2020 embedded vision summit

Deep Learning for Manufacturing Inspection: Case Studies

Dr. Stephen Se
FLIR Systems, Inc.
September 2020

FLIR

Outline

- Manufacturing Inspection Applications
- Deep Learning Workflow
- Case Studies
- Best Practices

Manufacturing Inspection

FLIR

Deep Learning Revolution

- Machine Vision (MV) industry shifting from traditional rule-based approach to DL
  - Traditional methods require handcrafted algorithms by skilled engineers with programming & vision expertise
  - DL can adapt to new examples without re-programming
  - Tedious programming for identifying different defects no longer necessary

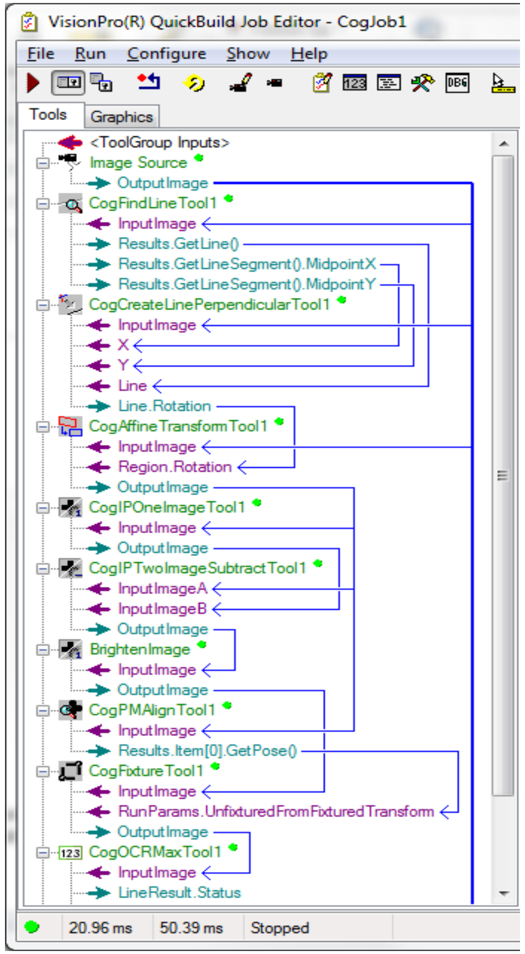

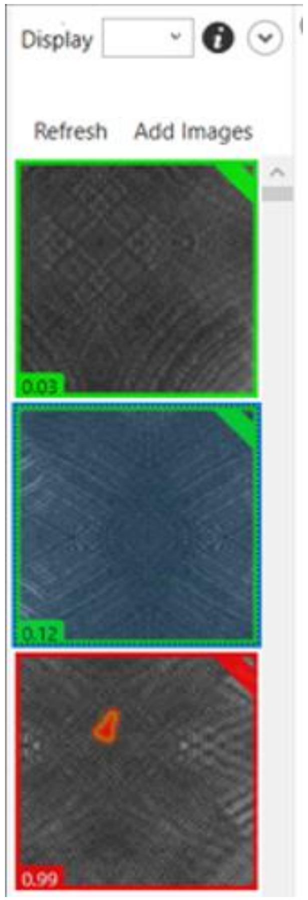

Source: Cognex

FLIR

Deep Learning MV Applications

- Manufacturing inspection, defect detection, classification, segmentation
- Textile inspection example
    - Traditional methods do not perform well after years of development
    - DL outperforms them after a few minutes of training

Textile inspection

Machined part inspection

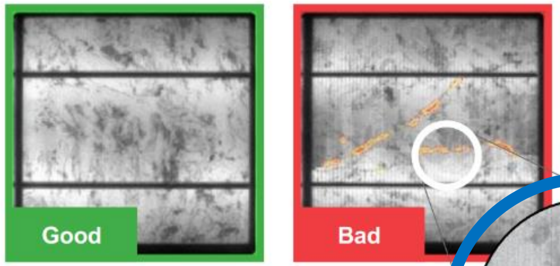

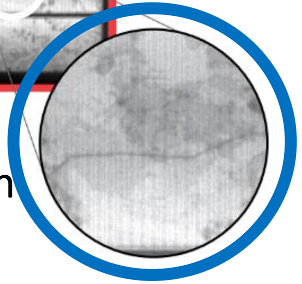

Solar panel inspection

Source: Cognex

Manufacturing Inspection Applications

- Well-controlled setup and lighting
  - Relatively few training examples
  - Training on standard laptop/desktop
- Commercial software
  - End-to-end solution: annotation GUI, training & inference tools
  - Cognex ViDi, MVTec HALCON, Adaptive Vision Studio
- Open source tools
  - Tensorflow, Caffe

Source: Cognex

Source: EPIC Vision Systems

Deep Learning vs Traditional Methods

	Deep Learning	Traditional Methods
Typical applications	• Surface inspection (cracks, scratches) • Food, plant, wood inspection • Textile inspection • Plastics, injection moulding	• Dimensional measurement • Code reading • Robot guidance • Precision alignment
Typical characteristics	• Deformable objects • Variable orientation • Vague specification with examples of good & bad parts	• Rigid objects • Fixed orientation • Formal specification with tolerances

