harsh factory lighting, deep folds create dark shadows. To a deterministic contrast threshold tool, a shadow across white cotton has the exact same pixel intensity drop as a grease stain. The system triggered false positives constantly, ejecting perfectly clean sheets simply because they were wrinkled.

2. The Shape Problem: Blob Analysis Failures

The previous provider's engineering team then attempted to use Blob Analysis to filter out the shadows. Blob tools look for continuous regions of foreground pixels and measure their properties (aspect ratio, perimeter, circularity). A "good" part in traditional manufacturing produces blobs within expected geometric ranges.

Stains on bunched sheets have absolutely no defined geometry. A makeup smear might be stretched out, while a chemical burn might be a perfect circle. Furthermore, because the textiles were constantly morphing, the background against which the blobs were measured was entirely unpredictable. The explicit logic required to account for every permutation of a stain's shape was impossible to code.

3. Vendor Lock-In and the Proprietary Deep Learning Wall

Realizing that rule-based tools could not handle the natural variability of textiles, the laundry team evaluated upgrading to the vendor's proprietary deep learning suite.

The financial and architectural reality of this upgrade was staggering. Utilizing proprietary deep learning vision platforms in this ecosystem required:

- **Heavy Hardware Overhead:** Each inspection station required a localized Windows PC equipped with expensive NVIDIA GPUs (10GB+ VRAM) to run inference.
- **Aggressive Licensing:** Software training and runtime licenses cost thousands of dollars annually per line, completely destroying the unit economics of the laundry facility.
- **Physical Constraints:** The software required physical USB security dongles on every machine on the factory floor, creating a severe hardware failure point and eliminating any possibility of a modern CI/CD software deployment model.

The traditional fixed solution proved to be fundamentally misaligned with the economic and physical realities of the facility. The team needed a solution that offered the semantic understanding of deep learning, but with the lightweight, hardware-agnostic flexibility of modern software engineering.

The Paradigm Shift to Edge AI: Enter Edge Impulse and FOMO

The engineering leadership, recognizing the need for a modern MLOps architecture, pivoted entirely. They abandoned proprietary smart cameras and heavy PC dependencies in favor of **edge AI**.

They turned to Edge Impulse to orchestrate the pipeline, leveraging a novel object detection architecture pioneered by the company known as **FOMO (Faster Objects, More Objects)**.



The FOMO algorithm can detect stains of various sizes and intensity, while avoiding false positives from shadows.

Unpacking the FOMO Architecture

Standard deep learning object detection networks (like MobileNet SSD or YOLO) are computationally heavy because they perform complex regression to determine the exact width, height, and bounding box coordinates of an object.

FOMO radically reimagines this process. It recognizes that for many industrial use cases — like finding a stain on a sheet or counting items on a belt — the exact dimensions of the bounding box are irrelevant. The PLC only needs to know where the defect is to actuate the diverter arm. Instead of outputting heavy bounding boxes, FOMO outputs a highly efficient spatial probability heatmap. It finds the **centroids** of target objects.

From a mathematical standpoint, FOMO drops the highly complex, fully connected layers found in traditional networks. It takes an input tensor (e.g., a 224 x 224 x 3 RGB image), processes it through a highly truncated backbone, and outputs a reduced spatial grid of 28 x 28 x C, where C is the number of classes (e.g., "stain", "tear"). Each cell in this grid represents a localized probability that an object's center exists within that spatial region.

This mathematical reduction provided three massive advantages over the previous traditional vision approach:

- **Semantic Understanding over Rigid Rules:** Unlike blob tools that rely on explicit geometry and pixel thresholds, the neural network backbone of FOMO learns what a stain actually looks like (the gradient, the texture) versus a cotton fold. It completely eliminated the "shadow vs. stain" false positive issue.
- **Micro-Computational Footprint:** FOMO requires up to 30x less processing power than standard object detection. It completely eliminated the need for heavy Windows PCs and NVIDIA GPUs. The model could run inference in under 10 milliseconds on a standard, low-cost Linux edge node or even a Cortex-M microcontroller.
- **Hardware Agnosticism:** Edge Impulse compiles the final model as an open C++ library. The facility was no longer locked into proprietary vendor hardware or USB licensing dongles. They could run the model on generic industrial hardware sourced globally.

After previewing FOMO's capabilities, it was clear to the laundry provider that it would solve the issues they had been experiencing.

Engineering the Solution: The Edge Impulse Workflow

With the right architecture identified, the team executed a highly streamlined workflow, treating the machine vision challenge as a modern software deployment rather than a rigid hardware installation.

1. Hardware Abstraction and Data Ingestion

Breaking Vendor Lock-In

The facility installed rugged, off-the-shelf global-shutter machine vision cameras directly above the sorting belts. These were tethered via GigE or USB 3.0 to generic, lightweight industrial PCs (Linux edge nodes). Using Edge Impulse's robust API, the team continuously ingested thousands of raw images directly from the active sorting lines into Edge Impulse Studio, to build a robust dataset based on the actual hardware and environment.



Off-the-shelf edge AI cameras were mounted directly over the sorting lines for localized processing.

2. Rapid Annotation using AI-assisted labeling pipelines

Accelerated Labeling

The data labeling process was exponentially faster than traditional bounding-box annotation by leveraging Edge Impulse's AI-assisting labeling pipelines. Instead of painstakingly attempting to draw boxes around highly irregular, amorphous stains, operators were able to upload new images to an Edge Impulse upload portal and the labeling pipeline took care of auto labeling incoming new data from the field. This allowed the data engineering team to build a massive, highly varied dataset in days rather than weeks.

