AI-Driven Quality Control Using the Falcon Digital Twin Platform

Summary

Modern manufacturing faces the ongoing challenge of maintaining high product quality while reducing waste and improving efficiency. As production lines grow more complex and faster, traditional quality control methods—particularly manual inspection and sampling—struggle to keep pace. Human fatigue, subjectivity, and limited inspection coverage create bottlenecks that lead to undetected defects, costly rework, and diminished confidence in product reliability.

Duality AI has developed an AI-Driven Quality Control (AIDQ) framework that automates defect detection and classification across manufacturing processes. Using advanced computer vision and machine learning, AIDQ systems continuously monitor every unit in production with greater speed, consistency, and accuracy than human inspectors. Real-time feedback enables proactive process adjustments, reducing scrap, improving throughput, and ensuring consistent product reliability. While other vision systems in manufacturing exist, the datasets required to train them often must be obtained from the production lines themselves (e.g. collecting real world images of defective pills, in all varieties and lighting conditions that can be encountered in production). These datasets are expensive to collect and interfere with production, leading to increased costs.

To combat this and accelerate adoption of AIDQ, Duality leverages its high-fidelity digital twin platform, Falcon, to generate synthetic training data optimized for training AI vision systems. Falcon simulates real-world manufacturing environments with photorealistic accuracy, utilizing an extensive suite of virtual sensors to create labeled datasets that capture diverse scenarios, materials, and lighting conditions without interrupting live operations.

This approach overcomes one of the largest barriers to AI deployment, limited data, by enabling rapid, scalable, and controllable dataset generation tailored to specific products, components, and defect types. Moreover, this process can be applied to any AI vision model, running on any hardware system. Together, AIDQ and Falcon deliver an industry-proven, cost-effective solution for next-generation manufacturing assurance.

Technology Overview

The proposed solution to develop AIDQ for manufacturing and sustainment relies on synthetic data generated from digital twin simulations. Duality's custom digital twin platform, Falcon, enables realistic simulation of production lines, inspection systems, and sensor outputs to generate AI-ready datasets that accelerate development and validation of manufacturing detection AI models. This technology directly supports automated inspection, process optimization, and adaptive manufacturing control which are critical priorities in both defense and industrial applications.

Falcon Platform Architecture

The Falcon platform consists of two integrated components:

- FalconEditor**:** A real-time 3D environment built on Unreal Engine for designing digital twins. Engineers can import CAD models, define materials, lighting, and motion, and instrument scenes with virtual sensors that replicate real-world optics and noise characteristics.
- FalconSim**:** A lightweight simulator that enables automated, large scale data generation

Falcon was built specifically to generate sensor feeds that match any real world sensors found on a production line. Furthermore, all output data is automatically labeled with any information relevant to training the QA/QC AI vision model — a significant saving in cost and time.

Through procedural scene variation, Falcon can generate thousands of unique scenarios by adjusting lighting, materials, camera position, and defect attributes. This accelerates model retraining and validation without the cost and disruption of collecting new real imagery.

Manufacturing Applications

While current implementations of vision systems exist, these methods rely heavily on data collected on the production floor. Subsequently the systems are often fragile—with minor changes in lighting, camera angle, or material properties causing false detections and misclassifications.

Falcon mitigates this challenge by simulating the full range of environmental and process variability.Falcon allows engineers to generate datasets that include glare, shadow, motion blur, reflective surfaces, and variable illumination levels. For example, a component can be rendered across different times of day, factory lighting setups, or camera placements.

Whenever real-world production lines need to be reconfigured, upgraded, or set up for a new product, new synthetic data can be generated in advance, maintaining AI model performance through the transition. Furthermore, new data can be rapidly generated at any point to address any discovered QA/QC misses. This adaptability enables AI systems to remain robust under constant change, avoiding downtime and costly real-world data recollection.

Falcon's synthetic data generation capabilities enable AI-driven inspection systems across multiple dimensions of quality assurance and process control. These applications are central to the AIDQ framework and illustrate Falcon's scalability and versatility across manufacturing domains.

1.Item and Component Detection

Reliable item detection is foundational to manufacturing automation, ensuring that each component is correctly identified and tracked as it moves through an assembly line. Alongside general object detection more complex assemblies require recognition not only of finished products but also of subcomponents and their spatial relationships. Obtaining effective data for detection training requires dimensionally accurate models. Falcon allows direct import of CAD assemblies, automatically generating object-level metadata, segmentation masks, and material identifiers for every part.