Source: Adaptive Vision

Hybrid Method = Deep Learning + Traditional Method

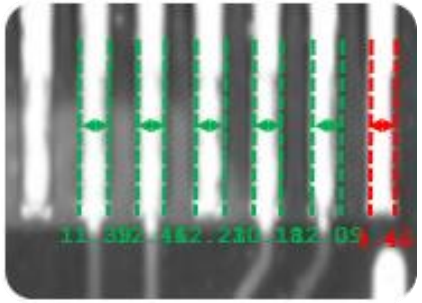

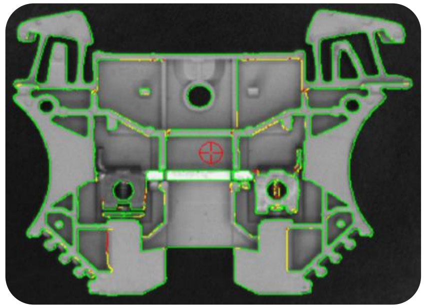

Source: Cognex

Deep Learning Workflow

Workflow and Open Source Tools

Training

Data collection & labeling → Data augmentation → Training (local / cloud) → Evaluate model performance

Caffe TensorFlow python

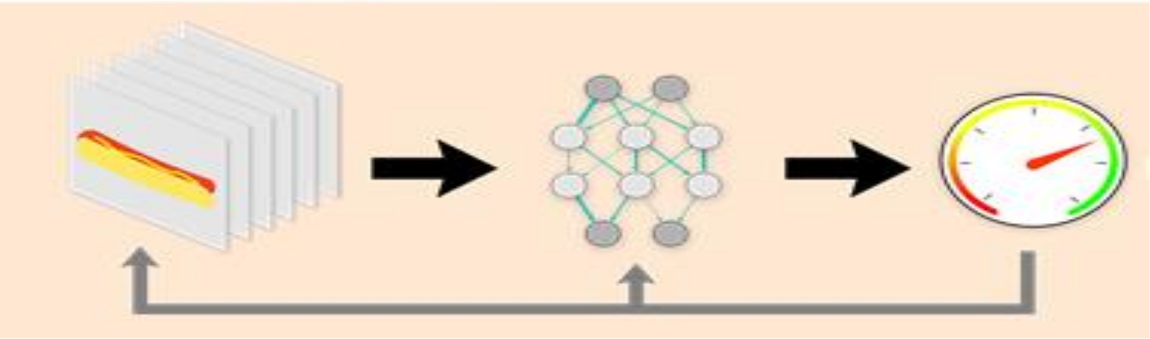

→ Convert model to Movidius graph format →

NCSDK OpenVINO

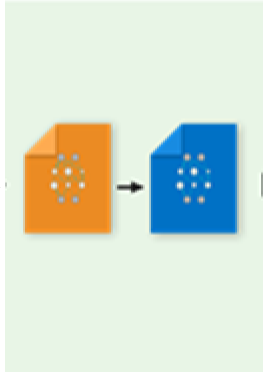

Deployment

Run inference → Output results (labels/boxes)

Intel Movidius Neural Compute Stick

Single-board computer

Data Collection & Labeling

- Data Collection
  - Capture images in similar setup as deployment such as lighting, camera, optics
  - Generate synthetic data – no labeling needed
- Accurate Data Labeling
  - For classification applications, divide training images into folders corresponding to the class
  - For detection applications, label the classes & bounding boxes using tools such as LabelImg

