



Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

Deep Learning-based Anomaly Detection on X-ray Images of Fuel Cell Electrodes

Jensen, Simon Buus; Moeslund, Thomas B.; Andreasen, Søren J.

Published in:

Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, Volume 4: VISAPP

DOI (link to publication from Publisher):

[10.5220/0010785400003124](https://doi.org/10.5220/0010785400003124)

Creative Commons License
CC BY-NC-ND 4.0

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Jensen, S. B., Moeslund, T. B., & Andreasen, S. J. (2022). Deep Learning-based Anomaly Detection on X-ray Images of Fuel Cell Electrodes. In *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, Volume 4: VISAPP* (Vol. 4, pp. 323-330). SCITEPRESS Digital Library. <https://doi.org/10.5220/0010785400003124>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Deep Learning-based Anomaly Detection on X-ray Images of Fuel Cell Electrodes

Simon B. Jensen¹^a, Thomas B. Moeslund¹^b and Søren J. Andreasen²

¹*Department of Architecture and Media Technology, Aalborg University, Aalborg, Denmark*

²*Serenergy, Aalborg, Denmark*

{sbje,tbm}@create.aau.dk, sja@serenergy.com

Keywords: Anomaly Detection, Deep Learning, Convolutional Neural Network, X-Ray, Data augmentation, Transfer Learning, Quality Control.

Abstract: Anomaly detection in X-ray images has been an active and lasting research area in the last decades, especially in the domain of medical X-ray images. For this work, we created a real-world labeled anomaly dataset, consisting of 16-bit X-ray image data of fuel cell electrodes coated with a platinum catalyst solution and perform anomaly detection on the dataset using a deep learning approach. The dataset contains a diverse set of anomalies with 11 identified common anomalies where the electrodes contain e.g., scratches, bubbles, smudges etc. We experiment with 16-bit image to 8-bit image conversion methods to utilize pre-trained Convolutional Neural Networks as feature extractors (transfer learning) and find that we achieve the best performance by maximizing the contrasts globally across the dataset during the 16-bit to 8-bit conversion, through histogram equalization. We group the fuel cell electrodes with anomalies into a single class called abnormal and the normal fuel cell electrodes into a class called normal, thereby abstracting the anomaly detection problem into a binary classification problem. We achieve a balanced accuracy of 85.18%. The anomaly detection is used by the company, Serenergy, for optimizing the time spend on the quality control of the fuel cell electrodes.

1 INTRODUCTION

Serenergy is a world-leading supplier of methanol-based fuel cell solutions with more than a thousand active units deployed globally. The fuel cells provide back-up power as well as temporary primary power or work in a hybrid system with renewable sources such as solar and/or wind. The core component of the fuel cell system, Serenergy, is a cell stack of 120 high temperature, polymer electrolyte membranes. Each cell contains 2 fuel cell electrodes that are coated with a platinum-based catalyst. Meaning the fuel cell system is made up of $120 \times 2 = 240$ fuel cell electrodes in total. Examples of fuel cell electrodes can be seen in figure 1. We will use the term fuel cell electrode or electrode interchangeably.

The quality of the platinum-based catalyst and the quality of how well it is coated on to each electrode is paramount to the overall conductivity of the fuel cell. The quality of an electrode is measured through

a semi-automatic/manual process where an X-ray image is manually captured of each electrode and the X-ray image is analyzed by an image analysis tool which outputs several quality parameters e.g., color histograms, box plots, standard deviation of the colors of the platinum coating, the minimum- and maximum colors of the platinum coating etc.

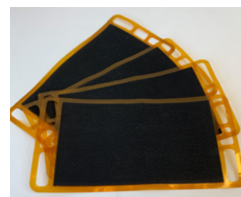




Figure 1: Fuel cell electrodes coated with a platinum catalyst

The process of analyzing the output of the image analysis tool is time-consuming and, in most cases, the electrodes have an acceptable quality. Serenergy wishes to optimize the quality control process by using a deep learning approach to perform automatic anomaly detection on the fuel cell electrodes,

^a <https://orcid.org/0000-0002-3217-1360>

^b <https://orcid.org/0000-0001-7584-5209>

by grouping them into two classes, normal or abnormal. Where a normal electrode can be used in the final fuel cell system and an abnormal electrode cannot.