Engineers can simulate partial assemblies, occluded parts, and orientation variations, producing highly accurate ground truth labels for component detection and categorization. Such datasets are critical for verifying multi-stage assembly processes and ensuring downstream components meet precise mechanical and visual standards.

2. Counting and Verification

Counting and verification tasks are essential for maintaining process control and ensuring assembly completeness. Falcon simulates full conveyor or robotic workflows, capturing the sequential motion of parts with deterministic precision. Each frame is traceable through time, allowing AI systems to learn temporal correlations and verify production sequence integrity.

This enables the development of vision models that can track parts across multiple stations, detect missing items in a sequence, or confirm correct batch quantities. Falcon's deterministic simulations can reproduce challenging real-world phenomena including motion blur, vibration, inconsistent spacing, allowing robust training without interrupting live operations.

By rapidly generating synthetic datasets for these scenarios, manufacturers can validate AI models for throughput counting, batch verification, and visual auditing without halting production.

3. Defect Detection

Defect detection is one of the most transformative AIDQ applications. Real defect data is rare, inconsistent, and expensive to collect. Falcon overcomes this by procedurally introducing realistic surface and structural defects directly onto canonical digital twins. Engineers can control variables such as severity, frequency, and placement to generate balanced datasets that include both normal and defective samples.

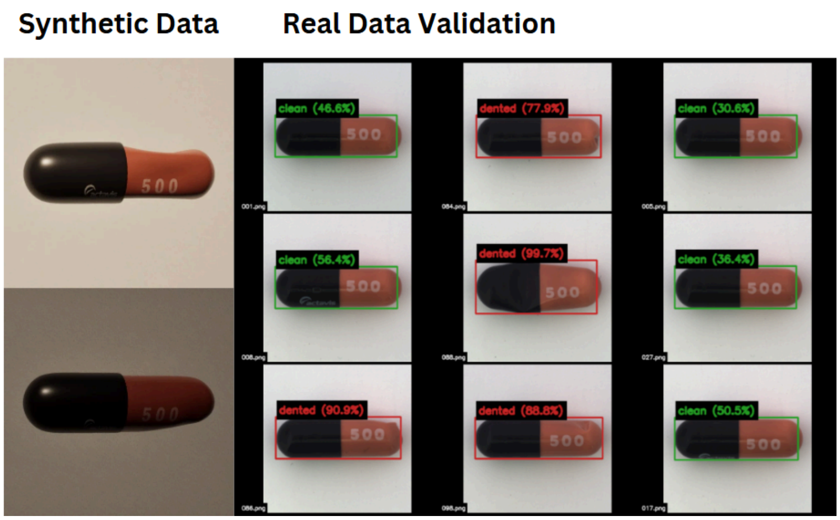

Fig 1. Results of model trained on synthetic data

Defects can include (though are not limited to):

- **Surface defects:** scratches, abrasions, or corrosion
- **Structural defects:** dents, cracks, or misalignments
- **Print and marking defects:** faded labels or misprints

Each defect is rendered with physically based materials that accurately capture texture, reflection, and depth. Falcon's parameterized workflow allows regeneration of defect datasets for new defect classes or product variants in hours rather than weeks. This enables continuous model improvement and deployment of highly granular inspection systems that move beyond simple pass/fail classification to localized defect mapping and root cause analysis.

Approach and Path to Maturity

Synthetic data enables systematic exploration of manufacturing variability at scale. Falcon's simulation pipeline, combined with the USD Digital Twin Encapsulation Standard, ensures every simulated scene encodes material, geometry, and process metadata for traceability and reproducibility. Virtual sensors emulate real-world hardware, while Falcon's Python API enables scripted parameter sweeps and automated edge-case generation.

Path to maturity:

1. Integrate Falcon synthetic data into any QA workflows.
2. Validate AI models using real-world inspection footage.
3. Iteratively optimize simulation realism using the Indistinguishability Score (detailed in the following section).