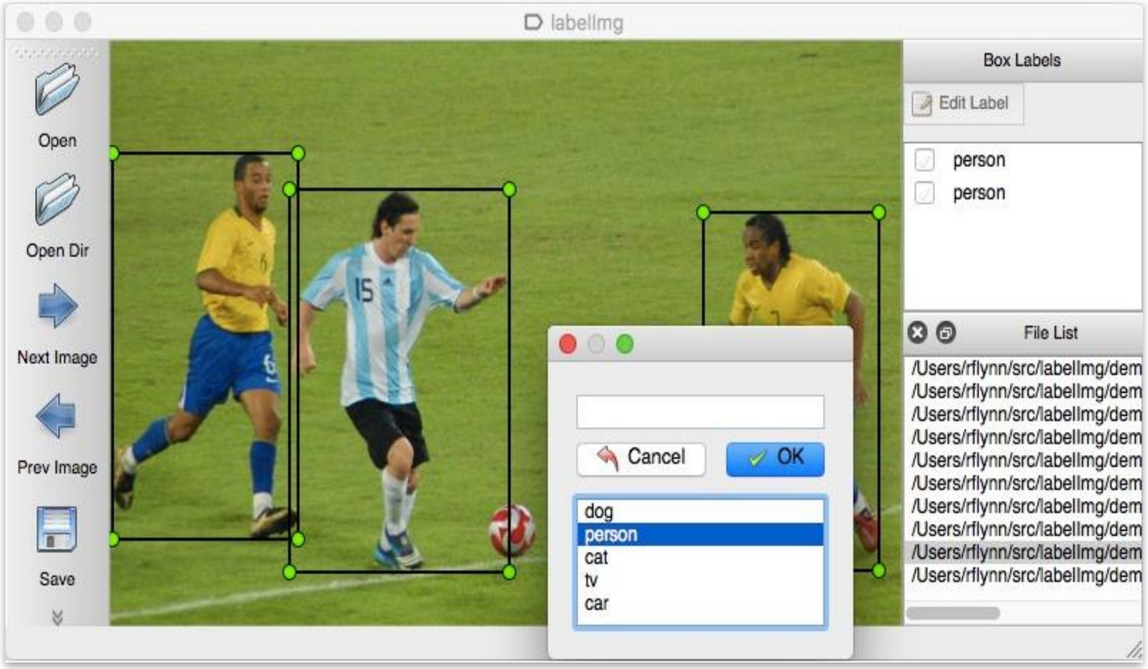

Screenshot from LabelImg
https://github.com/tzutalin/labelImg

Data Augmentation

- Expand and diversify training data by data augmentation
  - Perform image processing on collected images such as rotation, scaling, changing brightness, etc
  - Could improve inference accuracy without collecting more training images
  - Can be done before training, or on-the-fly during training

Augmentation examples

Transfer Learning

- Typical DL models have millions of parameters & take weeks to train
- Transfer learning is a technique to shortcut a lot of this work
  - Takes a fully-trained model and re-trains it for new tasks
  - Requires less training data, much faster to train

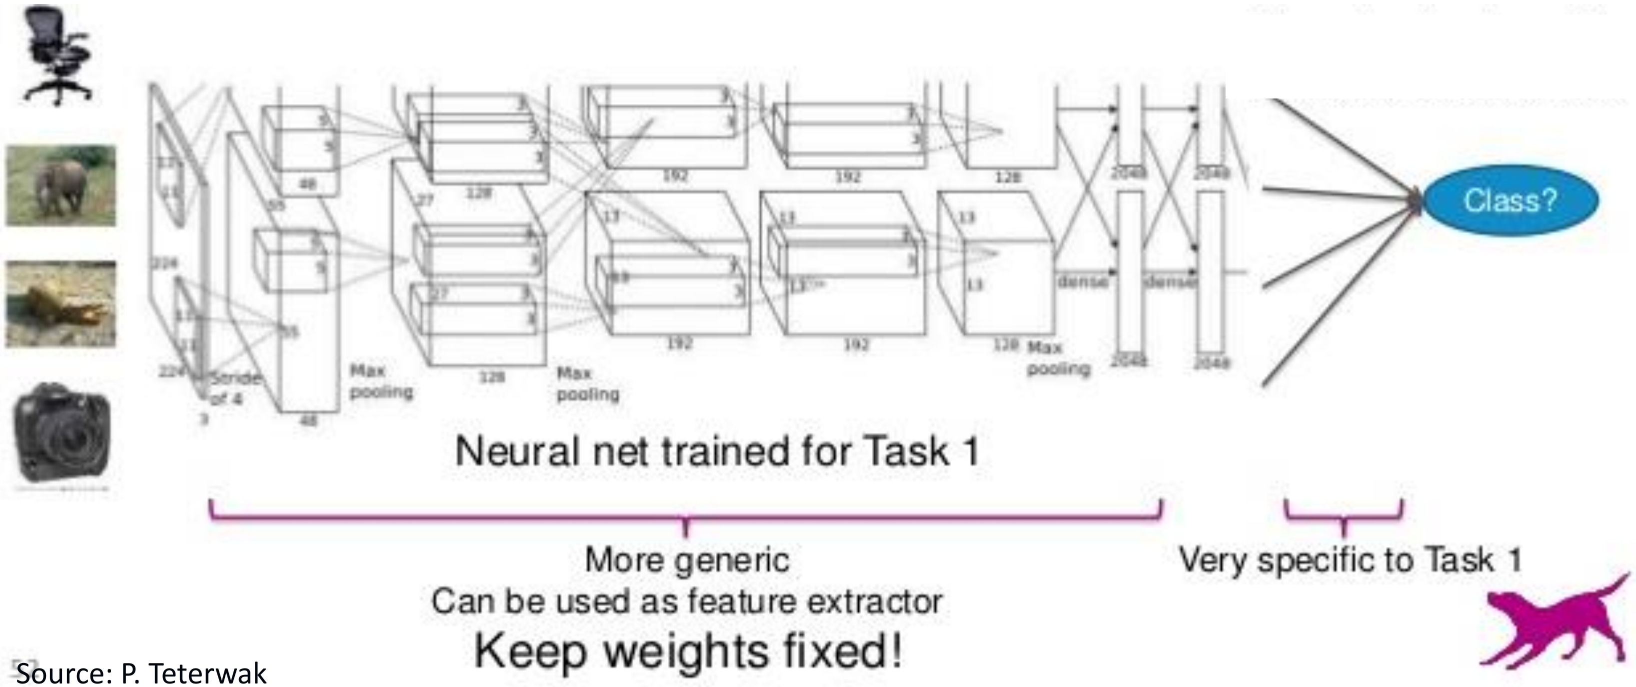

Source: P. Teterwak

Recent Trend

- Smaller CNN (Convolutional Neural Network) Models
  - More practical for mobile & embedded vision applications
  - Less parameters – trained with less training data
  - Less bandwidth to send a new model from cloud to edge
  - Achieve similar accuracy as larger models