In figure 2 examples of X-ray images of normal and abnormal electrodes are shown. This paper describes our approach to this important problem and the contributions are:

- 16-bit to 8-bit conversion methods for X-ray images of fuel cell electrodes using histogram equalization.
- A Deep Convolutional Neural Network classifier to perform anomaly detection of X-ray images of fuel cell electrodes.

1.1 Related Work

1.1.1 Deep Learning-based Classification on X-ray Image

Deep learning is a widely used technology for classification of images. More specifically Deep Convolutional Neural Networks (CNNs). The technological development of CNNs has been accelerated by image classification datasets and challenges such as pascal VOC (Everingham et al., 2010) and ImageNet (Russakovsky et al., 2014).

For X-ray images, there is a lack of large publicly available dataset, which means that few studies exist using CNNs for classification in X-ray images.

However, in the specific domain of pneumonia detection in chest X-ray images, a number of dataset has recently been made publicly available for example (Mooney, 2018), (Wang et al., 2017) and (Cohen et al., 2020). This has resulted in a large number of studies which use CNNs to classify X-ray images in this specific field. (Çallı et al., 2021) reviews and compares a large collection of these studies.

(Rahman et al., 2020) trains four well-recognized CNN models pre-trained on the ImageNet dataset, AlexNet (Krizhevsky et al., 2012), ResNet-18 (He et al., 2015), DenseNet201 (Wang and Zhang, 2020) and SqueezeNet (Iandola et al., 2016) for detecting pneumonia in the Pneumonia Chest X-ray dataset and compare their performance. Similarly, (Jiang, 2020) trains a CNN on the Pneumonia Chest X-ray dataset to detect pneumonia while utilizing the Dynamic Histogram Enhancement algorithm (Abdullah-Al-Wadud et al., 2007) as pre-processing method to improve the quality of X-ray images before training and evaluating the CNN model.

Other domains where deep learning-based classification has been applied to X-ray image data are e.g., classification of threat objects in X-ray security

imaging for baggage inspection (Akçay and Breckon, 2021) and classification of dental caries in bitewing X-ray images (Lee et al., 2021).

1.1.2 Transfer Learning for Image Classification

(Hussain et al., 2019) shows that transfer learning for image classification using deep CNNs is a valid and efficient method for achieving high performance in image classification tasks when dealing with datasets of limited size. They use the Inception-v3 (Szegedy et al., 2015) CNN pre-trained on the ImageNet dataset and re-train the model on the Caltech Face dataset consisting of only 450 images while achieving an accuracy of 65.7%.

(Rahman et al., 2020) validates transfer learning for image classification when the base dataset (ImageNet) consist of 3-channel RGB images and the target dataset consists of 1-channel grayscale X-Ray images. They do this by fine-tuning CNN models pre-trained with ImageNet on chest X-ray datasets.

Several transfer learning methods have recently been publicized which further optimizing the performance which can be achieved by the method. (Wang et al., 2019) proposes a method called attentive feature distillation and selection (AFDS), which adjusts the strength of transfer learning regularization and also dynamically determines the important features to transfer. They impose the method onto ResNet-101 and achieve state-of-the art computation reduction.

1.1.3 Anomaly Detection using Deep Learning

(Alloqmani et al., 2021) reviews twenty studies which utilizes deep learning for anomaly detection and identify challenges and insights in the domain. They identify the three main challenges for anomaly detection to be: (1) handling the class imbalance of normal and abnormal data, (2) the availability of labeled data and (3) the fact that there is often noise in the data that appears to be close to the actual anomalies and thus it becomes difficult to differentiate them.

(Ho and Wookey, 2020) proposes a solution to the class imbalance challenge, by introducing a new loss function for binary- and multiclass classification problems called Real-World-Weight Cross-Entropy loss function. Which allows direct input of real world costs as weights. This could prove useful for classification problems where there is a well-defined loss/cost for misclassified samples.