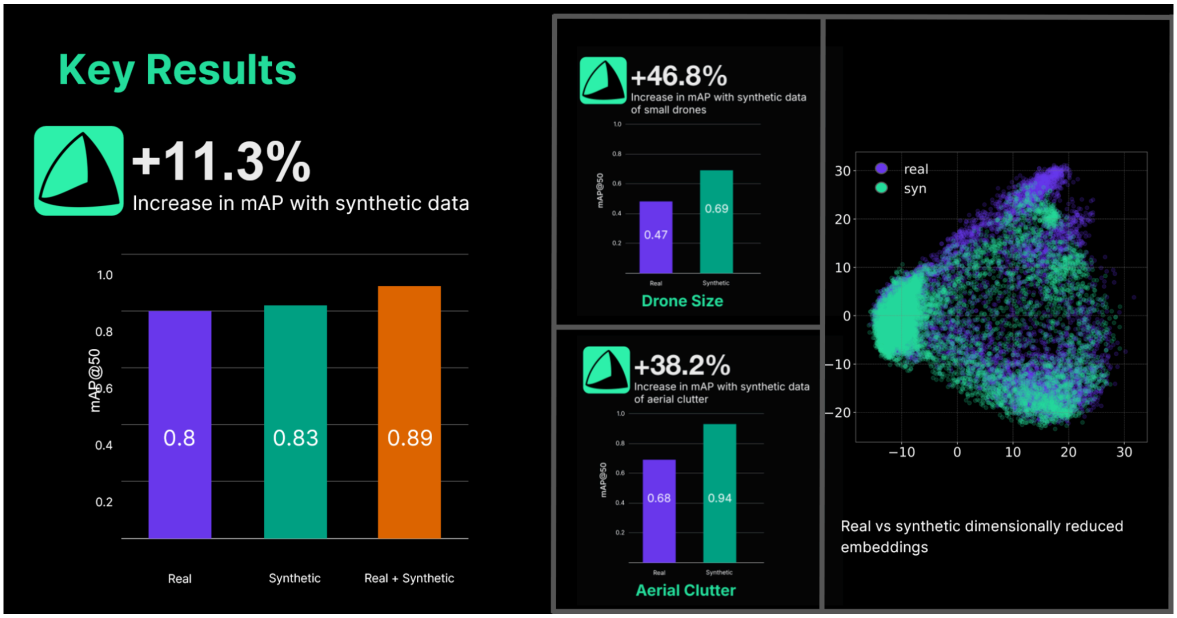

Fig 2. Results of training using synthetic data

Measuring Realism: The Indistinguishability Score

To ensure synthetic data matches real-world conditions, Duality employs the Indistinguishability Score, a quantitative metric that measures similarity between real and synthetic data distributions. The score compares feature embeddings from both datasets; higher overlap indicates higher realism. Scores above 0.8 have been correlated with significant improvements in model generalization. This feedback loop enables Falcon users to systematically increase realism until synthetic data achieves production-level reliability.

Validation and Impact

Falcon's synthetic data pipeline has been validated across multiple domains:

- **U.S. Army:** Drone detection models trained with Falcon data significantly outperformed those trained solely on real imagery.[2]
- **NASA-JPL:** Computer vision foundation models fine-tuned with Falcon synthetic data showed improved off-road perception robustness.[1]

Falcon's synthetic data pipeline has been validated across multiple domains, including defense, aerospace, and manufacturing. In high-volume manufacturing environments, Falcon has demonstrated transformative results. As stated by Procter & Gamble:

> "Falcon is a cost-effective tool to create synthetic images critical in building models that perform to the very high accuracy requirements for manufacturing. Precisely controlling features in synthetic data sets is key to success and is simply not practical with real data

-- we could not afford to physically generate every kind of possible defect with real materials and line time.

With the synthetic data created with Falcon, we have access to data sets needed for higher accuracy and improved robustness. In high volume manufacturing, the accuracy needed from these ML models exceeds 99.9%. Decreasing our error rate from 4% to 0.02%, as we saw with our pilot project, clearly showed the value of synthetic data enabled quickly with Falcon".

- ***Steve Varga, Principal Scientist, P&G, AI/ML Systems***

By combining AI-Driven Quality Control with high-fidelity synthetic data generation, Duality Robotics provides a scalable framework for next-generation manufacturing assurance. Falcon bridges the gap between simulation and production, enabling continuous improvement of AI systems that enhance quality, efficiency, and reliability across the industrial lifecycle.

[1] F. Mejia et al., "Fine-tuning foundation models for off-road autonomy with Digital Twin Simulation," Synthetic Data for Artificial Intelligence and Machine Learning: Tools, Techniques, and Applications III, p. 3, May 2025. doi:10.1117/12.3053776

[2] F. Mejia, S. Shah, P. C. Young, and A. T. Brunk, "Fine-tuning yolov8 for accurate unmanned aerial vehicle detection with Digital Twin Simulation," SAE Technical Paper Series, vol. 1, Sep. 2025. doi:10.4271/2025-01-0473