Model	Million Mult-Adds	Million Parameters	ImageNet Accuracy
AlexNet	720	60	57.2%
GoogleNet	1550	6.8	69.8%
VGG 16	15300	138	71.5%
MobileNet V1 (1.0, 224)	569	4.24	70.6%
ShuffleNet	524	5	70.9%
MobileNet V2 (1.0, 224)	300	3.47	71.8%
NasNet (4@1056)	564	5.3	74.0%

Case Studies

Pills Inspection

- Classify between good and bad pills
  - Bad – scratch, dirt, chip or visible contamination
- Data collection in a controlled environment
  - Black background to increase visibility of pills
  - Select lens and camera height to capture fine details on pills
  - Good and bad images in separate folders

Good examples

Bad examples

Pills Inspection

- Data augmentation to generate more training images from the collected data
  - Using open-source Augmentor library https://github.com/mdbloice/Augmentor
  - Rotation, shearing, elastic distortion, perspective transform, chaining operations, …

Augmentation examples

Pills Inspection

- Data collection: 37 good examples & 54 bad examples
- Data augmentation: ~10000 training images for each class
- Transfer learning with MobileNet V1
  - Last layer (classification) changed from 1000 to 2 classes
  - Proof-of-concept: 94% accuracy
- Large difference in confidence for same pill at different location
  - Could benefit from more data and/or better augmentation

Camera Case Inspection

- Develop a deep learning solution for camera case inspection

good

bad

not ready

Imaging Challenges

Camera exposure variation

Diffuse light vs directional light

Camera Case Inspection

Class	# of training images
Good	61
Bad	167
Not Ready	127

- Re-trained a pruned MobileNet V1 network
  - Keep the first 7 layers to retain low level image features
  - Classification layer changed from 1000 to 3 classes
  - Proof-of-concept: 97.3% accuracy, 24 fps

PCB Hand Soldering Inspection

- Develop a deep learning solution for PCB hand-soldering inspection
  - Classify good solder vs missing/bad solder
  - Inspect multiple regions of interests