(Elgendi et al., 2021) proposes a solution to the second challenge by introducing and comparing four data augmentation methods for artificially increasing the number of training samples of X-ray images,

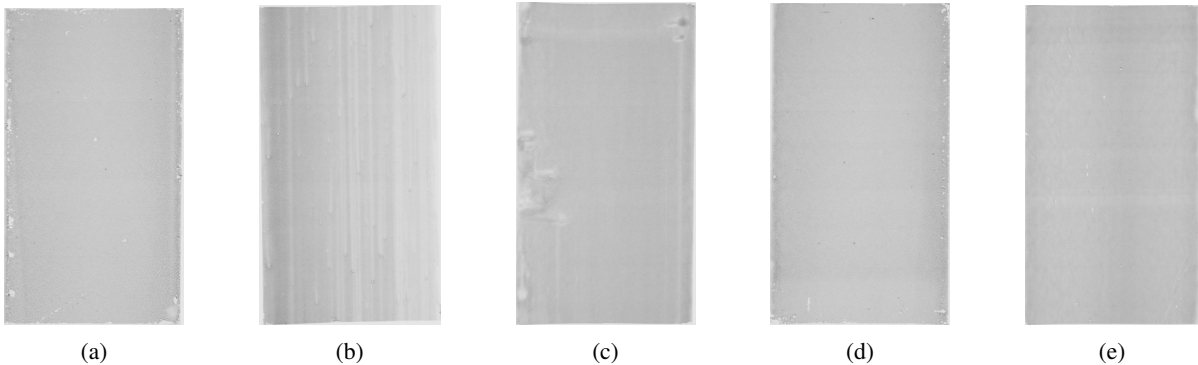


Figure 2: Examples of X-ray images of fuel cell electrodes. Figure 2a, 2b and 2c are examples of abnormal plates, while figure 2d and 2e are examples of normal plates. 2a is abnormal due to issues with scratches near its edges. 2b is abnormal due to issues with lines and bubbles and 2c is abnormal due to issues with smudges.

while performing Covid-19 pneumonia detection using a CNN. The methods use combinations of random rotations, shear, translation, horizontal- and vertical flipping among other data augmentation methods.

We discuss the three challenges and possible solutions further, in regards to anomaly detection in X-ray images of fuel cell electrodes in section 4.

2 APPROACH

2.1 Overview

An overview of the anomaly detection approach proposed in this paper is seen in figure 3. The approach is described by the following three steps:

1. Convert the X-ray images from 16-bit to 8-bit, as described in section 2.3.
2. Extract features using a pre-trained ResNet-34 CNN model as described in 2.4.
3. Classify anomalies in the feature-maps generated by the CNN model, using fully connected layers, as described in 2.4.1.

2.2 The Fuel Cell Electrode X-Ray Dataset

A real-world labeled anomaly dataset, consisting of 16-bit X-ray images of platinum catalyst coated fuel cell electrodes, is created for this work, as no existing dataset exists to the best of our knowledge, in this specific domain.

The fuel cell electrode X-ray dataset consists of 714 X-ray images. Each electrode belongs to 1 of 12 bundles, named bundle 1, bundle 2, ..., bundle 12. Where a bundle represents a collection of electrodes

which are coated with a platinum catalyst solution using the same coating method and mixture. The bundles in the dataset have varying number of X-ray images. The class balance varies for each bundle as well. The dataset is imbalanced with an over representation of normal samples. Across all 714 images 562 (78.71%) are labeled as normal and 152 (21.29%) are labeled as abnormal.

A dataset size of 714 X-ray images is considered to be a small dataset when utilizing a deep learning approach. We utilize transfer-learning to overcome this problem, using a pre-trained ResNet-34 model, as described in section 2.4. Due to the limited size of the dataset, creating a representative test set for the dataset proves difficult. We evaluated the X-ray images of each bundle through cross-validation. This is done by evaluating each bundle individually, while the remaining bundles are used as training samples and combining the results of each evaluation into a total score. This is further described in section 3.

2.2.1 Fuel Cell Electrode X-Ray Images

The X-ray images of the dataset have varying dimensions with a minimum width of 1119 pixels, maximum width of 1219 pixels and mean width of 1145 pixels. The minimum height of an electrode is 2053 pixels, the maximum height is 2115 pixels and the mean 2072,1 pixels. During training and evaluation of the anomaly detector, the X-ray images are transformed into a uniform size of 2000×1000 . Examples of normal and abnormal electrode X-ray images can be seen in figure 2.

