

ML in ITK Visual Inspection

Zhan Li

2022/9/13



Automated Visual Inspection of ITk Sensors


Chris Gubbels

Supervisor: Oliver Stelzer-Chilton

UBC/TRIUMF

June 19, 2019

Visual Inspection & Goal of Automation



Task: Identify defects on the sensor

Current Method: Use a microscope and scan over the sensor by eye

Issues: Time consuming, tedious, likely inconsistent

Solution: Take pictures with an automated image stitching microscope and use machine learning/other algorithms to detect defects from that image

Methods of Automation

Goal: Flag defective regions

Approaches:

- Outlier Detection
 - Learn probability distribution of “normal” sensor over pixels
 - Pixels/regions of sensors below some probability cut are flagged
- Segmentation
 - Split the image into regions based on some criteria (e.g. colour)
 - Regions with a certain label or which fail to meet certain criteria are flagged


Defect Detection as Outlier Detection

Goal: Learn a probability distribution over pixels for all pixels


Solution 1: Reed-Xiaoli Detector (RXD)

- Each pixel in the image \mathbf{X} follows the same multivariate Gaussian
- Calculate sample mean $\hat{\mu}$ and covariance \hat{C} of colours (RGB)

$$\delta(\mathbf{x}) = (\mathbf{x} - \hat{\mu})^T \hat{C}^{-1} (\mathbf{x} - \hat{\mu})$$



Reed-Xiaoli (RX) Detector



- Misses most of scratch
- Misses most of thread
- Frequent false positives

Defect Detection as Outlier Detection

Goal: Learn a probability distribution over pixels for all pixels

Solution 1: Reed-Xiaoli Detector (RXD)


- Each pixel in the image \mathbf{X} follows the same multivariate Gaussian
- Calculate sample mean $\hat{\mu}$ and covariance \hat{C} of colours (RGB)

$$\delta(\mathbf{x}) = (\mathbf{x} - \hat{\mu})^T \hat{C}^{-1} (\mathbf{x} - \hat{\mu})$$

Solution 2: Laplacian Anomaly Detection (LAD)

- Pixels are now also correlated with their neighbors
- L replaces \hat{C}^{-1} , it contains both spatial and spectral weights

$$\delta(\mathbf{x}) = (\mathbf{x} - \hat{\mu})^T L (\mathbf{x} - \hat{\mu})$$




Understanding LAD

- Choose a graph structure and assign weights, W_{ij} , to edges
- Using the degree matrix and weight matrix, we can compute the Laplacian
$$L = D - W_{ij}$$
- The Laplacian used as a metric, gives a weighted square sum of differences between pixels which we may use as our discriminant.

$$\delta_{LAD}(\mathbf{x}) = \sum_{i,j} W_{ij} (x_i - x_j)^2 = \mathbf{x}^T L \mathbf{x}$$


Degree Matrix

- Degree matrix is a diagonal matrix.
- The element is the degree of different vertex. Degree means the number of edges connected to the vertex.



$$\Delta(\mathcal{G}) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 3 \end{pmatrix}$$

Laplacian Anomaly Detection





Defect Detection as Image Segmentation

Goal: Learn criteria for segment splitting and background segment selection

Solution 1: Colour-Distance Segmentation with segment pruning

- Each pixel begins as a segment
- Calculate Euclidean distance between mean RGB vectors of any two segments
- If the distance does not exceed the given threshold, merge the segments
- Segments under size threshold are merged with the most similar segment
- The largest segment or that closest to the mean image colour, is the background

Colour-Distance Segmentation




- Performs quite well, but still makes continuity errors and has slightly more false positives than LAD

Segmentation with Structure: U-Net

Problem: Simple segmentation algorithms can't learn complex enough patterns


Solution: U-Net

- Regular neural networks require large amounts of data
- U-Net's architecture allows it to maintain pixel-level features despite compressing the information with each layer
- Adding many transformed (reflected, rotated) images and z-standardizing images allows U-Net to learn from even <100 images



U-Net – More Info


- Left/down/encoding part of the “U” is like a normal CNN:
 - Convolutions encode translational equivariance
 - Max-pooling reduces dimension keeping local maxima
- Right/up/decoding part is upsampling to convert low-level representation into a pixel map
 - Upsampling part uses high level features from earlier
 - Interpolation through nearest-neighbor, bilinear, or transposed convolution (or other methods)
- U-Net relies on data transformations to extend information learned from each sample
- Important regions like regions near segment boundaries can be weighted heavier in loss (\sim error) function



- This is the simple form of U-Net, other forms are possible (e.g. feeding the features from the bottom of the “U” into another algorithm, as in [Dong, Taylor, Cootes, 2018])


Contracting Path

Expanding Path



Issues with Fading (Being fixed presently)

- Fading of smaller class leads to poor results. This should be reparable through alteration of the loss as in the Grayscale membrane example
- Training running right now



Summary

- Sensors require visual inspection. Doing this manually is slow, tedious, and inconsistent
- This visual inspection can be automated using defect detection algorithms
- Outlier detection methods make fewer assumptions and improve significantly with growing complexity. However, while LAD shows the best performance of all the models it still has frequent false positives and much more data is required to extend outlier detection further
- Colour-Based Segmentation works well; however, false positive rates are too high and some complex defects are not fully identified despite heavy optimization
- U-Net having issues with fading of smaller-segment class. Fix being implemented as we speak