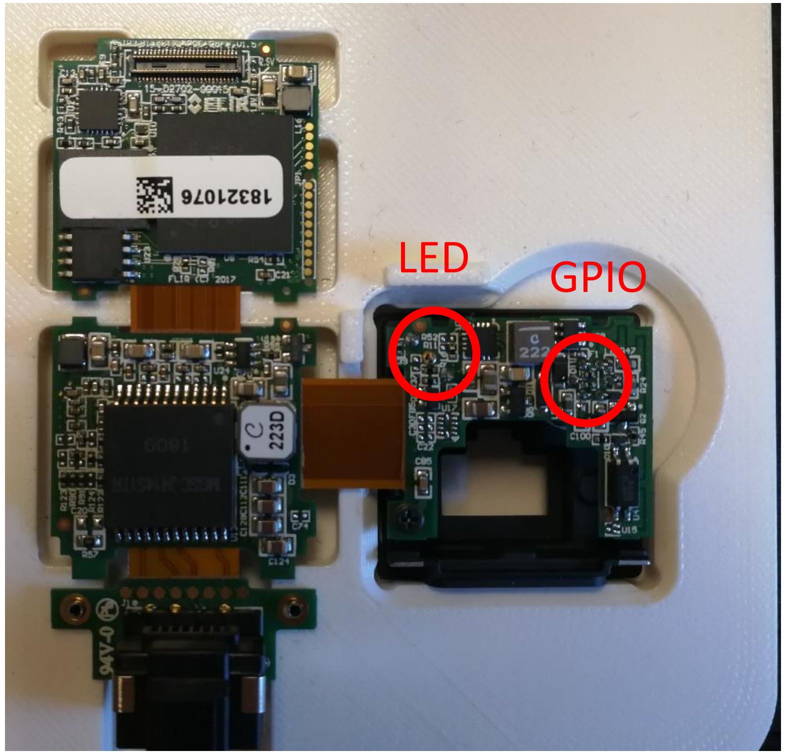

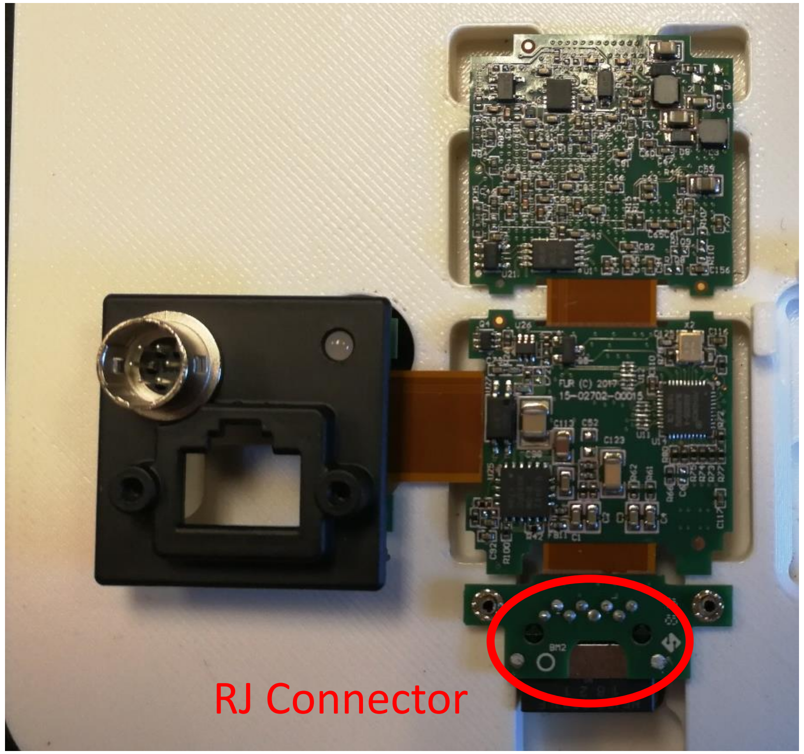

Common Soldering Problems

Solder bridge

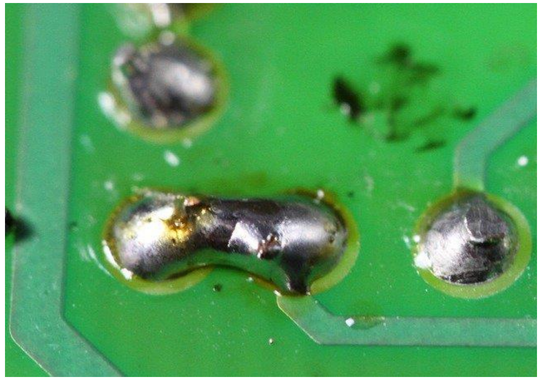

Stray solder spatters

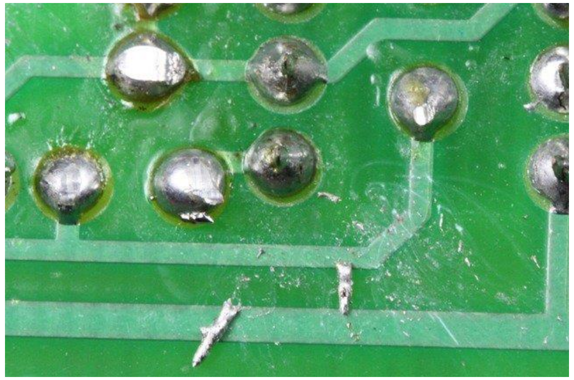

Lifted pad

Untrimmed leads

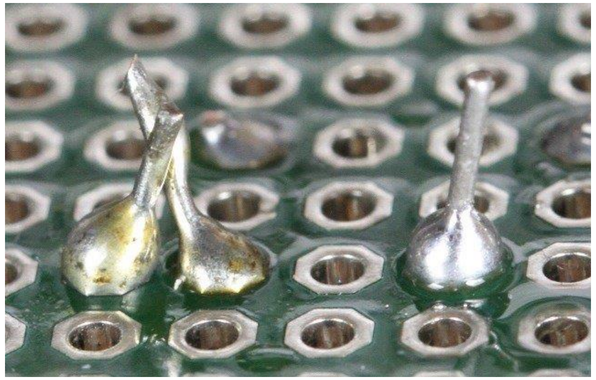

All of the above!