2.2.2 Anomalies

Serenergy has identified 11 common anomaly types, which are grouped into a single class called abnormal. The identified 11 common anomaly types are named:

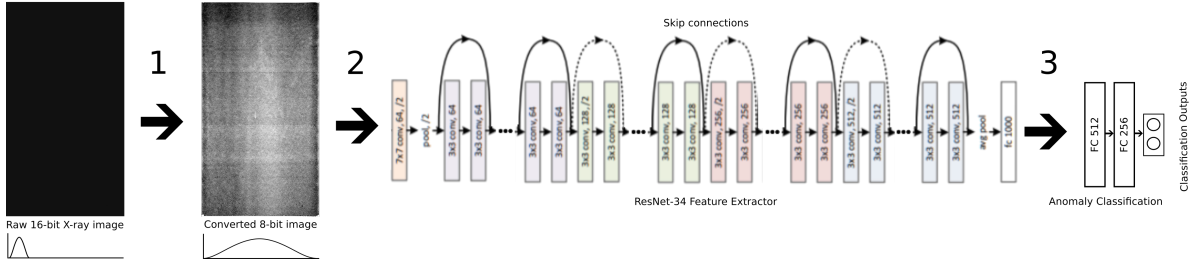


Figure 3: An overview of the anomaly detection approach proposed in this paper made up by three main steps which are described in section 2. Some of the layers of the ResNet-34 feature extractor are hidden for illustration purposes.

scratches, lines, edge cuts, edge tensions, smudges, edge ink flow, bubbles, missing ink, agglomerate, ink fluctuations, ink entry/exit.

The normal fuel cell electrodes can have minor representations of one or more anomaly type, as long as the severity is not too great. Figure 2d is an example of an electrode which have minor scratches, but the scratches are not sever enough to be classified as abnormal, while figure 2a is an example of a fuel cell electrode which is abnormal due to scratches near its edges.

2.3 X-ray Image Conversion Methods

To utilize pre-trained CNNs such as ResNet-34 as feature extractor in the anomaly detector, the depth of the electrode X-ray images is extended from 1 channel to 3 channels and the pixel values of the images are converted from a 16-bit values into 8-bit values. Thus, increasing the similarity of the electrode X-ray images to the images of ImageNet on which the CNN is pre-trained on. This is further described in section 2.4.

The 16-bit color range consists of 2 to the power of 16 (65536) colors and the 8-bit color range consists of 2 to the power of 8 (256) colors. A loss of information during the conversion is therefore inevitable.

We experiment with four methods for converting the X-ray images from 16- to 8-bits using histogram equalization and implemented using Python's OpenCV library (Itseez, 2015).

In section 2.3.1 a naive 16-bit to 8-bit conversion method is described which is used as baseline. In section 2.3.2 we use histogram equalization with global maximum and minimum bounds calculated across the entire dataset and in section 2.3.3 we use histogram equalization with local maximum and minimum bounds calculated for each individual X-ray image. Finally, we mix the methods in section 2.3.4. Examples of the resulting 8-bit images, for each conversion, are shown in figure 4.

2.3.1 Method 1: Naive Conversion

Method 1 is a naive conversion which is used as baseline. Method 1 simply scales each 16-bit pixel value and convert it to an unsigned 8-bit type.

An example of a resulting fuel cell electrode image after conversion can be seen in figure 4a.

Most pixels lie in the range 1700-2800 in the 16-bit X-ray images, as shown in section 2.3.2, which means method 1 will appear very dark. Pixel values in the 16-bit color range of 1700-2800 corresponds to pixel values of 6-11 in the 8-bit range using naive conversion.

Finally, the resulting 1-channel 8-bit image is extended with 2 additional channels, resulting in a 3-channel 8-bit image.

2.3.2 Method 2: Conversion by global min and max

For method 2 the global maximum- and minimum pixel value, G_{max} and G_{min} , of the 16-bit X-ray dataset is calculated and used as upper- and lower-bounds for histogram equalization, during 16-bit to 8-bit conversion.

