

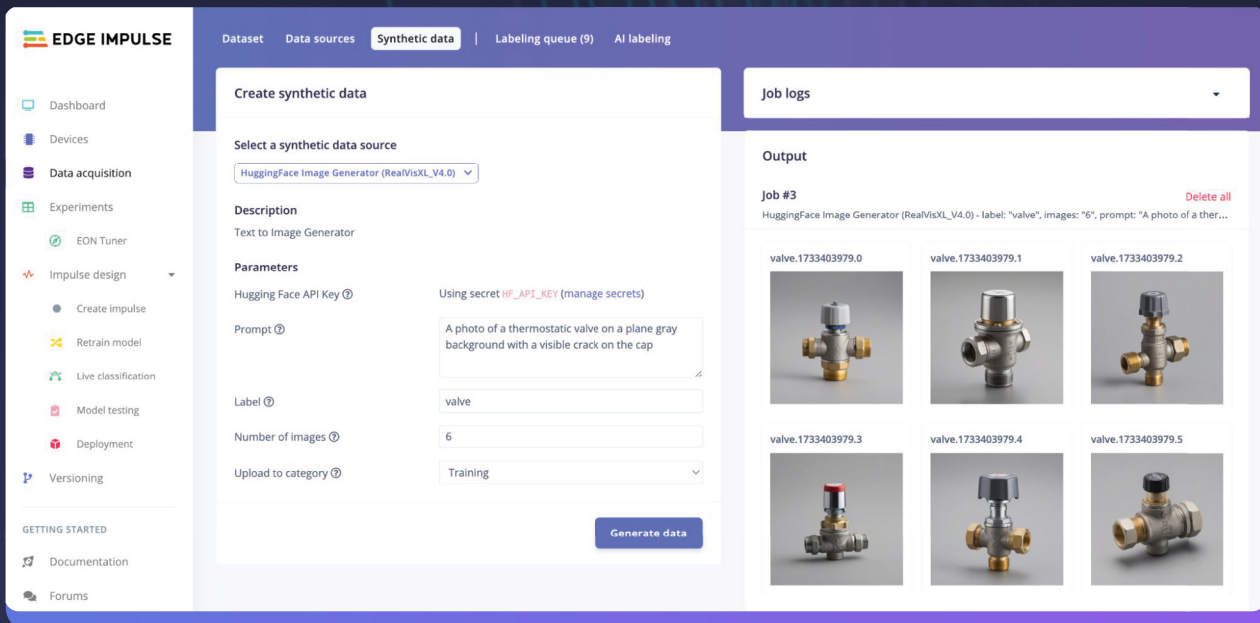
COMPUTER VISION AT THE EDGE

With the Edge Impulse platform, you can build, optimize, and deploy AI models on any edge system, from industrial gateways to ultra low-power cameras.

- Boost profit margins: Benefit from the lowest possible BOM costs by deploying AI to cost-efficient and ultra low-power cameras
- Iterate quickly and adapt to changing needs and environments
- Achieve long battery life (up to 10 years depending on use case)
- Accelerate and automate data collection, labeling, and validation without manual intervention

GenAI for Computer Vision at the Edge

Augment your datasets by leveraging the latest foundational models from OpenAI and Hugging Face to generate photorealistic images all from within the Edge Impulse platform.



The screenshot displays the Edge Impulse web interface. The top navigation bar includes 'Dataset', 'Data sources', 'Synthetic data', 'Labeling queue (9)', and 'AI labeling'. The left sidebar contains a navigation menu with options like 'Dashboard', 'Devices', 'Data acquisition', 'Experiments', 'EON Tuner', 'Impulse design', 'Create impulse', 'Retrain model', 'Live classification', 'Model testing', 'Deployment', 'Versioning', 'GETTING STARTED', 'Documentation', and 'Forums'. The main content area is titled 'Create synthetic data' and features a dropdown menu for 'Select a synthetic data source' set to 'HuggingFace Image Generator (RealVisXL_V4.0)'. Below this, the 'Description' is 'Text to Image Generator'. The 'Parameters' section includes a 'Hugging Face API Key' field with a note 'Using secret HF_API_KEY (manage secrets)', a 'Prompt' field containing 'A photo of a thermostatic valve on a plane gray background with a visible crack on the cap', a 'Label' field set to 'valve', a 'Number of images' field set to '6', and an 'Upload to category' dropdown set to 'Training'. A 'Generate data' button is at the bottom right. On the right side, the 'Job logs' section shows 'Job #3' with a 'Delete all' link and a list of six generated images, each with a unique ID (e.g., valve.1733403979.0) and a corresponding thumbnail image of a valve.



Sign up for a free Community Plan and get started today.

Reduce resource requirements with ultra low-power computer vision on edge devices

Latest Edge AI Vision Models

FOMO (Faster Objects, More Objects)

The only object detection model that can run on constrained devices, designed from the ground up for the edge. Able to achieve 30x better performance than MobileNet SSD and to run in <200K of RAM.

FOMO-AD

Visual anomaly detection models designed to scale and be optimized to run efficiently on any edge device, from constrained microcontrollers to powerful accelerators. Uses unsupervised learning techniques that don't require faulty data to train the model, enabling real-world applications where it's difficult to collect samples for every anomaly, especially for unanticipated defects.

YOLO-Pro

The first object detection model trained with extensive industrial datasets for optimal performance in real-world industrial scenarios. Optimized and available in various sizes to run efficiently on edge devices from CPUs to accelerators and GPUs.

Computer Vision on the Edge Use Cases

- Visual Inspection
- Anomaly Detection
- Predictive maintenance
- Asset Tracking
- Safety and Security
- Vehicle recognition
- OCR

We help companies scale and productize edge AI hardware, optimized for low power and cost



Sign up for a free Community Plan and get started today.

Unlock Edge AI

Experiments

Accelerate your project development by experimenting with different ML models simultaneously.

EON Tuner

Find and select the best computer vision model for your application within the constraints of your target device. The EON Tuner analyzes your data and all the possible neural network architectures - and gives you an overview of the optimal set of models that will fit your chosen device's latency and memory requirements.

Optimization and Deployment

Everything you develop will be able to be optimized and run on the largest variety of devices, thanks to our EON Compiler and our flexible deployment options.

Synthetic Data

Augment your dataset with images generated from photorealistic simulations, leveraging the power of NVIDIA Omniverse Replicator. Seamlessly integrated with Edge Impulse.

Edge AI Advantages

Faster time to market, with higher accuracy, and lower cost.

Real-Time: Instant feedback and insights

Privacy boost: On-device processing for extra security

Low latency: No cloud delays

Offline Use: Works anywhere, anytime

Battery life: Less cloud sync, more uptime, and longer life

High accuracy: Reliable and actionable insights