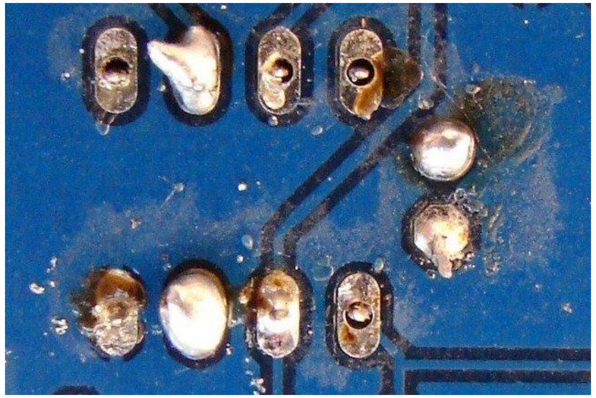

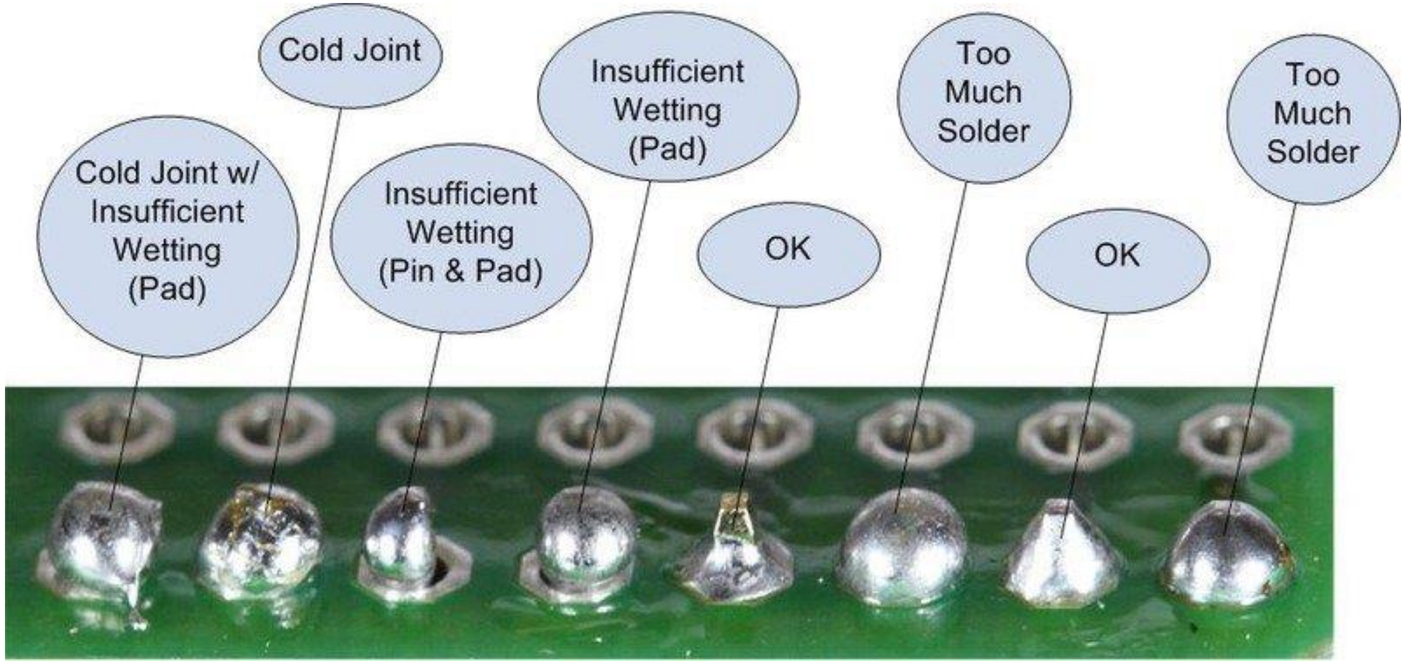

FLIR

PCB Hand Soldering Inspection

ROI	OK	NG		Total
		missing solder	all other defects	
LED	30	21	29	80
GPIO	31	21	28	80
RJ Connector	35	20	25	80