The global maximum pixel value is found by calculating the 99.99th percentile of the pixel values in all X-ray images in the dataset and taking the maximum pixel value found and round it to nearest hundred. Similarly, we calculate G_{min} by finding the 0.01th percentile. G_{max} and G_{min} are found to be 28000 and 1700. Method 2 is given by equation 1 and illustrated in figure 5.

$$f_2(x_i) = \text{usign} \left(\frac{\max(\min(x_i, G_{max}), G_{min})}{G_{max} - G_{min}} \times 256 \right) \quad (1)$$

The reason for calculating the 99.99th and 0.01th percentiles of the pixel values is to avoid that noise in the X-ray images will affect the values of G_{max} and G_{min} .

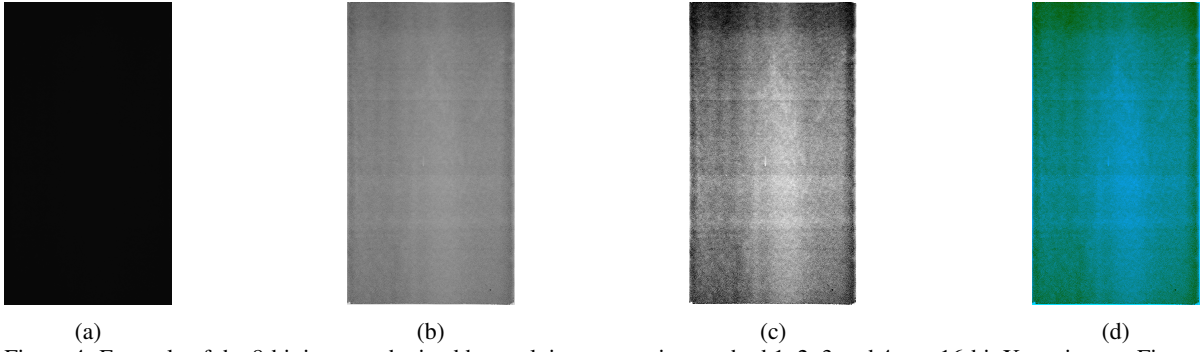


Figure 4: Example of the 8-bit images obtained by applying conversion method 1, 2, 3 and 4 to a 16-bit X-ray image. Figure 4a, 4b, 4c and 4d shows method 1, 2, 3 and 4 respectively.

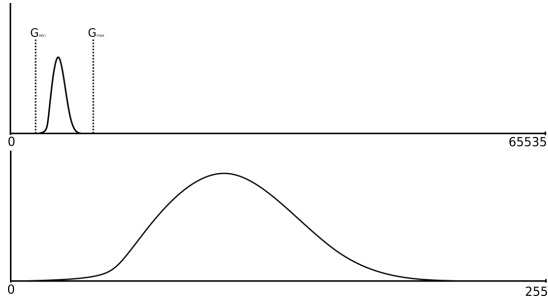


Figure 5: Method 2 uses the global minimum- and maximum pixel values, G_{max} and G_{min} , calculated across the entire fuel cell electrode X-ray dataset as upper- and lower bounds to convert from 16-bit to 8-bit through histogram equalization.

2.3.3 Method 3: Conversion by local min and max

Method 3 uses the local maximum- and minimum pixel value, L_{max} and L_{min} , found for each individual X-ray image, as upper- and lower-bound for histogram equalization.

Thereby, maximizing the contrast in the 8-bit color range for each individual X-ray image. Similarly to method 2, the 99.99th and 0.01th percentiles of the pixel values for each X-ray image are used to avoid that noise will affect the values of L_{max} and L_{min} .

The danger however, can be that the converted pixel values in one X-ray image lose their meaning relative to the converted pixel values of another X-ray image. Consider an X-ray image, image_{*i*} with a L_{min} value of 2200 and another X-ray image, image_{*j*} with an L_{min} value of 1800. After conversion, a value of 0 in the 8-bit color range will have corresponded to 2200 in image_{*i*} and to 1800 in image_{*j*}. Method 3 is given by equation 2.

$$f_3(k, x_i) = \text{usign} \left(\frac{\max(\min(x_i, L_{k,max}), L_{k,min})}{L_{k,max} - L_{k,min}} \times 256 \right) \quad (2)$$

Where $L_{k,max}$ and $L_{k,min}$ correspond to the local maximum and minimum of the *k*'th image in the electrode dataset.