3. Digital Signal Processing (DSP) and Training

Automated Augmentation

Within Edge Impulse Studio, the team constructed a custom impulse that routed input images through an image DSP block to normalized lighting variations and downsampled the resolution. Because Edge Impulse handles complex data augmentation under the hood — randomly rotating, scaling, and cropping images to simulate different textile folds — the FOMO network learned to generalize perfectly, achieving an F1-score of over 98%.

4. Direct Edge Deployment and PLC Integration

Real-Time Actuation

The optimized model, weighing less than 500 KB, was compiled directly from Edge Impulse as a highly optimized C++ library and was deployed to the edge nodes. When the system then detected a stain, the edge node instantly fired a Modbus TCP signal to the local Programmable Logic Controller (PLC). The PLC then triggered an I/O signal to reject the defective textiles.

The 3-Month Sprint:

6x Faster Time-to-Value

The contrast between the rigid, 18-month failure of the traditional machine vision project and the flexible precision of the Edge Impulse deployment was staggering. By utilizing a platform built specifically for the realities of edge AI and the highly efficient FOMO algorithm, the engineering team deployed a production-ready solution in a single fiscal quarter.

Month 1: Infrastructure Setup & Data Engineering

Days 1-30

- Installed localized global-shutter cameras and edge computing nodes. Began continuous data ingestion directly into Edge Impulse Studio. Utilized AI-assisted labeling pipelines to rapidly annotate images of textiles, capturing a vast array of folds, shadows, organic stains, and chemical burns under active plant lighting conditions.

Month 2: Model Architecture & Iterative Training

Days 31-60

- Configured the DSP and FOMO neural network blocks. Leveraged Edge Impulse's data augmentation to artificially expand the dataset and prevent overfitting. Trained and validated the model, optimizing for maximum recall (ensuring no stains escaped) while maintaining a latency footprint capable of sub-10 millisecond inference.

Month 3: PLC Integration & Production Go-Live

Days 61-90

- Compiled the optimized C++ library from Edge Impulse and deployed it directly to the generic edge nodes. Mapped the inference output to PLC logic to actuate the pneumatic sorting arms. Completed User Acceptance Testing (UAT) and fully automated the sorting line.

The Business Impact and ROI

The deployment of the Edge Impulse FOMO solution fundamentally transformed the facility's unit economics and operational reliability, establishing a new framework for how they approach automation.

1. 100% FTE Savings for Sorting Operations

Prior to the edge deployment, the facility required 12 full-time employees (four operators across three shifts) dedicated exclusively to manually sorting and inspecting textiles on this specific line.

Upon successful integration of the automated pneumatic diverters guided by the FOMO model, the facility achieved a 100% FTE reduction in manual sorting. In a notoriously tight labor market, these workers were seamlessly reallocated to higher-value, less physically taxing roles within the facility, such as complex textile repair and dispatch logistics. The immediate labor cost savings generated a positive ROI on the generic hardware investment in less than five months.

2. Minimizing Defect Returns and SLA Penalties

Human inspectors, plagued by eye fatigue, frequently missed subtle defects, particularly pale makeup stains on white sheets.

Furthermore, the previous traditional machine vision system caused massive bottlenecks by ejecting perfectly clean, wrinkled sheets.

The FOMO algorithm solved both. Operating at a consistent sub-10 millisecond inference speed, the system accurately identified defects as small as a dime while the textiles passed at 1.5 meters per second, entirely ignoring the shadows caused by folds. Within the first month of operation, the facility noted a near-total elimination of stained sheets reaching their premium hospitality clients. This drastically reduced return-shipping logistics costs, eliminated expensive SLA breach penalties, and significantly strengthened client trust.

3. Rapid, Cost-Effective Scalability

Because the Edge Impulse solution compiles into standard C++ and does not rely on expensive per-seat software licensing or heavy GPU environments, the marginal cost of scaling is exceptionally low. To automate the next sorting line, the facility simply purchases another off-the-shelf camera and a \$300 industrial PC, clones the compiled library, and mounts the hardware. This decoupled, vendor-agnostic architecture allows for infinite horizontal scaling across their remaining facilities nationwide.

Conclusion

The transition from deterministic, traditional machine vision to edge AI is not just a technology upgrade; it is a fundamental shift in how complex physical automation is approached.

The previous 18-month failure proved that high-speed industrial environments handling natural, deformable materials cannot be solved with explicit geometric rules, pixel contrast thresholds, or rigidly licensed proprietary hardware and software.

By leveraging the Edge Impulse platform and the specialized architecture of the FOMO algorithm, the engineering team bypassed these archaic limitations entirely. They delivered a solution that not only achieved 100% labor automation for sorting and vastly improved product quality, but did so six times faster than traditional methods, establishing a blueprint for modern industrial machine vision.

Further Reading

FOMO is one of several computer vision algorithms available in the Edge Impulse platform. Edge Impulse also offers its own state-of-the-art YOLO algorithm, **YOLO-Pro**, for use cases that require traditional bounding-box outputs, as well as **FOMO-AD**, which is designed to identify visual anomalies using only “normal” training data.

Together, FOMO, FOMO-AD, and YOLO-Pro provide **a versatile suite of computer vision algorithms for addressing sophisticated business challenges and helping developers build solutions for demanding vision use cases.**

Beyond computer vision, Edge Impulse supports machine learning pipelines for audio and time-series sensor data, enabling applications such as keyword spotting, sound and motion classification, sensor fusion, regression, and anomaly detection.

Learn more about Edge Impulse at www.edgeimpulse.com, explore additional **case studies**, or create a free developer account at studio.edgeimpulse.com.