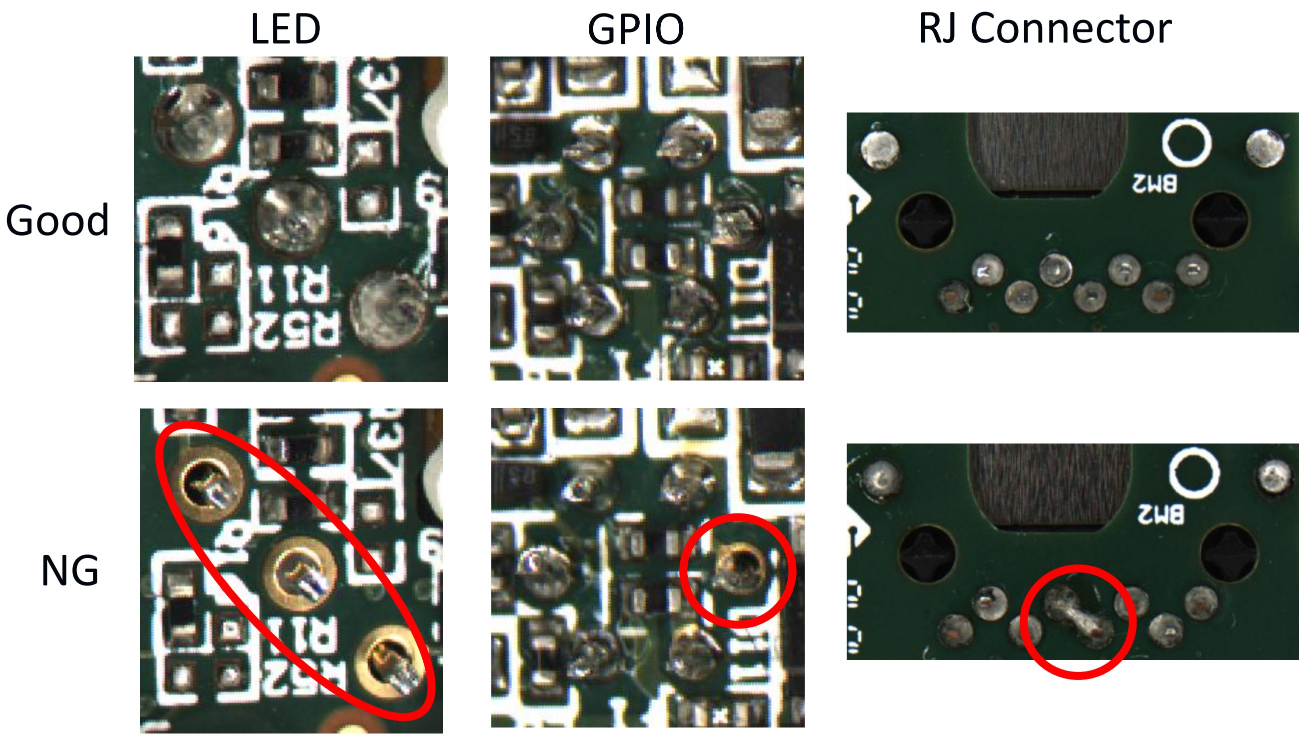

PCB Hand Soldering Inspection

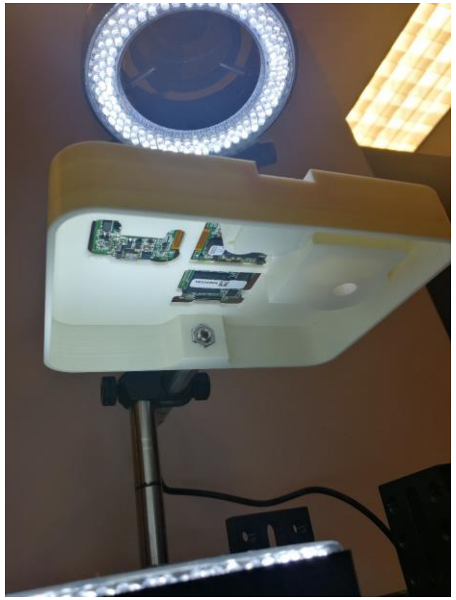

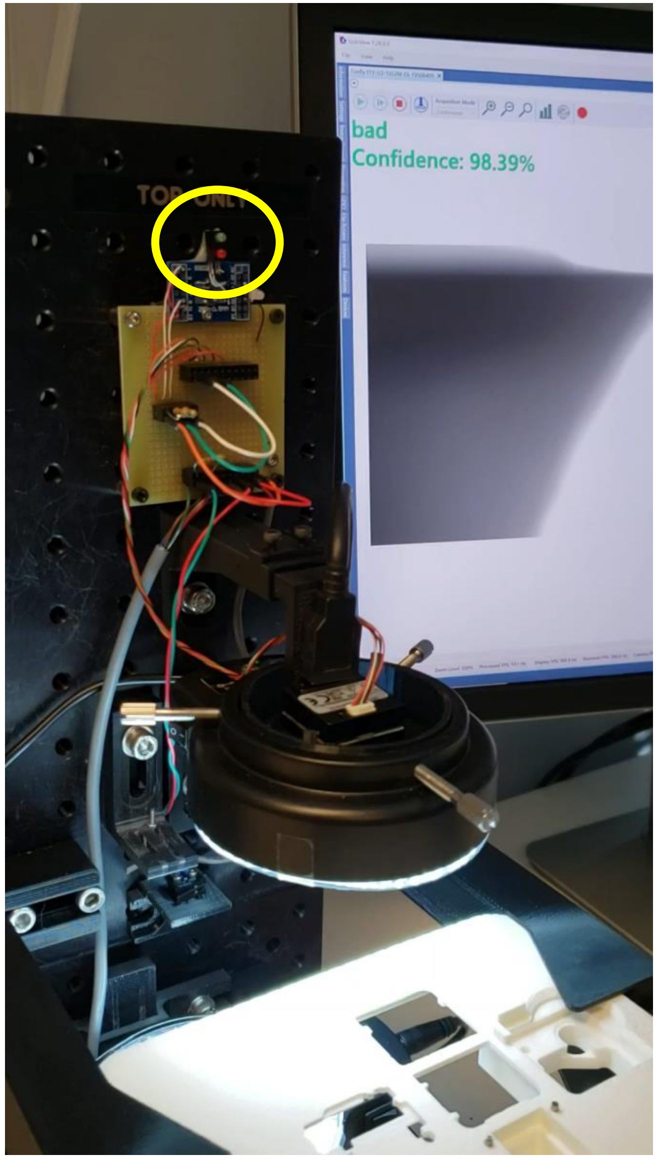

- Re-trained a MobileNet V1 network for each ROI
  - Last layer (classification) changed from 1000 to 2 classes
  - Proof-of-concept: 95%+ accuracy

Classification vs Anomaly Detection

- Classification approach (supervised)
  - Considerable training data for different classes
  - For inspection application, there are typically many good examples but very few bad examples
  - Class imbalance even after augmentation
- Anomaly detection approach (unsupervised)
  - Only good examples needed to train a model
  - Segment pixels/regions that differ from the model
  - More complex than classification approach