2.3.4 Method 4: Mixing the methods

Finally, method 4 mixes the conversion methods from method 1, 2 and 3, such that the resulting image will contain the 8-bit pixel values from method 1 in its 1. channel, the pixel values from method 2 in its 2. channel and the pixel values from method 3 in its 3. channel as seen in figure 6.

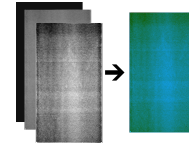


Figure 6: Method 4 mixes the conversion methods from method 1, 2, 3 into a 3-channel image, with 1 channel corresponding to the values achieved by each of the methods.

2.4 Anomaly Detector Architecture

This paper proposes a fuel cell electrode anomaly detector which uses a Convolutional Neural Network as feature extractor. For this purpose, a PyTorch (Paszke et al., 2019) implementation of the ResNet-34 (He et al., 2015) CNN model is utilized.

The ResNet-34 model is pre-trained on a base dataset (ImageNet (Russakovsky et al., 2014)). The reason for using a pre-trained CNN, also referred to as transfer learning, is that training CNNs from random initializations, usually requires a large amount of data, to achieve a high performance. For many real-world applications it can be both time-consuming and expensive to collect the required amount of data, as is also the case for this research.

The ResNet-34 model is extended with two fully connected layers of size 512 and 256 and finally with

a softmax layer of size 2, to get an output for each class, normal and abnormal. The model architecture is illustrated in figure 3.

2.4.1 Training

The anomaly detector is trained using a single Nvidia GeForce GTX 1080 GPU. We train an instance of the anomaly detector for each of the 4 conversion methods, which are each trained for 50 epochs. We choose 50 epochs to avoid overfitting the anomaly detector to the small training set. All bundles of fuel cell electrodes are used as training set except one, which is used for evaluation, as described in section 3. All layers in the ResNet-34 CNN are frozen, except for the last fully connected layer of size 1000 (as seen in figure 3), which is fine-tuned together with the fully connected layers of the anomaly classifier. We use the cross-entropy loss and an initial learning rate of 0.001 which decays every 10 epochs.

We augment the dataset with random horizontal and vertical flips (set to 50% probability) and resize the electrode images to height of 2000 pixels and width of 1000 pixels.

3 RESULTS

In this section we present the results achieved by the anomaly detector when applying the 4 conversion methods described in section 2.3 to the fuel cell electrode X-ray dataset. The anomaly detectors are evaluated using the balanced accuracy metric described in section 3.1.

3.1 Balanced accuracy

We use the balanced accuracy metric to evaluate the anomaly detector, as it proves useful for binary classification problems on datasets with class imbalance, as is the case for this project. Whereas, accuracy can be misleading if the class imbalance is great. The balanced accuracy metric overcomes this issue by weighting the positive and negative samples equally significant despite one class being more numerous than the other. This is done by adding the true positive rate (TPR) with the true negative rate (TNR) and dividing them by 2, as can be seen in equation 3.

$$\text{BALANCED ACC} = \frac{\text{TPR} + \text{TNR}}{2} \quad (3)$$

For a dataset with 98 positive samples and 2 negative samples, a classifier will achieve an accuracy score of 98% by simply classifying every sample as

positive. The balanced accuracy score will only be 50% in this case.

3.2 Cross-validation Evaluation

The dataset is evaluated through cross-validation, due to the limited size of the dataset and the nature of the dataset where each bundle of images is coated with a platinum catalyst solution using the same coating method and solution mixture. This means images from the same bundle will have a high similarity to one another. Utilizing samples from the same bundle as both training and testing samples will therefore inevitably occlude the performance of the anomaly detectors.

We train the anomaly detector as described in section 2.4.1 using all but one bundle, which is used as test set. The same procedure is replicated until each bundle has been evaluated individually for each of the 4 conversion methods. Meaning a total of 4×12 training/evaluations are performed. For each training/evaluation run, the balanced accuracy score is calculated. A combined balanced accuracy score is then calculated for each conversion method by adding the true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN) obtained when evaluating each of the 12 bundles, for the given conversion method. The results can be seen in table 1.

We find that the best anomaly detection performance is achieved by conversion method 2, with a balanced accuracy of 85.18%. Surprisingly, we find that method 4, which combines conversion method 1, 2 and 3 achieves the worst overall anomaly detection performance. A possible explanation for this might be that the dissimilarity between the combined 3-channel image created by method 4 and the features of RGB images in ImageNet, which our feature extractor CNN is pre-trained on, is too great.