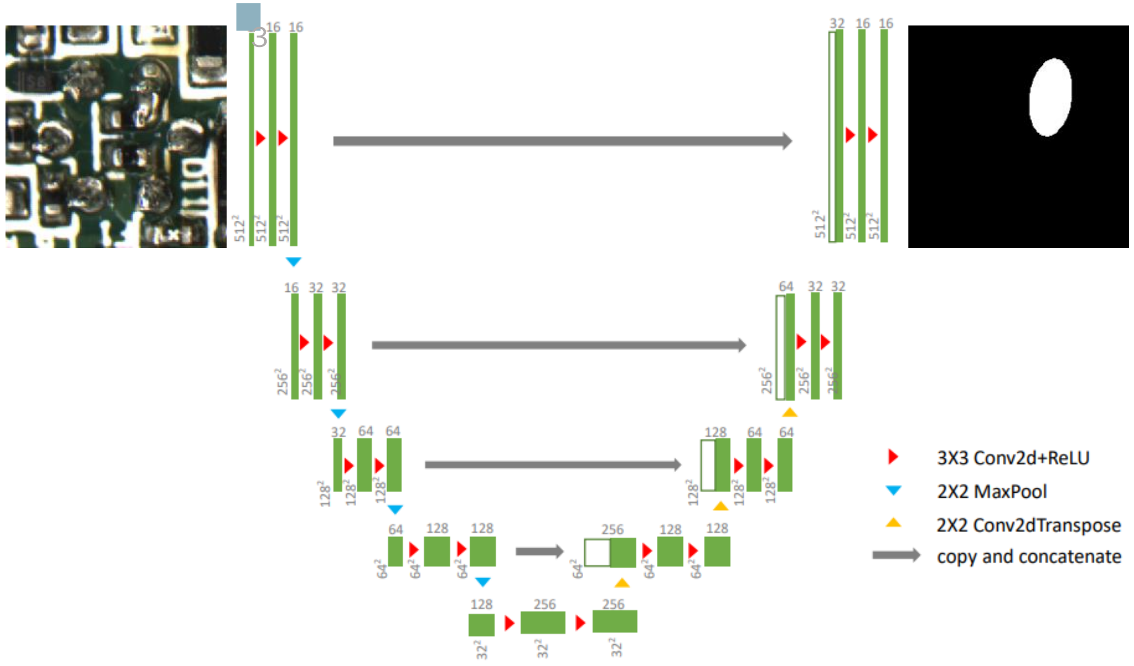

Anomaly detection using U-Net

Best Practices

- Training images must resemble inference images
  - Consistent object positioning, lighting, camera, optics
- Optimizing physical setup can simplify problem
  - Color-related problem could be solved using monochrome camera
  - Camera with resolution & dynamic range to capture fine details
  - Lighting could help highlight differences between classes
  - Smaller network requires less training data & can train faster
- Consistency in image processing is crucial

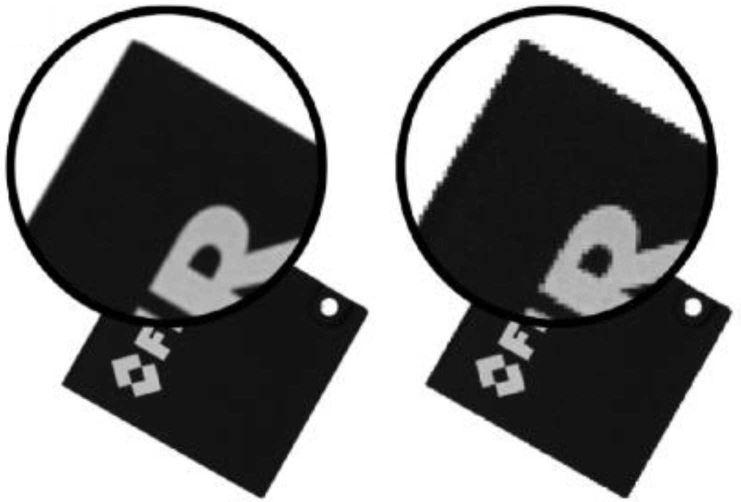

Resizing image using bilinear method with & without anti-aliasing

Lessons Learned

- DL is revolutionizing machine vision industry
  - DL vs traditional methods
- DL is highly suitable for manufacturing inspection
  - Well-controlled setup and lighting
  - Relatively few training images and short training time
- DL inference on the edge is feasible
  - Smaller networks are practical for embedded vision applications

Resources

- Manufacturing inspection software and use cases
    - Cognex VisionPro ViDi
    - MVTec HALCON
    - Adaptive Vision Studio
    - SUALAB SuaKIT
- Tutorial: “How to build a deep learning classification system for less than $600”
- Tutorial: “Tips for creating training data for deep learning neural networks”
- FLIR Firefly DL camera
    - Inference on the edge with integrated Intel Movidius VPU
    - https://www.flir.com/products/firefly-dl/

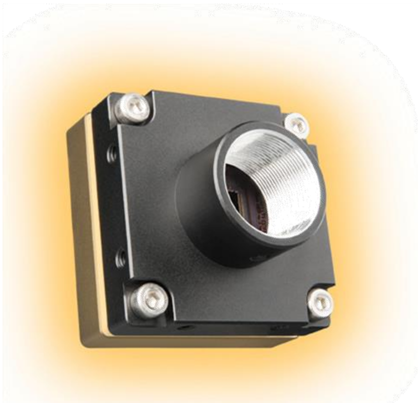

Questions?