4 DISCUSSION

Labeling anomalies in X-ray images of fuel cell electrodes is a difficult and time-consuming task, which requires expertise in the specific domain and knowledge about the severity and consequences different anomaly types impose to the conductivity of the fuel cell systems in which the electrodes will be used.

When labeling samples for a binary classification problem, the person labeling must decide on which class the sample belong to, for this work, whether an electrode is normal or abnormal.

We find that making this decision, for some samples, is a non-trivial task prone to subjectivity. While

Table 1: The balanced accuracy (%) achieved by the anomaly detector when applying conversion method 1, 2, 3 and 4 to the fuel cell electrode X-ray images. The performance achieved by each conversion method is measured through cross-validation, where each bundle is used as test set individually while the remaining bundles are used as training set. The green colors highlight the method(s) which achieved the best balanced accuracy score for each bundle. The olive-green color highlights the method which achieved the best overall balanced accuracy score when combining the evaluations of each bundle, which was conversion method 2.

Test Bundle	Method 1	Method 2	Method 3	Method 4
Bundle 1	50.00	96.88	50.00	40.63
Bundle 2	75.00	50.00	75.00	75.00
Bundle 3	50.00	50.00	50.00	50.00
Bundle 4	90.54	58.65	67.30	67.30
Bundle 5	75.00	62.50	75.00	62.50
Bundle 6	66.15	87.05	77.56	85.13
Bundle 7	82.39	89.53	73.89	72.17
Bundle 8	50.00	50.00	100.00	50.00
Bundle 9	68.53	50.00	90.00	70.00
Bundle 10	50.00	50.00	54.17	54.17
Bundle 11	75.00	99.40	75.00	75.00
Bundle 12	53.33	80.70	75.52	67.76
Overall	82.31	85.18	85.00	81.02

one expert might label the sample as normal another expert might label the same sample as abnormal. One solution, which was chosen for this work, is to let the more experienced expert make the final decision.

A second solution, could be to distribute the labeling task across many experts or non-experts and let the label with most votes represent the sample. The method has been described and evaluated by (Younesian et al., 2020) and it has the potential to remove subjectivity from the labels. The drawback to this solution is that it can be expensive and time-consuming and in some cases a sample might end up having equally many votes for each class in which case an additional solution for these cases need to be found.

Other solutions could be to simply exclude such samples from the dataset or to introduce a third class which represents samples which are undecidable, if one or more experts disagree. Such a class would in our case be small and cause a highly imbalanced dataset.

5 CONCLUSIONS

This paper proposed an anomaly detector using a Deep Convolutional Neural Network, an extended ResNet-34 model, for detecting anomalies in X-ray images of fuel cell electrodes. For this purpose, a dataset with normal and abnormal fuel cell electrodes X-ray images was created. The dataset consists of 12 bundles of images with a total of 714 X-ray images. The anomaly detector is used by the company Serenergy for automatizing a time-consuming manual quality control of the fuel cell electrode X-ray images. The anomaly detector was trained and evaluated

through cross-validation where a single bundle of images is used as test set and the remaining bundles are used as training set. The proposed anomaly detector was trained and evaluated using 12×50 epochs with a Nvidia GeForce 1080 GTX GPU and the PyTorch deep learning framework. We compared 16-bit conversion methods for pre-processing the X-ray images. We find that performing histogram equalization with upper- and lower bounds set by the maximum and minimum pixel values calculated across the entire dataset achieves a better performance than when using local maximum and minimum as upper- and lower bounds calculated for each individual image. We achieve a balanced accuracy of 85.18%.

In the future, we will continue to explore approaches for performing more accurate anomaly detection in X-ray images. Potential improvements could be achieved by using variations of weighted cross-entropy loss and data augmentation to cope with the imbalanced dataset and by utilizing different histogram equalization methods e.g., the DHE algorithm. Further, we see a potential in using CNNs pre-trained on large-scale gray-scale image datasets for classifying X-ray images. Whereas most CNNs today are pre-trained on RGB image datasets e.g., ImageNet.

ACKNOWLEDGEMENTS

We would like to thank Serenergy for their contributions, collaboration and dataset for this paper. We would also like to thank Ambolt Aps for initiating and facilitating the collaboration with Serenergy.

REFERENCES

- Abdullah-Al-Wadud, M., Kabir, M. H., Akber Dewan, M. A., and Chae, O. (2007). A dynamic histogram equalization for image contrast enhancement. *IEEE Transactions on Consumer Electronics*, 53(2):593–600.
- Akcay, S. and Breckon, T. (2021). Towards automatic threat detection: A survey of advances of deep learning within x-ray security imaging. *Pattern Recognition*, 122:108245.
- Alloqmani, A., B., Y., Khan, A., and Alsolami, F. (2021). Deep learning based anomaly detection in images: Insights, challenges and recommendations. *International Journal of Advanced Computer Science and Applications*, 12.
- Cohen, J. P., Morrison, P., Dao, L., Roth, K., Duong, T. Q., and Ghassemi, M. (2020). Covid-19 image data collection: Prospective predictions are the future. *arXiv 2006.11988*.
- Elgendi, M., Nasir, M. U., Tang, Q., Smith, D., Grenier, J.-P., Batte, C., Spieler, B., Leslie, W. D., Menon, C., Fletcher, R. R., Howard, N., Ward, R., Parker, W., and Nicolaou, S. (2021). The effectiveness of image augmentation in deep learning networks for detecting covid-19: A geometric transformation perspective. *Frontiers in Medicine*, 8:153.
- Everingham, M., Gool, L. V., Williams, C. K. I., Winn, J., and Zisserman, A. (2010). The pascal visual object classes (voc) challenge.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. *CoRR*, abs/1512.03385.
- Ho, Y. and Wookey, S. (2020). The real-world-weight cross-entropy loss function: Modeling the costs of mislabeling. *CoRR*, abs/2001.00570.
- Hussain, M., Bird, J. J., and Faria, D. R. (2019). A study on cnn transfer learning for image classification. In Lotfi, A., Bouchachia, H., Gegov, A., Langensiepen, C., and McGinnity, M., editors, *Advances in Computational Intelligence Systems*, pages 191–202, Cham. Springer International Publishing.
- Iandola, F. N., Moskewicz, M. W., Ashraf, K., Han, S., Dally, W. J., and Keutzer, K. (2016). Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size. *CoRR*, abs/1602.07360.
- Itseez (2015). Open source computer vision library.
- Jiang, Z. (2020). Chest x-ray pneumonia detection based on convolutional neural networks. In *2020 International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE)*, pages 341–344.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc.
- Lee, S., Oh, S., and Jo, J. (2021). Deep learning for early dental caries detection in bitewing radiographs. *Nature Scientific Reports*.
- Mooney, P. (2018). Chest x-ray images (pneumonia).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. *CoRR*, abs/1912.01703.
- Rahman, T., Chowdhury, M. E. H., Khandakar, A., Islam, K. R., Islam, K. F., Mahbub, Z. B., Kadir, M. A., and Kashem, S. (2020). Transfer learning with deep convolutional neural network (cnn) for pneumonia detection using chest x-ray. *Applied Sciences*, 10(9).
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M. S., Berg, A. C., and Li, F. (2014). Imagenet large scale visual recognition challenge. *CoRR*, abs/1409.0575.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2015). Rethinking the inception architecture for computer vision. *CoRR*, abs/1512.00567.
- Wang, K., Gao, X., Zhao, Y., Li, X., Dou, D., and Xu, C.-Z. (2019). Pay attention to features - transfer learn faster.
- Wang, S. and Zhang, Y.-D. (2020). Densenet-201-based deep neural network with composite learning factor and precomputation for multiple sclerosis classification. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 16:1–19.
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., and Summers, R. M. (2017). Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. *CoRR*, abs/1705.02315.
- Younesian, T., Hong, C., Ghiassi, A., Birke, R., and Chen, L. Y. (2020). End-to-end learning from noisy crowd to supervised machine learning models.
- Çallı, E., Sogancioglu, E., van Ginneken, B., van Leeuwen, K. G., and Murphy, K. (2021). Deep learning for chest x-ray analysis: A survey. *Medical Image Analysis*, 72:102125.